Thèse n° 10 228

Exploring policy change through agent-based simulation

Présentée le 19 mars 2021

Collège du management de la technologie
Chaire La Poste en management des industries de réseau
Programme doctoral en management de la technologie

pour l'obtention du grade de Docteur ès Sciences

par

Raphaël KLEIN

Acceptée sur proposition du jury

Prof. G. J. A. de Rassenfosse, président du jury
Prof. M. Finger, directeur de thèse
Prof. K. Ingold, rapporteuse
Dr I. Nikolic, rapporteur
Prof. T. Schmidt, rapporteur

École
polytechnique
fédérale
de Lausanne

2021

Acknowledgements

As I write this dissertation, completing three years of my life, I would like to thank some of the people that have helped me get to this point. The work would not have been what it is right now without their assistance and support.

I would like to first thank Matthias for all opportunities he gave me. Thanks to him, I learnt a lot both academically and practically. This was enabled by the different tasks that he made me responsible for, opening doors in a country I had never been to before. This helped greatly with the completion of this thesis but also my personal enrichment.

I would also like to thank my colleagues starting with Reinier who showed me the ropes at the beginning and later on Paul, who helped me with everything related to the Swiss electricity system. Just as important were my office colleagues who provided well-needed respite during moments of high stress.

Further away, I would like to thank all the people that helped me advance my research. This includes Amit who kept answering my calls long after I had left Delft. It also includes Igor, Jill, Karin, Paul and Chris, who helped to give me feedback on the papers I drafted, something I was more grateful than they can ever imagine.

On a more personal level, I would like to thank Thilini, who was always there regardless of the distance and the time. And finally, I would like to thank Arun for keeping me going even in the worst moments. This dissertation would not have been completed without his presence in my life.

Lausanne, January 8, 2021 R. Klein

Summary

Policymaking is a complex process that has been studied using policy process theories almost exclusively. These theories have been built using a large number of qualitative cases. Such methods are useful for theory building but remain limited for theory exploration and policy advice. On a different front, socio-technical system (STS) simulations are often used to test the impact and effectiveness of policies. This is done by using policy scenarios. This manner of dealing with deep uncertainty disregards the dynamic aspect of the policy process and of the way STS interact with policies.

The objective of this dissertation is to present an approach that can be used to systematically model and simulate the policy process. This dissertation demonstrates how such a model can be used in these two widely different applications to explore the policy process theories and to better account for deep uncertainty in STS simulations.

In the first part of the dissertation, I present a common language. It considers four building blocks using concepts from prominent theories as requirements to any model of the policy process: time, the policy arena, actors and the environment, and the actor interactions. Additionally, in this part, I argue why agent-based modelling is best suited to simulate the policy process. Finally, I present the hybrid modelling approach that is used to incorporate the policy process model into already existing STS models.

The second part focuses on the simplest implementation model. This model is the first use of the common language and is entitled 'simplest implementation' as the common language is used to build the simplest model possible capable of emulating the policy process. It is created to demonstrate the potential of the common language and is tested with three different STS. A predation model is used to present the possibilities stemming from this simplest implementation model. It is then added to a model of the Swiss electricity market to demonstrate what benefits the endogenisation of the policy process can have on the study of a complex STS. As a third example, the policy process model is added to a system dynamics model of flood safety in The Netherlands to demonstrate its versatility.

The third part of this dissertation consists of the advocacy coalition framework (ACF) implementation model. In this part, I present a model, also developed using the common language, that is aimed at emulating the ACF as closely as possible. The ACF being a complex framework, this implementation is done in steps of increasing complexity. The first step introduces policy learning to the model, the second coalitions, and the third and fourth steps aspects of partial knowledge and imperfect information. This ACF implementation model is also demonstrated using first a predation model followed by a model of the Swiss electricity market.

In conclusion, this dissertation shows that it is possible to model the policy process. However, it was shown that the policy process theories could benefit from additional and more focused studies so that behaviours and mechanisms can be better understood. This work can be the first step to such studies. It was also shown that introducing the policy process in a STS simulation can be useful to understand the said system better. This approach allows for easy integration to already existing models. Finally, this work can be extended by looking at other theories or different applications.

Résumé

L'élaboration des politiques est un processus complexe qui a été étudié en utilisant presque exclusivement les théories des processus politiques. Ces théories ont été construites à partir d'un grand nombre de cas qualitatifs. Ces méthodes sont utiles pour l'élaboration de théories mais restent limitées pour l'exploration de la théorie et pour donner des conseils politiques. Dans un autre domaine, les simulations de systèmes sociotechniques (SST) sont souvent utilisées pour tester l'impact et l'efficacité des politiques. Cela se fait à l'aide de scénarios de politiques. Cette manière de faire face à une incertitude profonde ne tient pas compte de l'aspect dynamique du processus politique et de la façon dont les SST interagissent avec les politiques.

L'objectif de cette thèse est de présenter une approche qui peut être utilisée pour modéliser et simuler systématiquement le processus politique. Cette thèse montre comment un tel modèle peut être utilisé dans ces deux applications très différentes pour explorer les théories du processus politique et pour mieux tenir compte de l'incertitude profonde dans les simulations SST.

Dans la première partie de la thèse, je présente un langage commun. Il considère quatre éléments constitutifs, utilisant des concepts de théories importantes comme exigences pour tout modèle de processus politique : le temps, l'arène politique, les acteurs et l'environnement, et les interactions des acteurs. De plus, dans cette partie, j'explique pourquoi la modélisation basée sur les agents est la mieux adaptée pour simuler le processus politique. Enfin, je présente l'approche de modélisation hybride utilisée pour incorporer le modèle de processus politique dans les modèles SST déjà existants.

La deuxième partie se concentre sur le modèle d'implémentation le plus simple. Ce modèle est la première utilisation du langage commun et est intitulé mise en œuvre la plus simple car le langage commun y est utilisé pour construire le modèle le plus simple possible capable d'émuler le processus politique. Il est créé pour démontrer le potentiel du langage commun et est testé avec trois SST différents. Un modèle de prédation est utilisé pour présenter les possibilités issues de ce modèle d'implémentation le plus simple. Il est ensuite ajouté à un modèle du marché suisse de l'électricité pour démontrer les avantages que l'intégration du processus politique peut avoir sur l'étude d'un SST complexe. Comme troisième exemple, le modèle de processus

politique est ajouté à un modèle de dynamique des systèmes de sécurité contre les inondations aux Pays-Bas pour démontrer sa polyvalence.
La troisième partie de cette thèse aborde le modèle de mise en œuvre du cadre de coalition de plaidoyer (ACF). Dans cette partie, je présente un modèle, également développé en utilisant le langage commun, qui vise à émuler l'ACF le plus fidèlement possible. L'ACF étant un cadre complexe, cette implémentation se fait par étapes de complexité croissante. La première étape introduit l'apprentissage des politiques sur le modèle, la seconde présente les coalitions, les troisième et quatrième étapes présentent les aspects de la connaissance partielle et de l'information imparfaite. Ce modèle de mise en œuvre d'ACF est également démontré en utilisant d'abord un modèle de prédation et ensuite un modèle du marché suisse de l'électricité.
En conclusion, cette thèse montre qu'il est possible de modéliser le processus politique. Cependant, il a été démontré que les théories du processus politique pourraient bénéficier d'études complémentaires et plus ciblées afin de mieux comprendre les comportements et les mécanismes. Ce travail peut être la première étape de telles études. Il a également été démontré que l'introduction du processus politique dans une simulation SST peut être utile pour mieux comprendre ledit système. Cette approche permet une intégration facile aux modèles déjà existants. Enfin, ce travail peut être étendu en examinant d'autres théories ou différentes applications.

Contents

Acknowledgements **i**

Summary **iii**

Résumé **v**

List of Figures **xiii**

List of Tables **xxi**

Glossary **xxiii**

I The preliminaries **1**

1 Introduction **3**
- 1.1 The state of the literature . . . 3
- 1.2 The limitations of current approaches . . . 4
  - 1.2.1 The policy process . . . 5
  - 1.2.2 Modelling and simulation . . . 5
- 1.3 The research objectives . . . 7
- 1.4 Research design . . . 8
  - 1.4.1 Addressing the limitations . . . 8
  - 1.4.2 The methodological approach . . . 9
  - 1.4.3 Operationalising the theories . . . 10
  - 1.4.4 The model selection . . . 11
- 1.5 The research process . . . 14
- 1.6 The dissertation structure . . . 16

2 Literature review **19**
- 2.1 Modelling and simulation . . . 19
- 2.2 The modelling paradigm . . . 21
- 2.3 Simulation in political sciences . . . 22

2.4 Policy process theories . . . 26
2.4.1 The policy cycle . . . 27
2.4.2 The multiple streams framework . . . 28
2.4.3 The advocacy coalition framework . . . 29

3 The methodological approach 31
3.1 The common language . . . 31
3.1.1 Time . . . 32
3.1.2 The policy arena . . . 33
3.1.3 The actors and their environment . . . 34
3.1.4 Actor interactions . . . 40
3.2 Hybrid simulation . . . 42
3.2.1 Hybrid simulation in the literature . . . 42
3.2.2 Endogenising the policy process . . . 44
3.3 Discussion . . . 45
3.3.1 The policy steps . . . 45
3.3.2 The belief system . . . 46
3.3.3 The policy entrepreneurs . . . 47
3.3.4 The streams . . . 47
3.3.5 The coalitions . . . 47
3.3.6 The common language potential . . . 48
3.3.7 Model complexity . . . 49
3.3.8 Data requirements . . . 50
3.3.9 Model validation . . . 50
3.4 Conclusions . . . 51

II The simplest implementation model 53

4 Conceptualisation 55
4.1 Time . . . 55
4.2 The policy arena . . . 56
4.3 The actors . . . 56
4.4 The interactions . . . 56
4.5 Discussion . . . 57
4.6 Conclusions . . . 58

5 Formalisation 59
5.1 The process . . . 59
5.2 The agents . . . 60
5.2.1 Active agents . . . 60

5.2.2 Passive agents . . . 62
5.2.3 Belief hierarchy . . . 62
5.3 The actions and interactions . . . 63
5.3.1 Preference calculation - secondary issues . . . 64
5.3.2 Policy instruments and preference calculation . . . 64
5.3.3 Policy instrument selection and implementation . . . 65
5.3.4 Electorate passive action on policy makers . . . 65
5.4 Conclusions . . . 66

6 The predation model **67**
6.1 The model . . . 67
6.2 The hybrid integration . . . 68
6.2.1 The belief system . . . 68
6.2.2 The policy instruments . . . 69
6.3 The research questions . . . 70
6.4 The scenarios . . . 71
6.5 The simulations . . . 73
6.6 The results . . . 73
6.6.1 With and without policy emergence . . . 74
6.6.2 Changing policy core preferred states . . . 75
6.6.3 Changing secondary preferred states . . . 77
6.6.4 Changing understanding of the world . . . 79
6.6.5 The impact of the electorate . . . 81
6.7 Discussion . . . 83
6.7.1 Hybrid integration . . . 83
6.7.2 The policy process simulation potential . . . 84
6.7.3 The agent distribution . . . 86
6.7.4 The inherent stability of the policy process model . . . 87
6.7.5 The causal beliefs and policy instruments assessment . . . 88
6.8 Conclusions . . . 89

7 The electricity market model **91**
7.1 The model . . . 91
7.2 The hybrid integration . . . 92
7.2.1 The belief system . . . 92
7.2.2 The policy instruments . . . 94
7.3 The research question . . . 95
7.4 The scenarios . . . 95
7.5 The simulations . . . 96
7.6 The results . . . 97

7.7 Discussion . . . 99
7.8 Conclusions . . . 103

8 The flood safety model **105**
8.1 The model . . . 105
8.2 The hybrid integration . . . 107
8.2.1 The belief system . . . 108
8.2.2 The policy packages . . . 108
8.2.3 Practical integration . . . 110
8.3 The research questions . . . 110
8.4 The scenarios . . . 111
8.5 The simulations . . . 112
8.6 Results . . . 112
8.7 Discussion . . . 115
8.8 Conclusions . . . 115

III The ACF implementation model **117**

9 Conceptualisation **119**
9.1 The simplest implementation model . . . 120
9.2 Emulating policy learning . . . 120
9.3 Adding the coalitions . . . 124
9.4 Including partial knowledge . . . 127
9.5 Considering imperfect information . . . 128
9.6 Discussion . . . 130
9.7 Conclusions . . . 131

10 Formalisation **133**
10.1 Policy learning extension (+PL) . . . 133
10.1.1 The process . . . 133
10.1.2 The agents . . . 135
10.1.3 The agent interactions . . . 137
10.2 Coalition extension (+Co) . . . 138
10.2.1 The process . . . 139
10.2.2 The agents . . . 140
10.2.3 The agent interactions . . . 140
10.3 Partial knowledge add-on (+PK) . . . 143
10.3.1 The agent interactions . . . 143
10.4 Imperfect information add-on (+PI) . . . 144
10.4.1 The process . . . 145

10.4.2 The agents . . . 145
10.4.3 The agent interactions . . . 147
10.5 Discussion . . . 149
10.5.1 The +PL model . . . 150
10.5.2 The +Co model . . . 151
10.5.3 The +PK add-on . . . 153
10.5.4 The +PI add-on . . . 154
10.6 Conclusion . . . 155

11 The predation model **157**
11.1 Emulating policy learning (+PL) . . . 158
11.1.1 The research question . . . 158
11.1.2 The scenarios . . . 158
11.1.3 The results . . . 161
11.2 Adding coalitions (+Co) . . . 169
11.2.1 The research question . . . 169
11.2.2 The scenarios . . . 170
11.2.3 The results . . . 171
11.3 Including partial knowledge (+PK) . . . 176
11.3.1 The research question . . . 176
11.3.2 The scenarios . . . 176
11.3.3 The results . . . 177
11.4 Considering imperfect information (+PI) . . . 179
11.4.1 The research question . . . 179
11.4.2 The scenarios . . . 179
11.4.3 The results . . . 182
11.5 Discussion . . . 188
11.5.1 The +PL model . . . 188
11.5.2 The +Co model . . . 189
11.5.3 The +PK model . . . 190
11.5.4 The +PI model . . . 191
11.5.5 General remarks . . . 193
11.6 Conclusion . . . 198

12 The electricity market model **199**
12.1 The research question . . . 199
12.2 The scenarios . . . 200
12.3 The results . . . 204
12.4 Discussion . . . 210
12.5 Conclusions . . . 215

IV Conclusions **217**

13 Discussions **219**

14 Conclusions **231**

15 Further research **235**

A The Swiss electricity market model **239**

A.1 Purpose of the model 239
A.2 Entities, state variables and scales 239
A.3 Process overview and scheduling 240
A.4 Design concepts 241
A.5 Initialisation 242
A.6 Input data 243
A.7 Submodels 243

B The Swiss electricity market model hybridisation equations **247**

C The simplest implementation predation model results **251**

D The simplest implementation electricity model results **261**

E The ACF implementation predation model results **267**

F The ACF implementation electricity model results **303**

Curriculum Vitae **337**

List of Figures

1.1 The research process. 15

3.1 Architecture of the hybrid model. 44

5.1 Steps followed by the policy process for the simplest implementation model. 61

5.2 Belief hierarchy architecture. The number of issues on each layer is provided as an example. 63

6.1 Comparison of the population of sheep and wolves between a simulation of the predation model alone and the benchmark scenario (scenario 0) incorporating the policy emergence model. 74

6.2 Predation model populations for the benchmark scenario and scenario 1. 76

6.3 Policy instrument selection for the benchmark scenario and scenario 1. 76

6.4 Predation model populations for the benchmark scenario and scenario 2. 78

6.5 Policy instrument selection for the benchmark scenario and scenario 2. 78

6.6 Predation model populations for the benchmark scenario and scenario 3. 80

6.7 Policy instrument selection for the benchmark scenario and scenario 3. 80

6.8 Predation model populations for the benchmark scenario, scenario 2 and scenario 4 (50% influence only). 81

6.9 Predation model populations for the benchmark scenario, scenario 2 and scenario 5 (50% influence only). 82

6.10 Policy instrument selection for scenario 1, 2, and scenarios 4 and 5 (20% and 50% influence only). 83

7.1 Monthly averaged electricity prices, split according to the three electorate influence rates and for the three demand growth scenarios. . . . 98

7.2 Overall demand with the electricity supplied by solar, wind and CCGT sources for all scenarios. 99

7.3 Cumulative policy instrument selection for a 3% electricity demand growth scenario for all fifty repetitions for each simulation. 100

8.1 System diagram of the system dynamics model where the non-policy dependent exogenous factors are presented in green, the endogenous factors in grey and the policy dependent exogenous factors are in magenta (adapted from (Klein et al., 2016)). 106
8.2 Causal loop diagram representing the system dynamics model conceptualisation. (1) is the maintenance repair loop, (2) and (3) are sufficient safety loops acting in opposite directions and (4) is the broken levee loop. 107
8.3 Outputs from the flood safety hybrid model showing the levee length, the safety parameters and the policy choices for both flood scenarios. . 113

10.1 Steps followed by the policy process for the ACF model with the first level of complexity +PL (with details for the agent interactions). 136
10.2 Steps followed by the policy process for the ACF model with the second level of complexity +Co (with details for the coalitions). 141
10.3 Steps followed by the policy process for the ACF model for the initialisation with the imperfect information add-on. 146

11.1 Policy instruments selected for the SIM and benchmark scenarios. . . . 162
11.2 Predation model results for wolf and sheep counts for the SIM and benchmark scenarios. 162
11.3 Policy core issues selected per affiliation for the SIM and benchmark scenarios. 163
11.4 Secondary issue preferred states for the benchmark scenario. 164
11.5 Policy core issue preferred states for the benchmark +PL scenario and scenario 2. 165
11.6 Policy core issues selected per affiliation for the benchmark +PL scenario and scenario 2. 165
11.7 Policy instruments selected for the benchmark scenario, and scenario 3 and 4. 167
11.8 Predation model results for wolf and sheep counts for scenario 4. . . . 168
11.9 Policy core issue preferred states for scenario 3. 168
11.10 Secondary issue preferred states for scenario 4. 168
11.11 Policy instruments selected for the benchmark scenarios. 172
11.12 Policy core issue preferred states for the benchmark scenarios. 172
11.13 Number of coalitions in the simulations for the +Co benchmark scenario and scenario 6 and 7. 174
11.14 Policy core issue 1 states for the scenario 8 with the colours representing the different levels of share resource parameter. 175
11.15 Secondary issue 1 states for the scenario 8 with the colours representing the different levels of share resources parameter. 175

11.16 Policy core issue preferred states for scenarios 1, 9, 5 and 10. 178
11.17 Actual beliefs of secondary issues of scenarios 1, 11, 12 and 15. 183
11.18 Actual beliefs of the policy core issues of scenarios 13. 185
11.19 Policy instruments selected for scenario 13. 185
11.20 Actual beliefs of secondary issues of scenarios 11, 14 and 16. 186

12.1 Policy core issue preferred states for scenario 2 and a demand growth scenario of 0%. 205
12.2 Secondary issue selected per affiliation for scenario 0 and a demand growth scenario of 3%. 206
12.3 Secondary issue preferred states for scenario 0 and a demand growth scenario of 3%. 206
12.4 Cumulative policy instrument selection for each demand growth for all fifty repetitions for scenario 1 and a demand growth scenario of 0%. . . 208
12.5 Cumulative policy instrument selection for each demand growth for all fifty repetitions for scenario 2 and a demand growth scenario of 0%. . . 208
12.6 Monthly averaged electricity prices, split according to the three policy process scenarios and for the three demand growth scenarios. 209
12.7 Overall demand with the electricity supplied by solar, wind and CCGT sources for all policy scenarios and for the 3% demand growth scenario. 210

C.1 Predation model populations for the predation model and the benchmark scenario. 252
C.2 Predation model populations for the benchmark scenario and scenario 1. 253
C.3 Policy instrument selection for the benchmark scenario and scenario 1. 253
C.4 Predation model populations for the benchmark scenario and scenario 2. 254
C.5 Policy instrument selection for the benchmark scenario and scenario 2. 254
C.6 Predation model populations for the benchmark scenario and scenario 3. 255
C.7 Policy instrument selection for the benchmark scenario and scenario 3. 255
C.8 Predation model populations for the benchmark scenario and scenario 4. 256
C.9 Predation model populations for the benchmark scenario and scenario 5. 257
C.10 Predation model populations for the benchmark scenario and scenario 4. 258
C.11 Predation model populations for the benchmark scenario and scenario 5. 259
C.12 Policy instrument selection for the benchmark scenario, and scenarios 4 and 5. 260

D.1 Cumulative policy instrument selection for a 3% electricity demand growth scenario for all fifty repetitions for each simulation. 262
D.2 Cumulative policy instrument selection for a 1.5% electricity demand growth scenario for all fifty repetitions for each simulation. 262

D.3 Cumulative policy instrument selection for a 0% electricity demand growth scenario for all fifty repetitions for each simulation. 263
D.4 Monthly averaged electricity prices, split according to the three electorate influence rates and for the three demand growth scenarios. . . . 263
D.5 Overall demand with the electricity supplied by solar, wind and CCGT sources for all scenarios. 264
D.6 Overall demand with the hydro electricity supplied for all scenarios. . . 264
D.7 Overall demand with the electricty supplied by NTC and LTC sources for all scenarios. 265

E.1 Predation model results for wolf and sheep counts for scenarios SM, 0, 1, 2, and 3 (+PL implementation). 268
E.2 Predation model results for grass patch count for scenarios SM, 0, 1, 2, and 3 (+PL implementation). 269
E.3 Policy instruments selected for scenarios SM, 0, 1, 2, and 3 (+PL implementation). 270
E.4 Policy core issues selected per affiliation for scenarios SM, 0, 1, 2, and 3 (+PL implementation). 271
E.5 Policy core issue preferred states for scenarios SM, 0, 1, 2, and 3 (+PL implementation). 272
E.6 Secondary issue selected per affiliation for scenarios SM, 0, 1, 2, and 3 (+PL implementation). 273
E.7 Secondary issue preferred states for scenarios SM, 0, 1, 2, and 3 (+PL implementation). 274
E.8 Predation model results for wolf and sheep counts for scenarios 0, 5, 6 and 7 (+Co implementation). 275
E.9 Predation model results for grass patch count for scenarios 0, 5, 6 and 7 (+Co implementation). 276
E.10 Policy instruments selected for scenarios 0, 5, 6 and 7 (+Co implementation). 277
E.11 Policy core issues selected per affiliation for scenarios 0, 5, 6 and 7 (+Co implementation). 278
E.12 Policy core issue preferred states for scenarios 0, 5, 6 and 7 (+Co implementation). 279
E.13 Secondary issue selected per affiliation for scenarios 0, 5, 6 and 7 (+Co implementation). 280
E.14 Secondary issue preferred states for scenarios 0, 5, 6 and 7 (+Co implementation). 281
E.15 Number of coalitions for scenarios 0, 5, 6 and 7 (+Co implementation). 282

E.16 Predation model results for wolf, sheep and grass counts for scenario 8 with the colours representing the different levels of the shared resources parameter. 282

E.17 Predation model results for wolf, sheep and grass counts for scenarios 8 with the colours representing the different levels of the coalition creation threshold parameter. 283

E.18 Policy core issue states for the scenario 8 with the colours representing the different levels of the shared resource parameter. 283

E.19 Policy core issue states for the scenario 8 with the colours representing the different levels of the coalition creation threshold parameter. 284

E.20 Secondary issue states for the scenario 8 with the colours representing the different levels of pooled resources from the agents joining the coalitions. 284

E.21 Secondary issue states for the scenario 8 with the colours representing the different levels of thresholds required to join coalitions. 285

E.22 Predation model results for wolf and sheep counts for scenarios 0, 5, 9 and 10 (+PK implementation). 286

E.23 Predation model results for grass patch count for scenarios 0, 5, 9 and 10 (+PK implementation). 287

E.24 Policy instruments selected for scenarios 0, 5, 9 and 10 (+PK implementation). 288

E.25 Policy core issues selected per affiliation for scenarios 0, 5, 9 and 10 (+PK implementation). 289

E.26 Policy core issue preferred states for scenarios 0, 5, 9 and 10 (+PK implementation). 290

E.27 Secondary issue selected per affiliation for scenarios 0, 5, 9 and 10 (+PK implementation). 291

E.28 Secondary issue preferred states for scenarios 0, 5, 9 and 10 (+PK implementation). 292

E.29 Predation model results for wolf and sheep counts for scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation). 293

E.30 Predation model results for grass patch count for scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation). 294

E.31 Policy instruments selected for scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation). 295

E.32 Policy core issues selected per affiliation for scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation). 296

E.33 Policy core issue preferred states scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation). 297

E.34 Policy core issue actual beliefs scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation). 298
E.35 Secondary issue selected per affiliation for scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation). 299
E.36 Secondary issue preferred states scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation). 300
E.37 Secondary issue actual beliefs scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation). 301

F.1 Monthly averaged electricity prices, split according to the three electorate influence rates and for the three demand growth scenarios. . . . 304
F.2 Overall demand with the electricity supplied by solar, wind and CCGT sources for all policy scenarios and for the 0% demand growth scenario. 304
F.3 Overall demand with the electricity supplied by solar, wind and CCGT sources for all policy scenarios and for the 1.5% demand growth scenario. 305
F.4 Overall demand with the electricity supplied by solar, wind and CCGT sources for all policy scenarios and for the 3% demand growth scenario. 305
F.5 Overall demand with the hydro electricity supplied for all policy scenarios and for the 0% demand growth scenario. 306
F.6 Overall demand with the hydro electricity supplied for all policy scenarios and for the 1.5% demand growth scenario. 306
F.7 Overall demand with the hydro electricity supplied for all policy scenarios and for the 3% demand growth scenario. 307
F.8 Overall demand with the electricty supplied by NTC and LTC sources for all policy scenarios and for the 0% demand growth scenario. 307
F.9 Overall demand with the electricty supplied by NTC and LTC sources for all policy scenarios and for the 1.5% demand growth scenario. 308
F.10 Overall demand with the electricty supplied by NTC and LTC sources for all policy scenarios and for the 3% demand growth scenario. 308
F.11 Cumulative policy instrument selection for each demand growth for all fifty repetitions for all policy scenarios and for the 0% demand growth scenario. 309
F.12 Cumulative policy instrument selection for each demand growth for all fifty repetitions for all policy scenarios and for the 1.5% demand growth scenario. 309
F.13 Cumulative policy instrument selection for each demand growth for all fifty repetitions for all policy scenarios and for the 3% demand growth scenario. 310
F.14 Policy core issues selected per affiliation for all policy scenarios and for the 0% demand growth scenario. 310

F.15 Policy core issues selected per affiliation for all policy scenarios and for the 1.5% demand growth scenario. 311
F.16 Policy core issues selected per affiliation for all policy scenarios and for the 1.5% demand growth scenario. 311
F.17 Policy core issue preferred states for all policy scenarios and for the 0% demand growth scenario. 312
F.18 Policy core issue preferred states for all policy scenarios and for the 1.5% demand growth scenario. 312
F.19 Policy core issue preferred states for all policy scenarios and for the 3% demand growth scenario. 313
F.20 Secondary issue selected per affiliation for all policy scenarios and for the 0% demand growth scenario. 313
F.21 Secondary issue selected per affiliation for all policy scenarios and for the 1.5% demand growth scenario. 314
F.22 Secondary issue selected per affiliation for all policy scenarios and for the 3% demand growth scenario. 314
F.23 Secondary issue preferred states for all policy scenarios and for the 0% demand growth scenario. 315
F.24 Secondary issue preferred states for all policy scenarios and for the 1.5% demand growth scenario. 315
F.25 Secondary issue preferred states for all policy scenarios and for the 3% demand growth scenario. 316

List of Tables

6.1 Preferred states for the policy makers on a the interval [0,1] for all scenarios. 72
6.2 Causal beliefs for the policy maker for all scenarios. These causal relations can be read as: an increase of 1 in S2 will lead to a decrease of 0.75 in PC1. They are all given on the interval [-1,1]. 73

7.1 Preferred states for the policy makers and the electorate on a the interval [0,1] for all scenarios. 95
7.2 Causal beliefs for the policy makers for all scenarios. They are all given on the interval [-1,1]. 96

8.1 Table presenting the different policy packages. 109
8.2 Preferred states for the policy maker on a the interval [0,1]. 111
8.3 Causal beliefs for the policy maker. They are all given on the interval [-1,1]. 112

11.1 Preferred states for the policy agents on the interval [0,1] for all scenarios. 160
11.2 Causal beliefs for the policy agents on an interval [-1, 1] for all scenarios. 160
11.3 Biased values for the external parties for the scenarios 11 to 16. For a value of 5, the external party informs itself directly from the truth agent. For a value of 0, the external party does not change its beliefs. 180
11.4 Initial actual beliefs on an interval [0,1] for the scenarios 11 to 16. . . . 181

12.1 Preferred states for the policy agents on a the interval [0,1] for all scenarios. 203
12.2 Causal beliefs for the policy agents on an interval [-1, 1] for all scenarios. 203

B.1 Minimum and maximum values used to map the indicators from the electricity market model to the secondary issues on an interval [0,1]. . 250

Glossary

ACF	Advocacy coalition framework
aff	Political affiliation
as	Agenda setting
B	Actual belief
Bo	Policymaker bonus
C	Causal belief
CCGT	Combined-cycle gas turbine
CLD	Causal loop diagram
+Co	Coalitions (model with)
DC	Deep core issue
Dom. em.	Domestic emissions
diff	Difference
G	Preferred state (Goal)
GDP	Gross domestic product
I	Instrument impact
IAD	Institutional Analysis and Development
IP	Investment priority
Imp. em.	Imported emissions
MAIA	Modelling Agent systems based on Institutional Analysis
MSF	Multiple streams framework
OLL	Old levee length
OLS	Old levee safety
P	Issue preference
PC	Policy core issue
+PK	Partial knowledge (model with)
+PL	Policy learning (model with)
PI	Policy instrument
+PI	Imperfect information (model with)
pf	Policy formulation

Continued on next page

Continued from previous page

R	Resources
REI	Renewable energy investments
RES	Renewable energy production
S	Secondary issue
SIM	Simplest implementation model
SFOE	Swiss Federal Office of Energy
SLL	Standard levee length
SLS	Standard levee safety
TA	Truth agent
W	Weighing constant

The preliminaries Part I

1 Introduction

Studying policy change consists of understanding the interactions between political actors, the public and the environment that might lead to the selection of a new policy. Broadly defined, the aim of policy change is to influence a continually evolving environment to satisfy the objectives of the policy actors and the public. This thesis presents a tool that can be used to build models of the policy process that can then be used to study the policy process and policy change.

1.1 The state of the literature

Policy change is omnipresent throughout the policy sciences literature. The literature has greatly evolved since it was initiated in the 1950s. Early on, the goal was to define the field itself. This included the definition of decision making but also included the creation of simple rational models such as incrementalism, the rational model and stages heuristics (Lindblom, 1959; John, 2018). Studies focused on the bureaucracy, policy networks or early models of the policy process. This includes the work of Hofferbert and his funnel causality model where the policy process is approached as a narrowing funnel composed of a series of decision making steps that include governmental institutions, the socioeconomic composition of the system or elite behaviours (Hofferbert, 1974; Wilder, 2016). The outcome of the funnel causality model being a policy output which is the result of a formal policy conversion.

In a second phase, the focus of the literature has been to understand the complexity of decision-making better. The description of agenda-setting was the first framework to do so Kingdon (1984). The development of several theories followed, including the advocacy coalition framework (Sabatier, 1987) and the punctuated equilibrium theory (Baumgartner and Jones, 1993; Baumgartner et al., 2014). Each of these theories was originally developed as a means to understand the policymaking and decision-making processes better. Since then, the theories have been enriched and advanced through a panoply of cases, validating elements of the theories in the context of different political and geographical situations. Further theories have also been developed to complement them and address some of their limitations or oversights (Nowlin, 2011; Petridou, 2014). Examples include the institutional analysis and development framework (McGinnis, 2011) or the diffusion theory (Stokes Berry and Berry, 1999). It is based on these policy process theories that the present dissertation is written.

Policy analysis is a part of the literature that has been used to understand better and to explain policy change (Bankes, 2011; Lempert, 2002; Thissen and Walker, 2013). Policy analysis, a discipline that also finds its origin in the 1950s, has for goal to assist policymakers in their selection of policies (Walker, 2000). This is done by consistently analysing policies based on a set of clearly defined parameters and goals. Over the years, the analyses have become more complex as our understanding of large and complex systems has increased. Policy analysis also had to take into account non-quantifiable aspects. Processes have been established to systematise policy analyses and deal with its associated complexity (Bots, 2013).

A variety of methods have been used for policy analysis, including modelling and simulation in more recent years (Berger, 2001; Chmieliauskas et al., 2012; De Vries et al., 2013). The use of simulation has helped address a growing number of scenarios that need to be considered in complex systems, often socio-technical systems, within which policy needs to be assessed. A range of different modelling approaches can be used to this purpose, including optimisation models (Mustapa and Bekhet, 2016), system dynamics models (Stave, 2003) or agent-based models (Ghorbani et al., 2014; Karslen et al., 2019). These models have increased the opportunities for researchers to understand the systems they are studying while providing more information on the impact of policies on decision-makers and political actors.

Policy analysis also plays a crucial role in policy change as it is, more often than not, used to inform decision-makers. The results obtained from these analyses can subsequently go on to influence the policy process itself depending on how agents within the policy arena receive these results. Therefore, it is paramount to get accurate results that reflect the potential of the policies assessed.

1.2 The limitations of current approaches

To better understand the focus of my dissertation, I look at the limitations both within the policy process literature and within the modelling and simulation literature which is often used for policy analysis. The goal of the dissertation will be to contribute to both of these fields at the same time, although they very rarely overlap.

1.2.1 The policy process

Currently, contributing to the theories used to study policy change consists of conducting qualitative studies. In the field, this translates in extensive data collection campaigns through surveys and a large number of interviews with key actors ranging from policymakers to organisations or even the public (Kammermann and Strotz, 2014; Sutter, 2011). Such efforts are time-consuming. Additionally, the knowledge gained through these methods is specific to the context of the case. This makes the conclusions difficult to generalise to other cases on different topics, or where the political regime or geographical areas are different. The combination of these factors can make the advancement of the theories a slow and tedious process. For every new concept or idea about a theory proposed, a new data collection campaign might be needed. This can be the case when the political landscape changes and some of the assumptions used do not hold anymore. It also prevents the use of theory exploration methodologies, an effort that can help speed up the advancement of the theories or at least make it easier to widen the scope of the research (Cairney and Weible, 2015).

Another limitation relates to the understanding gained from qualitative policy process studies. Most of these studies are descriptive, using the theories as a framework to structure their analysis. However, as John (2018) argues, they are not providing "explanation[s] to understand change in complex decision-making environments". These studies often do not provide the causes that might have led to the specific decision made by the policy actors. Gaining such insights could be very valuable for the process theories but also for the researchers, providing a more thorough analysis.

Finally, the current literature focuses on past events. All cases relate to events that have happened anywhere from a year to several decades before the study (Kriesi and Jegen, 2001; Jenkins-Smith et al., 1991). It is the nature of such studies as they cannot be performed without data. This approach is helpful to advance and validate the theories while understanding what happened (Alink et al., 2001). Nevertheless, it can only provide limited insights on what might happen in future, often rendering these studies theoretical and limiting their usefulness for current policymaking advice.

1.2.2 Modelling and simulation

Testing and analysing policies using modelling and simulation also suffers from limitations (Bankes, 2002). The largest of these relate to the way these methods deal with deep uncertainty. It is an aspect inherent to the simulation of models that can be linked to a multitude of external factors ranging from environmental changes to economic development or actor behaviours (Walker et al., 2010). Appropriately considering deep uncertainty can be critical to the results of a study. For example, to support long-term water management planning, there is a need to account for future climate change and sea-level rise, elements that are impossible to predict precisely (Kwadijk et al., 2010). This extends to a wide array of other fields where planning for policies requires to take into account many deeply uncertain aspects.

Several approaches have been developed to deal with this deep uncertainty. They vary from simple, such as the construction of scenarios, to more complicated, like exploratory modelling or dynamic robust planning (Kwakkel and Pruyt, 2015; Lempert, 2002; Pruyt, 2015; Walker et al., 2013). The more advanced methods have focused on adaptive policies (Swanson et al., 2010). These policies are more robust as they can adapt based on the information learned by the actors over time (Hallegatte et al., 2012). Swanson et al. (2010) has presented many ways one can design these adaptive policies. The literature goes further by differentiating adaptive policies themselves. For example, some studies use the adaptive policymaking approach (Walker et al., 2001; Kwakkel et al., 2010; Hamarat et al., 2013) while others have opted to use an approach referred to as the tipping point and policy pathways approach (Kwadijk et al., 2010; Haasnoot et al., 2011, 2012; Offermans, 2016). These methods have even been combined into a so-called dynamic policy pathway approach Haasnoot et al. (2013); Kwakkel et al. (2015).

Despite these numerous approaches, limitations are still present, relating to the implied assumptions built into them. For example, this is the case for the policy objectives. In early approaches, these were static, but adaptive approaches now consider dynamic objectives (Haasnoot et al., 2013). Despite such considerations, the objectives are still set by the modeller, or by decision-makers involved in the study. They do not take into account the evolution that the decision-makers will undergo throughout decades of influence and interactions in the policy arena. These changes are likely to affect the objectives of the policy actors concerning what is happening in the environment.

The same issue can be found for the creation or selection of policies to be tested in these simulations. Some studies consider the social system but never go as far as considering the decision-makers as part of the simulation (Haasnoot et al., 2011). In reality, the policy process literature, which is not used or considered within the modelling and simulation literature, has shown that policy selection is a product of negotiations or cooperation (Schneider et al., 2003). For example, the MSF assumes that many existing policies are floating in the policy stream, and it is up to the policy entrepreneurs to select one. The ACF considers coalitions and their role in policy change. Current simulations do not take into account these aspects detailed in the policy process theories. Such omissions can affect the results obtained while including them would help make the results more representative.

Finally, even if the policy introduced in a simulation is called ’adaptive’, it has not-so-far been genuinely adaptive. It only remains semi-adaptive where simple policies would have been fully rigid. By this, I mean that current adaptive policies can adapt only to the extent of the scenarios, i.e., the modeller-defined pathways and decision trees. The policies are not implemented based on a continually evolving decision-making process. As a result, and due to emergent properties present within any simulation, it cannot take into account all potential outcomes of a model.

A different literature has attempted to deal with this issue by considering socio-technical simulation as control systems. Depending on the objectives set by the modeller, a controller - which is a proxy for the policy actors - is used to introduce policies so that the model achieves the modeller’s objectives. This approach has been suggested in Lempert et al. (2009) as an adaptive control approach. In this case, the controller bases its decision on the feedback and adaptation rules set by the modeller. The policies are assessed using a risk-benefit-cost assessment. This method does not consider policy agents or how their objectives might change over time. This re-introduces several of the limitations encountered in the previous approaches.

1.3 The research objectives

Based on the current state of research and the limitations highlighted, I formulate the following research question:

To what extent can simulating policymaking be used to explore the policy process theories and help study policy impacts in socio-technical systems?

Three additional sub-research questions are formulated to provide a roadmap to answer this larger research question.

1. Can a systematic approach for the creation of policy process models be formulated?
2. In what way can a policy process model be simulated within a socio-technical system?

3. What insight can be gained from the simulation of the policy process within a socio-technical system?

The overall objective of the thesis is to address the limitations mentioned may it be for the policy sciences literature or the modelling and simulation literature. The contribution of this thesis will be entirely methodological and consists of the creation of a tool that can be used to create policy process models to address these limitations.

1.4 Research design

A modelling and simulation approach is used to answer the research questions presented above. Such approach can help deal with limitations outlined in both the policy sciences - theory exploration possibilities, and the modelling and simulation - policy-related deep uncertainty - literature.

1.4.1 Addressing the limitations

Modelling opens the door for theory exploration when it comes to the political sciences. With current methods, as was mentioned earlier, theory exploration is possible but is a lengthy process. Using modelling and simulation, one can explore the theories by building models of these theories and tweaking the assumptions depending on new hypotheses being considered. This allows the researcher to test different assumptions either to better match empirically obtained results or to help discover new emergent phenomenon in the simulation. The hope with this thesis is that the approach provided here will be used to test more assumptions for varying policy process theories and hence help advance the field.

Models of the policy process can also be useful when it comes to decision-maker advice. Currently, studies using the policy process theories focus on understanding and explaining past behaviours. Through simulations of the policy process, researchers will be able to gain new insights that they can use to inform the decision-making process of real actors. This includes better understanding how other actors are likely to behave, how their beliefs might evolve and what policy instruments are most likely to be selected in the future. This does not mean that the results obtained from these simulations should be used as prediction. Instead, they should be approached as explanation and used for exploration purposes.

When looking at the field of modelling and simulation, the use of a policy process model can be seen as an additional tool for researchers dealing with policy-related deep uncertainty. Currently tools dealing with deep uncertainty tend to avoid policy aspects and focus on more environmental phenomenon as was shown above. Being able to endogenise the policy process in current simulations will help better take into account the impact of policies introduced as a result of prior events within the model simulation itself.

The use of policy process models can also benefit the field of policy analysis. By combining a model of the policy process with a socio-technical model, it will be possible to explore the feedback effects that are the results of the interactions between the policy subsystem - the policy process model, and the socio-technical system - the environment model[1]. This combination is already present in the literature in its qualitative form (Edmondson et al., 2018). It has been referred to as a co-evolution, and several feedback effects have already been categorised. Currently, policy-linked feedback effects are often missed because of the way deep uncertainty is dealt with.

1.4.2 The methodological approach

Modelling and simulation does not spell out the approach but only provides a direction to be followed. There are still different options possible. In the present case, two approaches are considered most viable and both answer the main research question formulated above. One approach is empirical, the other is methodological. For the empirical approach, one would select a specific case study. The goal would be to build one, or more, policy process models that can best fit the case and that are best suited to answer additional, case-related research questions. For example, such models could be used to see the impact of different policy process theories on one case study. These are models that would require thorough validation as they would be used for one case study, and would in turn have limited predictive power.

For example, two models could be considered. One could be based on the advocacy coalition framework while the other on the feedback theory, if these were theories that best fit the case. Such an empirical approach allows to gain more insight into the case studied while answering the research question formulated in this thesis. The drawback is that it would be heavily tailored to the case selected, making the development of a systematic approach more difficult. It would be a one-off study and additional cases would require the creation of new policy process models.

The alternative approach is a methodological one. This is the approach that is selected for this thesis. Note that the empirical approach presented above would is still feasible

[1]Throughout the dissertation, I will refer to socio-technical systems as environment models and vice versa.

following the work presented in this thesis and can benefit strongly from it. The goal with the methodological approach, beyond answering the main thesis' research question, is to make it easier for researchers from both the policy sciences and the modelling and simulation fields to use one another's theories and approaches. At the moment, the connections between both research fields are limited as is demonstrated in the literature review in Chapter 2. In this thesis, I propose a tool that aim at lowering the barrier to entry into either fields for researchers, thus helping bridge both disciplines. This tool provides a consistent approach to the construction of policy process models.

This tool, the common language, provides a set of requirements that need to be met by any model for it to be considered a model of the policy process. Using these requirements, researchers can build models of the policy process that can be tailored to specific theories or cases. This thesis consists of the presentation of this tool and its subsequent demonstration through the construction of two policy process models.

1.4.3 Operationalising the theories

Before the presentation of these two models, it is important to consider the issues that arise with the conceptualisation and formalisation of the policy process theories. To model the policy process, there is a need to reconcile concepts that are not necessarily considered compatible within the literature (Cairney, 2013). This is needed because each theory has a specific focus and no single theory encompasses enough of the policy process to be simulated alone. There is also a need to operationalise the policy process theories, something that will require additional assumptions currently not present in the literature.

For example, when considering agenda setting, no theory explains how, in an operationalised way, the policy entrepreneurs bring the streams together. A loose description is made of this action, but this description remains insufficient for its simulation (Zahariadis, 2014). There is a need to specify what such action exactly entails to construct a model. This can also be said for several other concepts such as for the belief system - and the way it helps actors' decision-making, or for coalitions - and the way they are created and they strategise in the policy arena (Sabatier and Weible, 2007).

Issues relating to the conceptualisation of models do not necessarily have to be seen negatively and can instead be leveraged as advantages. They put a spotlight on parts of the theories that are overlooked or intentionally avoided in the current literature. Hence, it forces researchers to deal with these elements and can be beneficial to the literature. The downside is that the assumptions formulated to deal with these issues require validation, a task that require a lot of work and often empirical cases. In the present thesis, the validation of the individual assumptions will not be performed as it is considered out of the scope. The focus is placed on the approach as a whole. However, this dissertation is written with the hope that future researcher will focus on the validation of such assumptions to advance the theories.

1.4.4 The model selection

The two policy process models

The first policy process model built, referred to as the simplest implementation model, is literally the simplest possible implementation of the common language. It is selected for its simplicity and to ease the results' exploration. It contains the minimum number of elements that are needed to simulate the policy process using the present approach - hence its name. The second model that is demonstrated is an emulation of the advocacy coalition framework (ACF). The goal with that model is to demonstrate that it is possible to build models that are based directly on the theories. For this, a number of elements from the framework itself are incorporated into the model including policy change through policy learning and coalitions (Sabatier and Weible, 2007).

Many more models can be constructed with the common language but because of the amount of time required to properly conceptualise, formalise, code and test each model, only these two were considered. The simplest implementation model was selected because of its simplicity while the ACF model was selected because of its wide use in the literature, arguably the policy process theory that is the most used. Other notable theories that could also have been selected are the multiple streams framework, the diffusion theory or the feedback theory (Stokes Berry and Berry, 1999; Jones et al., 2016; Mettler and SoRelle, 2014).

The three environment models

The policy process models created using the common language cannot be tested in a vacuum. A context needs to be provided for them to be simulated. It takes the form of an environment model. For this reason, a number of environment models are considered to demonstrate the policy process models. Two avenues were privileged for each policy process model here: demonstrating how theory exploration can be performed and how policy-related deep uncertainty can be endogenised in a simulation. For this two environment models were selected, each with different requirements.

On the one hand, for the theory exploration, the goal is to study the policy process model. There is therefore a need to have as simple and as well studied as possible an environment model. Numerous agent-based models have been used over the years but the model needs to have aa somewhat stable periodic behaviour as well. For example, a forest-fire model was initially considered but such system can enter irreversible behaviours, such as the destruction by fire of the entire forest, leaving the potential policy instruments helpless (Wilensky, 1997). For this reason, a predation model is considered (Wilensky and Reisman, 2006). It is a well-studied, simple and easy to understand model. Its behaviours are cyclical and even if one population of agents dies off, the model can still be impacted by policy instruments.

On the other hand, for deep uncertainty, the goal is to study the complex interactions between the policy process model and environment models. To do this, there is a need for a large complex environment model. Any complex socio-technical model can be used for this purpose. Because the focus of the thesis is to test these models and not build them, an already existent model of the Swiss electricity market was considered (van Baal, 2016). It is a fairly complex model where several emergent behaviours can be expected and for which a number of policy instruments can be designed.

A third and final model was also sought to illustrate the versatility of the modelling approach. The common language is suitable to build agent-based policy process models only. These can then be connected to environment models using a multitude of modelling paradigms. To demonstrate this, a different modelling paradigm was sought and was found in the form of a flood safety system dynamics model developed by the author in previous work (Klein et al., 2016). Because the goal with this model is only to illustrate that it is possible to connect different modelling paradigms, this is only done with the simplest implementation policy process model.

Data gathering and analysis

The outline of the different models above leads to the consideration of five hybrid models in total throughout the dissertation. Three hybrid models that on the one hand include the simplest implementation policy process model and on the other either the predation, the electricity market or the flood safety models. And two hybrid models that on the one hand include the ACF policy process model and on the other the predation or the electricity market models. To simulate each of these models, data is required and the results will have to be analysed.

The sourcing of the data required to simulate these models differs from one model to the other. For the predation models, the data used is dummy data because this model is not directly linked to a specific real life system. It is used as a simple and tested model to explore the potential of the policy process model. The initial data is therefore chosen so as to best display interesting dynamics - as opposed to dynamics that lead to systems where one animal population or the other crashes rapidly after initialisation.

The electricity market used being a real model, the approach is different. The model comes with its own data as presented in van Baal (2016). This same data is used for the present work as it has already been thoroughly validated and leads to an acceptable outcome. This is data that sees the Swiss electricity market simulated from the year 2015 onwards.

For the flood safety model, the data used is dummy data once again. This is data that comes with the original model and that was used in Klein et al. (2016). The flood safety model is a conceptual model and therefore does not use realistic data. The goal of this model is to observe specific behaviours and trends between important parameters such as feedback loops or stable equilibria for examples.

Finally, there are the two policy process models. For most implementations, the data that is used to initialise these models is dummy data. This relates to the goals of the dissertation but also to the fact that these models are not representative of any real political system. For the simplest implementation model, the resulting policy process is so simple that it would not be possible to validate its use with real life data. The data used is therefore a wide interpolation of potential real trends that are used to help show the potential of the models and its results. As an example this means considering agents that are either pro or against wolves and sheep.

For the ACF model combined with the electricity market model, a different approach is used. Because this model is closer to reality - it considers coalitions and is coupled to a model representing the Swiss electricity market - and because work has already been performed to obtain the beliefs of the two Swiss coalitions in the electricity market in (Markard et al., 2016), the data used to initialise that model is inspired by the results obtained in Markard et al. (2016). However, as a result of the limitations of the ACF policy process model and because the research done in this paper was not done specifically with the present models in mind, the data used is sometimes interpolated. If one were to conduct a proper complete policy study, then research similar to Markard et al. (2016) would need to be redone so that the beliefs of agents included in the models be more accurate along with the compositions of the coalitions.

For the data analysis of the results from the simulation of all five models, the approach is limited. It consists solely of graphical visualisations. This is considered sufficient to answer all of the smaller research questions asked throughout the testing of the different models. In the case of a policy study, a more thorough analysis would be needed, including but not limited to statistical modelling. This is not considered in the present dissertation because of the nature of the goals of the dissertation.

1.5 The research process

The research process is designed with the single goal of demonstrating the use of the common language and the possibilities brought by the modelling and simulation of the policy process. The entire research process is illustrated in Figure 1.1. The goal with this process is to first present the common language and then, in a deliberately repetitive manner, demonstrate how this common language can be used to build specific policy process models, and how this model can be useful for either theory exploration or deep uncertainty exploration. Once this has been completed, the results obtained are discussed, concluding the work and outlining a number of possible avenues for further work.

The first part of the dissertation is dedicated to the development of the common language. This part therefore contains an in-depth literature study. This study is wide and encompasses both the topics of the policy process theories and the literature surrounding modelling and simulation. Such a wide review is required because of the dissertation's multi-disciplinary nature. The review performed is methodical, first introducing modelling and simulation before focusing specifically on modelling and simulation in the political sciences. This is then complemented by a in-depth review of the three policy process theories that are used to build the common language, namely the multiple streams framework, the policy cycle theory and the advocacy coalition framework. This literature review is followed by the presentation of the common language using elements reviewed in the literature itself. The common language is accompanied by the presentation of the modelling approach that is to be used in the dissertation and that is proposed for other researchers that would want to use the common language. This consists of a hybrid modelling approach which simplify the coupling of almost any system dynamics or agent-based model with a policy process model formalised using the common language.

The second and third parts of the dissertation share a similar structure. The reason behind this shared structure is the need to test two different policy process models using the same method and steps. The first policy process model is a simplistic model while the second is one that closely emulates the advocacy coalition framework. Each model is first presented with a conceptualisation followed by a detailed formalisation.

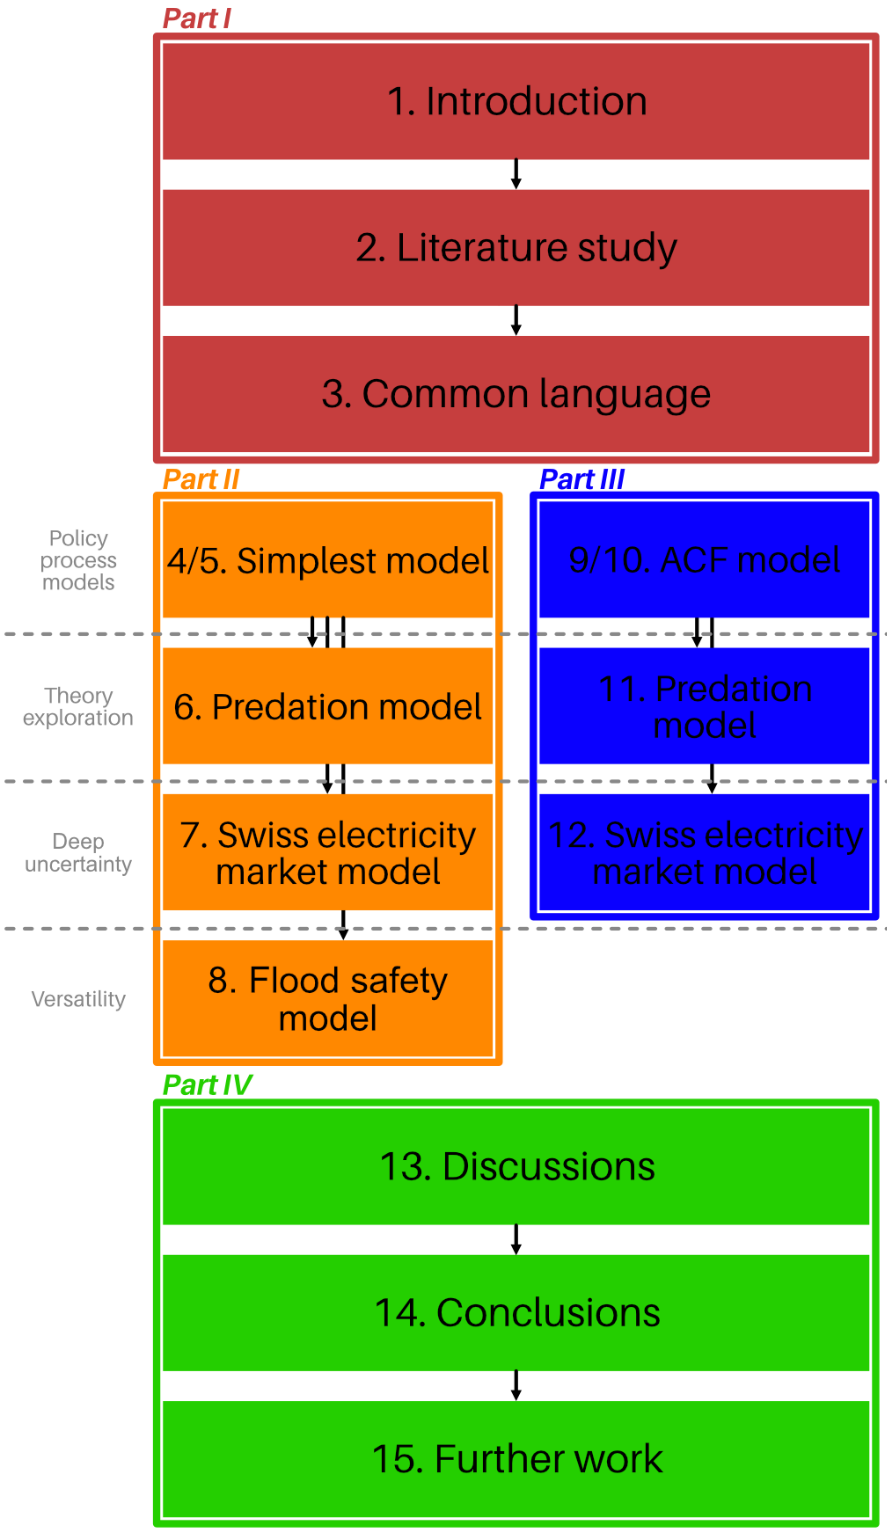

Figure 1.1 – The research process.

This helps highlight the assumptions taken and how other researchers might differ in their choices. Second, each model is simulated with a predation model with the goal of demonstrating how theory exploration can be performed. Third, each model is simulated with the electricity market model. In this case, the goal is to demonstrate how such simulation can help better understand deep uncertainty within the combined resulting model. Finally, and this is the case only for the simplest policy process model, it is simulated with a flood safety model to demonstrate the versatility of the hybrid modelling approach. This does not require two policy process models and is therefore only done once.

By having two parts with the same structure, it makes it makes it easier to illustrate how different models can be built using the common language and how they can be simulated using the same process. It highlights how one can fairly easily use different environment model to answer specific research questions may it be relating theory exploration or the exploration of deep uncertainty. The structure used also highlights the steps that have to be followed to go from the common language to simulation results. This includes the conceptualisation and formalisation of a policy process model, the hybridisation of the environment model, the selection of a research question, the design of the simulation experiments and finally the simulation of the hybrid model. Regardless of the environment or policy process model used, the steps will mostly remain the same with varying levels of depth.

The fourth part of the dissertation consists of the analysis of the dissertation's approach. That is, the use of a common language to formulate a policy process model, its hybridisation with an environment model and ultimately its simulation. This part outlines a number of strengths inherently present with this approach and the simulation of the policy process. It also highlights key weaknesses that can affect the results of the simulations or render the entire approach futile. Through this reflexion exercise, a number of possible changes are also highlighted. They are subsequently formulated into potential avenues of further research for researchers that might be interested in further developing the common language or further pursue policy process modelling research.

1.6 The dissertation structure

This dissertation is split into four parts as outlined previously. The first part consists of the preliminaries and is made of three chapters. This includes the introduction, which has highlighted the motivation for this thesis, along its research goals and design. This is followed by a literature review of both the policy process theories and the field of modelling and simulation in Chapter 2. The literature explored is then used to develop a common language in Chapter 3. This language is the main tool developed for this thesis to be used as a systematic approach to build models of the policy process, establishing requirements on what building blocks such model should incorporate. This chapter also presents the modelling approach to be used in subsequent chapters.

The second part of this dissertation consists of the simplest implementation model. In this part, I demonstrate the construction of the simplest model that can be created using the common language, and that can emulate the policy process. First, the conceptualisation of this model is described in Chapter 4. This is followed by its detailed formalisation in Chapter 5. Then comes a series of three examples with three different environment models. First, a predation model is used to highlight theory exploration in Chapter 6. Second, in Chapter 7, a Swiss electricity market model is used to demonstrate what benefits this endogenisation can bring to dealing with policy-related deep uncertainty. Finally, a flood safety system dynamics model is used to demonstrate the versatility of the approach used in Chapter 8.

The third part of the thesis consists of another implementation of the common language. This time, the goal is to simulate a specific theory: the advocacy coalition framework or ACF. This parts follows a similar structure to the previous part with a conceptualisation in Chapter 9, its formalisation in Chapter 10 and two examples. Similarly, the predation model is used to outline how this more complex model can be used for theory exploration in Chapter 11. The Swiss electricity market model is then used to assess whether some of the limitations found with the simplest implementation model can be bridged with the ACF implementation model. This is shown in Chapter 12.

The final part of this dissertation is more reflective. First discussions are presented in Chapter 13. This chapter reflects on the work that has been done, looks at the bigger picture, and outlines limitations that might have been overlooked. In Chapter 14, I outline the conclusions that can be drawn from this work, answering specifically the research questions that are presented in this introduction. Finally, in Chapter 15, I propose several avenues that can be explored for further research. This includes work that can be done to extend the present dissertation, work that has for goal to deal with some of its limitations and work that can simply use the present work to start new research branches.

2 Literature review

The policy process theories' literature is barely cited in the modelling and simulation literature. The opposite is also the case with only a very limited set of modelling and simulation literature within the policy sciences. This gap between the two fields can be a hurdle for researchers that consider using elements from both literatures together, in their goal to answer their research questions. One of the goals of this thesis is to show that both literature can benefit from the other. The present chapter is a review of efforts such as this one that have been made in the past. It outlines how far each of the literature has come and where the overlap currently exists.

First, this chapter goes over the research that has been done in the field of modelling and simulation. The review starts broadly with a description of the field and the role of models. There, the argument is made for the use of modelling and simulation to answer the thesis' research questions. This is then complemented by an outline of the different modelling paradigms that are present in the literature. Finally, the focus is placed solely on the political sciences and how modelling has been used in the field.

The chapter is completed with a review of the policy process theories, with the focus on three in particular. There are many more theories present in the literature but these three - the policy cycle theory, the multiple streams framework, and the advocacy framework framework - are theories that will form the basis of the common language developed in the following chapter. They are therefore the focus of the review of this chapter.

2.1 Modelling and simulation

Models are tools that can be used for a range of different purposes (Calder et al., 2018; Edmonds et al., 2019; Epstein, 2008; Saltelli et al., 2020). Models can be used to predict.

Prediction, in this case, is often considered in the weak sense of the term as it is difficult to predict what will happen in a complex system. Because of these difficulties, models are often used for explanation and exploration instead of outright prediction (Edmonds et al., 2019). This is useful for complex systems where it is difficult to understand the overall behaviour and where modelling allows the researcher to gain a better understanding of the causal chains present within the system. The explanatory power can extend to the understanding of theories. This understanding can be gained through the conceptualisation of a phenomenon, or through repeated simulations and tweaks of a model to obtain the empirically observed behaviour. This can ultimately be used for theory building (Henry and Brugger, 2018). Models can also be used as ways to visualise and illustrate complex systems (Edmonds et al., 2019). Finally, models can be used as thinking aids through analogies.

In the present case, the approach being entirely methodological, the work presented in this thesis will not be following the prediction route. Instead the goal will be to provide a tool for researchers to explore the policy process theories or deep uncertainty. This means the role of the model will be most explanatory and exploratory (Edmonds et al., 2019; Epstein, 2008). Through the construction of the tool, and the subsequent models, important insights and contributions are expected to be made.

One of the key aspects impacting the usefulness of a model is the consideration of what should be in the model. Models are an abstraction of reality, a simplification that leaves much of the described phenomenon out (Jager and Edmonds, 2015). They can be very complicated and include a copious amount of details. Models can also be elementary and only depict the most straight forward parts of a system. The choice between complexity and simplicity is one that has to be made by the modeller, and that should be made based on the research question that the model will be used to answer. More complex models do not necessarily lead to better results. They can often lead to a large number of results that are then impossible to interpret and explain. On the other hand, a model that is too simple might not be sufficient to answer a research question and therefore make the model useless.

When modelling, it is also essential to consider the data that is available to the modeller and the data that will is ultimately used to simulate the model (Nikolic et al., 2019). The data available has an impact on the model parametrisation. The model should avoid parts for which data is not available or not collectable. This can have an impact on the scope of the model that is built. Beyond the impact on the model itself, the data used to simulate the model can have an impact on the model results. If the data that is used as input to the model is incorrect, then the result outputs of the model will also be incorrect. Therefore, it is important to check sources and justify the use of data throughout a simulation.

Uncertainty also plays a crucial role in data inputs. All of the data that is introduced in the model is uncertain. The modeller needs to have an understanding of the margin of error introduced for each of the variables considered. The input data can also be a scenario related to specific parameters. Scenarios are used to explore the model in specific directions. These scenarios usually entail the consistent modification of a set of pre-defined parameters. These changes are performed to answer a specific research question. For example, a model can be used to understand better the effects of different growths of the GDP on a country. For each different growth profile, a different scenario can be constructed.

A model and its results cannot be taken at face value without performing a validation. The goal is to show that what is modelled is consistent with reality. Validation can be challenging to perform depending on the model being considered (Nikolic et al., 2019). It can take different forms. The standard validation is to look at whether the model represents reality. This can be done by comparing the results of an experiment and the results of the model. The difference between the results provides a measure for the validity of the model. Social systems cannot be validated this way as they involve subjective and normative perspectives. Validation can be performed by using expert or face validation. However, this kind of validation is prone to selection bias. Another common validation method is historical replication where historical data is compared with the results of the simulation over the same period. Such validation is however not necessarily sufficient to infer on the validity of results obtained after the historical period tested.

2.2 The modelling paradigm

Several different modelling paradigms are used throughout the literature. These vary from economic models to optimisation models, dynamic simulations, statistical models, technical and design performance models, and data-driven models (Nikolic et al., 2019). Each modelling paradigm can be seen as a tool to answer a research question. There is no correct or incorrect modelling paradigm, merely a need to justify why a particular paradigm is chosen and how it will be suitable to answer the research questions.

The policy process has been described within the different theories as a social system where actors interact with one another, and where the ultimate output of this process is a set of policy instruments. Most modelling paradigms are not capable of modelling social systems. Dynamics simulations can handle social systems taking into considerations their subjective and normative perspective. Agent-based modelling, more specifically, is the paradigm that is selected (De Marchi and Page, 2014). Agent-based simulations can emulate actors. They allow agents with individual preferences, strategies and knowledge, and interactions with the environment. This allows for patterns to emerge. Agents can also adapt their plans of action through learning, interactions and institutions (Johnson, 1999).

Because of its unique ability to both address agency aspects and the process, an agent-based model architecture can provide an approach that can adequately emulate the policy process. Beyond this, agent-based modelling is also most appropriate for any further work that would be performed. It allows for the exploration of the theories. This can be done, for example, through different hypotheses on actor behaviour specific to different theories. Additionally, agent-based modelling can be used to test other more encompassing phenomena such as the impact of imperfect information or bounded rationality on the outcome of the policy process.

2.3 Simulation in political sciences

The idea of using simulations as a formal research method in political science is not new. In his review of the evolution of simulations in political science, Johnson argued that computer models of political behaviour have been around since the 1960s when William McPhee and Robert Smith proposed a strategy similar to that of agent-based modelling, mainly to base models on knowledge of how individuals behave (Johnson, 1999). However, subsequent generations of researchers did not answer their call to understand the influence of micro-scale processes on macro-scale political phenomena, and instead adopted mathematical and analytical modes of analysis rather than furthering their simulation efforts until the 1990s. While informal logic, empirical research methods and data collection and mathematical modelling have dominated the field for the past 50 years, Johnson saw a growth in the number of political scientists using simulations even as he wrote in the late 1990s. He mentions, for example, the application of simulations to ascertain the probability of voting cycles, simulations of individual members of congress' positions and decisions based on the expected effects of a bill on their constituents or alternatively on their interaction with other members of Congress, votes for presidential nominations, simulations of conflict between nations, Robert Axelrod's simulation of the prisoner's dilemma.

Modelling and simulation has been used sporadically prior to Johnson's article in the domain of world politics. A number of studies involved interactions between humans and computers (Guetzkow et al., 1963), and more recently, these have been using full computational simulations (Cederman, 1997, 2003). Pepinsky (2005) warned future researchers on the use of modelling and simulation within the context of world politics simulation. He highlighted more fundamental limitations of this methodology: "its epistemological foundations and its ontological presuppositions". However, in his paper, Pepinsky only focused on areas of the literature that he believed lacked research. His goal was not to discourage future researchers from using these novel methods.

In Earnest (2008), this advice is followed. Through a model of international negotiations that investigates cooperation problems, an agent-based model is highlighted for its advantages compared to results that would have been obtained from a game theoretic approach. Similarly, in Bloodgood and Clough (2008), the influence of NGOs in world politics is investigated through an agent-based model discovering interesting and then unexplored implications. Geller (2011) also outlines the use of complexity-based models, such as agent-based models, to better explore and understand international relations. In its discussion, Geller mentions a number of models including a model focusing on resources, ethnicities and conflicts (Bhavnani et al., 2008). In this specific model, the researchers attempted to better understand the link between natural resources and civil wars.

The political sciences at large have also shown interest in the use of modelling and simulation, and not only agent-based models. Initially, this mostly took the shape of game theoretical models (Kollman et al., 2003). However, simple game theory models are often considered insufficient. They have therefore been enhanced with a number of add-ons. This is case in Bendor et al. (2003) where behavioural aspects are brought into the model, including the use of a basic learning cycle. The paper tests how the introduction of such add-on might change the outcomes observed in classical games such as the prisoner's dilemma game. In Baron and Herron (2003), game theory is also used but this time in a more applied manner, to understand the formation of governments and changes in entitlement programs and regulations.

The study of party competition has also benefitted extensively from the use of modelling and simulation (Lin, 2008). Studies can be found throughout the literature over the past two decades. In Jackson (2003) the focus is placed on party elites and the studies of preference evolution among the mass public. In De Marchi (2003), it is voter sophistication and ideology that are made essential parts of the model. This is another study loosely focused on party competition but this time with a focus on campaign financing and how this impacts diffuse interests. Bailey (2003) uses probabilistic voter theory to investigate the impact of campaign spending and, using a model, identifies optimal policies and campaign choices. Using game theory, Laslier and Goktuna (2016) goes further and studies the incentives that are present institutionally for politician to switch parties within the context of party and electoral competition.

In Kollman and Page (2005), the argument is made for the use of modelling, and more specifically agent-based modelling, for the study of politics. It is argued agent-based model can provide insights in the actors' behaviours. A number of domains are outlined such as studying adaptation - selective pressures and learning, understanding differences, looking at externalities, focusing on path dependence, considering geography, studying networks and exploring emergence.

Political scientists have used formal modelling techniques to study the ins and outs of the political system - from the effect of allocation of responsibilities and appointments to ministries on a Prime Minister's ability to implement her policies (Dewan and Hortala-Vallve, 2011), to the ways "civil service systems of personnel management interact with bureaucratic discretion to create expert bureaucracies populated by policy-motivated agents" (Gailmard and Patty, 2007). Researchers have also been using agent-based models to represent organisations, connecting the two levels of agent and organisation (Chang and Harrington Jr, 2006). Campaign dynamics have also been explored using agent-based models (Gulati et al., 2011).

Within the field, research on polarisation has also benefitted from the use of modelling and simulation. In the case of Li and Xiao (2017), the researcher has focused on opinion polarisation as part of a collective behaviour phenomenon. In this case, social judgment theory was used within the context of an agent-based model. Zhang et al. (2020) focused more on group polarisation and understanding the different interaction modes and their effect on group attitude polarisation. These models are derived from the classic polarisation D-W and J-A models.

Models are also used to study and understand specific processes. This is the case for Scott et al. (2019) which looks into consensus and stakeholder agreements in governance processes using an agent-based model. This is needed to better account for the many stakeholders involved in reaching agreements following a range of different scenarios. This approach is one that was already used in Choi and Robertson (2014) in a similar context: to understand deliberation and decision in collaborative governance.

The literature has also focused on collective decision-making, using a range of modelling paradigms (Thomson et al., 2003). Two main approaches often cited are: the expected utility model from Bruce Bueno de Mesquita (Bueno de Mesquita, 2011) and the exchange model from Frans Stokman, Reinier van Oosten and colleagues (Bueno De Mesquita, 1984; Bueno De Mesquita et al., 1985). Each approach considers that to come to a decision, actors follow four steps but with varying the definition of content of these steps for each of the approaches. Both approaches use game theory as modelling paradigm and are closely or loosely derived from the median voter theorem in the way they solve the outcome of the games they set up (Black et al., 1958). The expected utility model uses non-cooperative game theory and has been used in a number of cases with the explicit purpose of forecasting (Bueno De Mesquita, 1984; Bueno De Mesquita et al., 1985; Bueno De Mesquita, 2006; Sprinz et al., 2016). The exchange model on the other hand uses cooperative game theory and has focused largely on network dynamics (Stokman and Van Oosten, 1994; Stokman et al., 2000). Both models have their advantages and drawbacks as the literature shows (Achterkamp and Akkerman, 2003; Torenvlied and Thomson, 2003).

Within the scope of the policy process, modelling and simulation has also been used to study policy networks and social networks in general. In (Henry and Brugger, 2018), an agent-based model is used with the goal of theory building to understand the impact of social network on firm choices. Agent-based modelling has also been used to test hypotheses related to a social network in Snijders et al. (2010). Statistical modelling is also used extensively to attempt to understand, explain or even model social networks with potential applications for the policy process (Snijders, 2001; Steglich et al., 2010; Snijders et al., 2010). For example, this approach has been used to look into innovation adoption in Ghana (Boahene et al., 1999).

One key study that has used agent-based modelling is the one made in Laver and Sergenti (2011) relating to party competition. It was developed over a number of years and through a number of iteration of their party competition model (Fowler and Laver, 2008; Laver and Benoit, 2003; Laver, 2005; Benoit and Laver, 2006; Laver and Schilperoord, 2007). This model is very different from previous models on party competition presented in Kollman et al. (2003). It has used agent-based modelling as a tool to explore party competition. It was first approached as a very simple model. To capture more of the complexity observed, complexity was also added to the model layer by layer. The result was a model that can be used to better understand how parties behave and why their beliefs evolve over time under the pressure of their electorate.

A few researchers have also looked at using policy process theories within their models. This is the case of the SimPol model - a generic model of a basic polity at three different resolutions. To build the model structure, Cioffi-Revilla (2009) combined selected features from the work of political science theorists such as D. Easton, R.A. Dahl, K.W. Deutsch, K.V. Flannery, G. Johnson, J.W. Kingdon, and H. Wright. However, in his description of the model, the author does not go into great detail regarding how each theory influenced model design. In subsequent work, Cioffi-Revilla (Cioffi-Revilla and

Rouleau, 2010) used agent-based modelling to extend the theoretical understanding of civil unrest through the explicit representation of micro-level dynamics between society and government, aiming to not only draw on political science theories but also enrich them.

One work that relates to the study of the policy process was first presented in Ghorbani et al. (2013). In it, the MAIA framework (Modelling Agent systems based on Institutional Analysis), built using agent-based modelling, focuses on the simulation of the Institutional Analysis and Development framework instead of looking at the theories of the policy process (McGinnis, 2011). The resulting framework is one where actors introduce new rules based on their strategies and specific evaluation criteria. This is analogous to the policy process where policies are implemented by actors based on their beliefs and preferences.

2.4 Policy process theories

In his book, Sabatier presents over half a dozen policy process theories, ranging from diffusion models to the network approach (Sabatier, 2007). Each of these theories has a different focus and is used with a different goal in mind. Some of these theories focus on the policy network, such as the network approach (Adam and Kriesi, 2007). Others focus on rules and institutions such as the case of the Institutional Analysis and Development (IAD) Framework developed in Ostrom (2007). Overall, each theory presented by Sabatier (2007) deals with different aspects of the process, and in some cases overlap, looking at the same parts of the process while describing different dynamics.

This wealth of theories, and accompanying studies, results in incoherences between the theories. It can be a strength for qualitative studies that require different and specific approaches and can help make room for the broadest possible range of cases. However, it becomes a problem when modelling and simulation is considered. For a system to be simulated, the modeller has to provide a consistent and detailed set of instructions. This must include steps, along with exact algorithms that define how actors make decisions and interact, and how the system proceeds from a one-time step to another.

Such model-building is not something that can be done by using either a single theory to simulate the policy process, or by considering all theories at the same time. Each theory has a different focus and uses a different language, leading to most theories being unable to represent the entire process on its own or to inconsistencies between the theories (Cairney, 2013). Instead, the models will have to consider a combination of theories that might only include a few assumptions from each of the theories. Arguments have been made to reconcile some of the theories into a unified framework (Howlett et al., 2015). However, this is not the topic of this research. With this in mind, there is a need to select a few theories that are wide-ranging enough to consider all parts of the policy process but that do not necessarily tackle all possibilities and all of the details that the process can sometimes entail.

For the common language, the focus is placed on three theories in particular: the policy cycle, the multiple streams framework and the advocacy coalition framework. The two latter frameworks are selected because of their extensive use in the policy process literature, more so maybe than all other theories. The two either describe the policy process, the policy subsystem, or the behaviours of the actors within the policy process. However, they both lack any aspect of time or concepts that help define an order to the process, elements required by model simulations. For this reason, the policy cycle is also considered as part of the three main theories used. Despite its well-known limitations, it is a theory that defines steps which the policy process follows (Cairney, 2015). Each of the theories is detailed below.

2.4.1 The policy cycle

There are numerous theories of the policy cycle (Hogwood and Peters, 1982; Jann and Wegrich, 2007). The main goal behind these theories is to provide a simplified model of the policy process. The policy cycle provides a linear set of steps or stages that the policy process follows (Simmons et al., 1974). Over the years, the steps considered have changed depending on the authors and the studies. Versions of the policy cycle have varied from five to nine steps. Throughout these versions, several steps are always considered, even if under different names. These include the agenda-setting, the policy formulation, the policy implementation and the policy evaluation steps. An extensive set of policy cycles also consider a step termed as the decision-making step.

Following the most common policy cycle approach, it is assumed that the policy process begins with the agenda-setting step. In this step, the problems are defined by the actors and placed on an agenda. Then, policies meant to deal with the problem on the agenda are developed. Policymakers go on to decide which policy is best to address the agenda. If they agree on the policy, they implement this policy. The implementation step is often considered as a stand-alone step as policy implementation can be influenced by political, geographical, institutional or other factors (Baier et al., 1986; Howlett, 2018). Later, the actors go back to the policy to evaluate it and to see whether it needs to be changed, removed or kept in place (Nakamura, 1987). This is the policy evaluation step. As mentioned before, such an approach comes in very handy for a model of the policy process, as it requires a linear sequential process.

The policy cycle has been criticised in the literature (Cairney, 2015). The main criticism is aimed at its principal assumption: the linearity and sequentiality of the policy process. Most of the other theories assume that there is no order in the policy process. Further than that, most theories assume that different steps, as they are identified in the policy cycle, can happen at the same time or in different orders, painting the picture of a much more complex system. Studies have shown that, though one might be able to identify stages, in some cases, policy formulation comes before agenda-setting, or there is no agenda-setting step at all. This is all dependent on the context. Some of the criticisms also focus on the concept of stages itself. Critics argue that it is not possible to identify clear-cut stages within the process (Cairney, 2015).

2.4.2 The multiple streams framework

The Multiple Streams Framework (MSF) is a framework that attempts to explain how policies are chosen at a high level and under a condition of ambiguity (Herweg et al., 2015; Kingdon, 2014). The framework originates with the 'garbage can' model, in which it is assumed that the decision-making process has no beginning or end (Cohen et al., 1972). This model assumes that no single actor is in control of the policy process and argues that a problem often becomes of importance due to a focusing event.

The MSF follows the same assumptions. In it, the policy process is seen as a combination of three streams: policy, problem and politics. The policy stream is made of the policies that are present in the process, as developed by specialists. Similarly, the problem stream is filled with problems. The politics stream is filled with actors, including policy entrepreneurs. The assumption is that for policy to change, a window of opportunity has to be opened. For that to happen, the streams need to come together.

Within the MSF, it is the role of the policy entrepreneurs to bring the streams together. The framework assumes that policy entrepreneurs seize either a problem or a policy that they find essential. This importance is set based on several parameters. They will then work to bring the streams together by influencing other actors, reframing the problems or selecting different policies. Ultimately, it is their role to open windows of opportunity. It is also assumed that actors within the MSF have only a limited attention span and, as a result, can only focus on one problem or one policy at a time.

What happens after a window of opportunity has been opened is more opaque. The framework specifies that the policymakers have to choose whether to implement the policy used to open the window. They do this based on the problem that is on the agenda and if the political stream is in accordance. In effect, it says that the three streams must come together for policy to change. It also specifies that this is time sensitive as the policymakers are under strict time constraints, an aspect that can also play a role in their decision about whether to implement a policy or not.

2.4.3 The advocacy coalition framework

The Advocacy Coalition Framework (ACF) was first developed by Sabatier and was meant to deal with wicked problems (Sabatier, 1988). It is a framework that is structured around the concept of coalitions, actor beliefs and resources. The assumption is made that actors are part of a policy subsystem and that they can come together in coalitions to advance their interests. The coalitions are created by actors sharing similar beliefs. Through the coalitions, actors can create strategies to influence the decision-making process of other actors and, ultimately, to influence the policy process as a whole.

The ACF looks at the policy process on a decades-long scale. This is linked to the assumption that the beliefs of the actors are inherently stable and only change very slowly over time. It results in coalitions being stable groups that change only very little and whose membership remains mostly constant. The ACF considers that a broad range of actors can participate in coalitions. It includes policymakers and ranges all the way to interest group leaders. The ACF also highlights the role of the media and the judicial system within the policy subsystem.

One of the cornerstones of the ACF is the belief system. The framework argues that all actors have a tiered belief system composed of deep core beliefs, policy core beliefs and secondary beliefs. They act based on these beliefs. The deep core beliefs rarely ever change over a generation. They are often norms or values. The policy core beliefs are slightly more malleable and can change over periods of several years. They are directly related to the policy subsystem. At the bottom of the hierarchy are the secondary beliefs which are more likely to vary on shorter periods. The literature remains vague on what they consist of specifically. Some work has been done to try to detail the belief system further. For example, this includes the introduction of the notion of causal beliefs (Weible et al., 2004).

Resources are also an essential concept within the ACF. They are used by the actors to advance their interests. These resources are termed as political resources, and they can encompass an extensive range of resources, including financial resources. Mostly, they consider the ability of certain actors to mobilise their supporters for causes that they support. This, in turn, might influence other actors and affect their decision-making. Often the resources considered within cases in the literature are not quantifiable, an aspect that might make it more difficult to operationalise in the context of a simulation.

The ACF specifies four different avenues that can lead to policy change: policy-oriented learning, negotiated agreements, external shock or perturbations, and internal perturbations (Sabatier and Weible, 2007). Policy-oriented learning is the process of belief evolution that occurs over time as a result of actor and coalition interactions (Bandelow et al., 2017). Through learning, actors' beliefs change, and this can lead to a change in their policy preferences. This is a pathway to change that was first explored in Hall (1993), where Hall argued that social learning is required for policy paradigm shifts. Negotiated agreements are discussed among coalitions. They often happen if a hurting stalemate is present, something that brings the actor to the negotiation table. The outcome of such a negotiation will be a policy that can be directly implemented. An external shock can span a wide array of phenomena that might impact the policy subsystem. It can be a crisis, like a flooding event in a policy subsystem that deals with flood safety, for example. Finally, an internal perturbation relates to a shock that happens within the subsystem. It can be a rapid change of political resources or the beliefs of some of the actors. Such shocks have the potential to lead to a different policy being implemented later on.

3 The methodological approach

A broad methodological approach was selected for this thesis. This translates with the presentation of an approach where the researcher first builds a model of the policy process and, subsequently, simulates this model within the scope of a certain context. This is therefore a two step process. In the first step, for the creation of the policy process model, I propose a tool to make the creation of a model easier. This tool is referred to as a common language as it provides a set of requirements that are needed for any model of the policy process.

For the second step, the researcher will want to simulate the model s/he has created. For this, a context is needed and the researcher will have to select choose an environment model, either to test the policy process model s/he has just built or because the environment model itself is the subject of the study. To connect, and subsequently simulate, the two models, a hybrid simulation approach is presented in this chapter. This allows to systematically connect a policy process model built using the common language with an environment model selected by the researcher.

3.1 The common language

The common language is a tool meant to provide a set of key requirements for researchers that would like to build policy process models. It can be useful if the researchers lack knowledge on either the policy process theories or on modelling and simulation. These requirements are mostly extracted from elements in three policy process theories: the theory of the policy cycle, the multiple streams framework (MSF) and the advocacy coalition framework (ACF). The four requirements form the four building blocks of the common language. They are time, the policy arena, the actors and their environment, and the actor interactions.

3.1.1 Time

Time is the first building block considered for the common language because a process is being considered. This process consists of a series of actions or steps. Time introduces a sense of order without which it would not be possible to simulate the policy process. Modellers have to provide the computer with a sequence of actions that it needs to perform. To be able to do this, the researcher needs to first identify what happens during the policy process and in what order. Introducing the element of times forces the researcher to do so. Often so, this will translate in a set of steps that are informed from the literature and that can include an agenda-setting step, a policy formulation step, a policy implementation step and so on.

The concept of time is present, both explicitly and implicitly in the theories. The phase model, for example, explicitly provides a set of distinct decision-making stages which are sequential in nature (Teisman, 2000). It is also the case for the rounds' model which provides rounds instead of stages. They are less narrowly defined but still provide a sense of time. This is also the case for the theory of the policy cycle. It specifically mentions a sequence of steps that the process follows. This often includes, but is not limited to, agenda-setting, policy formulation, policy implementation and policy evaluation as outlined in the literature review of Chapter 2 (Simmons et al., 1974). Kingdon has said that the policymaking process can be conceptualised as a set of processes with a predefined number of steps, not too dissimilar from the ones of the policy cycle (Kingdon, 2014).

Even the MSF makes a special mention of the temporal order (Zahariadis, 2007). Rüb (2016) outlines that agenda setting and policymaking are "shaped by the temporal rhythms of politics". The flow of the three streams can in itself be a representation of flowing time where different polities, problems or actors rise and fall depending on the time and the circumstances. Time is also a prominent factor for the window of opportunity concept. The outcome of the policy process can depend on the timing of a window of opportunity and on how long that window remains open. At the same time, in the MSF, time is not considered linearly. Instead, the concepts can be seen as timeless, as they can happen at any time and at the same time. The MSF often limits itself to the agenda-setting step as is seen in the literature (Mazarr, 2007; Rüb, 2016). By definition, choosing the agenda-setting step implicitly assumes that there are different steps within the policy process. The MSF literature does not seem to address that fact.

The ACF does not consider time explicitly (Sabatier and Weible, 2007). It assumes that the steps in the policy process happen simultaneously or simply does not consider any steps. However, this does not mean that there is no element of time within the framework. The concepts of external events or perturbations are inherently time based. These are key concepts theorised as ways to induce policy change. They are events that happen at very specific time and that can last a certain duration. Their timing can influence the course of the policy process, and can lead to policy change. As an example, Nohrstedt (2009) outlined out how nuclear accidents, which are external perturbations, have influenced the coalitions active in Sweden's nuclear energy policy.

3.1.2 The policy arena

The second building block is the policy arena. A policy arena is considered because of the need to limit the scope of any study and any model. One cannot simulate the entire world but has to consider a bounded system. The policy process as a whole is highly complex and can span a large number of fields, disciplines and actors. This can make for a daunting challenge when studying any specific case. By having to define a policy arena, the researcher will have to define a set of boundaries. This arena will often come naturally and be related to a specific policy system. However, in some cases and depending on the nature of the study, there might be a need to consider more than one policy system. This is a choice that the researcher will have to make. Once selected, the policy arena will help limit scope of the study, hence limiting the actors that need to be considered or their beliefs.

The policy arena is brought up in the ACF literature through the policy subsystem concept. The framework assumes that actors will specialise within a specific policy subsystem. It then characterises the policy subsystem as having "both functional/-substantive dimension and a territorial one" (Sabatier and Weible, 2007). The ACF's policy subsystem, hence, helps narrow down the actors that need to be taken into account in a study, along with the selection of a specific topic and a specific location or jurisdiction.

The MSF does not provide the same explicit theoretical boundaries to its framework. However, in practice, studies that use the MSF for their analysis are always limited to a specific field or policy. It is done by limiting the study to an "arena", as explained in Durant and Diehl (1989) for an example in foreign policy. In other studies, this is done by focusing on a particular policy, explaining how this policy came to be, and how it was adopted using the MSF. Such focus helps narrow down the study and can effectively be considered as the selection of a policy arena.

Setting boundaries through the use of a policy arena is also important modelling-wise and should be done before any model building exercise. Boundaries help with model creep, lower the number of elements that need to be considered and ultimately simplify the analysis of the results. This leads to a greater understanding of the results obtained.

3.1.3 The actors and their environment

The actors and their environment are the third building block considered for this common language. The actors are indispensable to the policy process. They make decisions that affect the evolution of the policy process, based on their preferences and beliefs. The way these decisions are taken are conceptualised differently depending on the theory considered. At their core, these are decisions as to whether a specific policy should be implemented or not. Models of the policy process developed by researchers are free to make the decision-making process as complex as they would like, depending on the goal of their research.

Considering actors is broad. In most of the theories, roles are defined for the actors. In the common language, it is decided not to specify any role for the actors, beyond the role of the policymaker. It is assume that the policymaker is the only actor that is capable of direct policy change through the selection of a policy. The selection of additional roles is a task that is left to researchers building the policy process models. The roles selected should consider the case being addressed and fit within the scope of the study. It is of course advised to keep any model of the policy process as simple as possible without introducing unnecessary additional roles. In the present case, it is assumed that the roles are differentiated from one another through the attributes and the actions that they are provided with.

The actors are more or less prominent, depending on the theory considered. Regardless, they are always present. One role that is also always present within all theories is the role of the policymaker for reasons outlined above. Beyond this role, the theories vary widely with the actor roles that they explicitly consider and the capabilities of the actors.

In the MSF, there is a strong focus on policy entrepreneurs (Zahariadis, 2007). It has even led to a new branch of research specifically focused on policy entrepreneurship (Mintrom and Vergari, 1996; Mintrom and Norman, 2009). The policy entrepreneur is a role that actors can take. In this role, the actors will attempt to further their respective interests. They do that by opening windows of opportunities. Beyond policy entrepreneurs and policymakers, the MSF fails to mention many other roles. It reflects a vision where, if an actor is not a policymaker, and it wants to act to facilitate policy change, it must turn into a policy entrepreneur.

The ACF follows a different path. It acknowledges many actors: "The set of policy participants includes not only the traditional 'iron triangle' of legislators, agency officials, and interest group leaders, but also researchers and journalists who specialise in that policy area [...] and judicial officials who regularly intervene in a policy subsystem" (Sabatier and Weible, 2007). The framework makes a special mention of the role that researchers and scientists can have in influencing the beliefs of other actors. When the ACF identify these actors or these roles, it does not necessarily link them to specific attributes or actions, something that would be necessary to model. Therefore, modelling them would require additional assumptions that define what the differences between them are.

The ACF research focuses intensely on two main concepts: coalitions and policy brokers (Sabatier and Weible, 2007). Coalitions are groups of actors with like-minded policy core beliefs. They are also approached as actors themselves through their goals and actions. Actors that are part of coalitions lose their agency in the broader policy subsystem. It is the coalition that will advance their interests, though each actor can also influence the coalitions' direction from the inside. Policy brokers search for stability and help mediate between opponents, within the policy subsystem, with the goal of policy change (Ingold and Varone, 2011). Though notable within the ACF, coalitions and policy brokers are not included within the common language, as they are not requirements to emulate the policy process. They are however elements that can be included in models of the policy process.

The belief system

The ACF is also the only framework that goes beyond specifying roles for actors. It also details the actors' decision-making process. They do so using a belief system. According to the framework, the actors have a three-tiered hierarchical belief system, as was detailed before (Sabatier and Weible, 2007). It includes deep core, policy core, and secondary beliefs. Other theories approach the decision-process of actors through utility such as the game theoretic models developed by Bueno de Mesquita or Stokman (Bueno de Mesquita, 2011; Stokman et al., 2000). The Values-Beliefs-Norms (VBN) theory has also been mentioned in the literature, stemming from social psychology (Henry and Dietz, 2012). In the modelling and simulation literature, the Belief-Desire-Intention (BDI) paradigm has also been used for the simulation of decision-making of agents (Taillandier et al., 2012; Singh et al., 2016). However, it is a paradigm that is far from the theories of the policy process.

For the common language, I advocate for using a decision-making framework that is belief-based and hence follows the one presented in the ACF. Following a belief-based approach has not been done before and can enhance the knowledge that can be gained from the theories. This does not mean this is the correct approach, it is merely one of many. This is also an approach that is well suited for agent-based modelling, the modelling approach selected for this work. Finally, in the way it is conceptualised, it is the most versatile approach possible when considering the coupling of a policy process model with an environment model.

The literature on the ACF belief system informs the belief system that is conceptualised for a model of the policy process. The goal of this belief system is to facilitate the decision-making progress of actors at the individual level. However, the literature is both limited when it comes to operational details related to the belief system and it is also inconsistent (Sabatier and Weible, 2007). These inconsistencies are both present when comparing the theory and the empirical studies, but also between empirical studies themselves. To model the decision-making process of the actors, there is a need for a coherent and detailed belief system. This is obtained through a mix method using both assumptions and hypothesis from the literature but also introducing assumptions where the literature is lacking or where it is inconsistent.

The central element of the ACF belief system is the belief. Therefore, it is important to define what a belief is and how it is collected. This is important for several reasons. For the purpose of modelling, there will be a need to provide researchers that would like to use the common language a clear way of identifying beliefs in their cases. One needs to keep in mind that these beliefs will need to be directly associated with quantifiable parameters within the environment model - the model providing context to the policy process. Therefore, there is also the need to provide a a way to quantify these beliefs so that they can be used within the decision-making process of the actors.

Throughout the literature, the definition of a belief varies, along with the way it is recorded or collected and analysed. This is most visible with the policy core beliefs. The definition of the deep core belief is consistent throughout the literature and so its application. These are normative beliefs that encompass all policy subsystems (Sabatier, 1987). Empirically, studies have applied this definition consistently. For example in Elgin and Weible (2013), deep core beliefs are identified as fiscal and social beliefs, while in Weible et al. (2004), a study related to the use of science, the deep core beliefs are identified as pro-collaboration and pro-scientific management. More similar examples can be found throughout the literature (Matti and Sandström, 2011; Pierce, 2011; Weible et al., 2004).

The secondary beliefs are barely considered in the literature beyond their definition in Sabatier (1987). In a majority of cases, they are not collected at all. They are defined in Sabatier (1987) as being "instrumental decisions and information searches necessary to implement policy core" beliefs, and as being "specific to policy arena/subsystem of interest". One study that considers secondary beliefs is Weible et al. (2004). It identifies these beliefs as being secondary as opposed to the other beliefs because they are "limited substantively" to the "policy subsystem" of the case studied.

More complications arise when looking at the policy core beliefs, the most studied belief tier throughout the literature. The definition of the policy core beliefs states that they should be "fundamental policy positions concerning the basic strategies for achieving normative axions of deep core" (Sabatier, 1987). This implies that the policy core beliefs are policy related. The literature and subsequent studies have shown that this is not always the way this definition is used. In fact, the theory varies from the empirical studies. On the theoretical front, Sotirov and Memmler (2012) states that eleven policy core belief topics have been proposed, including "groups whose welfare is at risk, commodity vs. amenity values, state regulation vs. market self- regulation, centralisation vs. decentralisation of authority, policy instruments preferences, and policy core policy preferences".

Looking at empirical studies, these beliefs were not widespread. In Elgin and Weible (2013), most of the beliefs are formalised as normative statements assessing a number of policy options. One belief is provided as "human behaviour is the principal cause" of climate change. This is the only belief that is not related to a policy option. The goal behind these beliefs was mostly to define the coalitions. In a number of other works, the policy core belief are often presented in the form of statements. These are not necessarily related to policies but instead can be linked to the current state of affairs (Kukkonen et al., 2017; Markard et al., 2016). in Dela Santa (2013), the number of beliefs identified are limited to three, one per coalition identified. Some studies, like in Pierce (2011), not only are the beliefs found but also components to these beliefs which are used to "operationalise each belief". Overall, the trend throughout these studies is to find beliefs through the identification of coalitions and not necessarily as a result of the original definitions provided in the framework.

In several studies, the policy core beliefs are belief thats consider or are linked to the goals of the actors. This is the closest the policy core beliefs come to considering the future, or at least what the actors would like to see happen in the future. This is the case in Brugger (2017) where beliefs related to the transition goals are collected. The literature remains inconsistent on how it treats these goals on whether they should be full blown beliefs within the belief system or whether they are elements that should be considered separately. One such example is Dela Santa (2013) where the lines between beliefs, goals and policies are blurred. Two beliefs can be seen as either policies or goals: "customary practice of hunting manta ray should be restored" and "whaling as a way of life should be continued". The third belief however is formulated as a statement "species conversation could benefit communities".

And then one can also consider the ways the beliefs are collected and/or quantified throughout the literature. This can help set requirements for the way the belief system should be implemented in a model. Unfortunately, there is no consistent method used for this. Some studies use likert type scales (Brugger, 2017; Herron and Jenkins-Smith, 2014; Markard et al., 2016; Matti and Sandström, 2011), others have used yes or no answers (Kukkonen et al., 2017), while some studies have used more complex approach differentiating categories (yes, no, neutral or has no knowledge) (Pierce, 2011). Finally, beliefs can also be inferred from transcripts and document analysis as was done in Albright (2011). The variation in approaches is linked to the variation in methods. Likert scales are preferred for surveys while interviews are used for the other methods.

One additional category of beliefs that is rarely explicitly discussed but that is often implied is the causal belief. Causal beliefs are implied throughout the definition of the belief system because both the secondary beliefs and the policy core beliefs are defined with respect to their higher tiers in the belief system (Sabatier, 1987). The secondary beliefs is defined such that is "implement the policy core" beliefs while the policy core beliefs are defined such that they "achieve [..] normative axioms of deep core" beliefs. These definition imply a link between the tiers of the belief system. One study has looked at the causal beliefs explicitly. Weible et al. (2004) goes as far as quantifying these causal beliefs using path analysis and looking at how beliefs on higher tiers of the belief system help explain the beliefs of the lower tiers.

All of these elements from the literature can be taken into account when providing an approach for the modelling of a belief system. A number of literature based assumptions are provided below:

1. The belief system should have three tiers - secondary, policy core and deep core.
2. The beliefs of the actors should be an integral part of the belief system - though a definition of what a belief is needs to be specified.
3. The deep core beliefs will span all subsystems while policy core beliefs only span one policy subsystem at a time.
4. Causal beliefs can be considered to link the different tiers of the belief system.
5. Beliefs can also consider the goals of the actors, hence considering the future.

6. it is assumed that actors have cognitive limitations and hence consider only one item at a time. This is an assumption that stems from the punctuated equilibrium theory (PET) (Baumgartner et al., 2014).
7. The formal model of the belief system should consists of elements that can all be quantified.

There are a number of additional assumptions present throughout the literature such as the devil shift, the speed at which beliefs can change, the impact of venues and resources or the salience of beliefs (Fischer et al., 2016; Nohrstedt, 2011). However, these were not used to create the belief system. This is because the goal of this belief system is to be as simple as possible and not a part of the model that can reconcile all hypotheses and assumptions present in the literature. Such elements can always be added in subsequence work. For example, a study on the devil shift could introduce it within the model to be studied in more details.

Based on the assumptions from the literature provided above, a belief system that can be used in a model of the policy process is formalised. The following assumptions are made for this belief system:

- The belief system is composed of three tiers - secondary, policy core and deep core.
- Each tier of the belief system is composed of issues, directly related to parameters in the environment model - the context within which the policy subsystem is anchored.
- Each issue is defined by three constituting parameters:
  1. The actual belief: this is the perception that the actor has of the issue at one moment in time based on information from the environment model.
  2. The preferred state: this corresponds to where the actor would like to see the issue be in the future. In effect, it refers to the goal of the actor for the issue.
  3. The preference: The preference is a calculated parameter that is used to define which issue the actor will seize on on each tier. It is assumed that each actor can only consider one issue at a time. It is calculated as the difference between the actual belief and the preferred state. On each tier of the belief system, the actor will seize on one issue, the one with the highest preference.

- The issues on different tiers of the belief system are linked together through causal belief. It can be understood as to how an actor understands the impact of the change of an issue's actual belief on the actual belief of issues in a higher tier of the belief system. As an illustration, a causal belief would specify what would happen to the state of the economy actual belief if the unemployment rate actual belief increased.

These assumptions all relate to the assumptions that are present in the literature, as they were presented above. They also introduce new concepts necessary for the modelling and simulation of the policy process. All elements in this belief system can be quantified based on the parameters that are present in the environment model. This will be further demonstrated within subsequent sections of this chapter and subsequent chapters.

3.1.4 Actor interactions

The last building block of the common language relates to actor interactions. Two types of interactions are distinguished: interactions between actors and the environment, and interactions between actors. The interactions between actors and the environment are mostly passive actions. The actors' actual beliefs are informed by what they observe in the environment - the information they obtain from the environment model. On the other hand, actors also influence the environment by implementing new policies.

Actor on actor interactions are left open ended in the common language, none are therefore specified specifically. Various implementations can be thought based on the wealth of options in the literature. This is enable by the coherent architecture of the belief system presented in the previous section. The goal of this common language is not to specify one implementation but only to foster the creativity of modellers and researchers to come up with different implementations based on the literature and to test them. Furthermore, actor-on-actor interactions are not strictly requirements for the policy process simulation; only the passive actor-environment actions are.

In this section, I propose one example of actor-on-actor interactions as an illustration of what can be conceived using the common language. This example is further expanded on in Chapter 9 for the demonstration of the common language as an ACF implementation.

To advance their interests, actors can influence others on their belief systems. This includes influences on the actual beliefs, the preferred states and the causal beliefs of other actors. The influence effectiveness of actors can take into account several important factors based on the literature. For example, an individual actor is more likely to engage with another actor if they both share an average level of conflict. If it is too high, they will not listen to one another, and if it is too low that is because they probably already share the same belief making that interaction pointless (Henry, 2011). Furthermore, the interactions are more likely to be positively received by actors of the same political affiliation or whom they trust (Calanni et al., 2014). Power difference and venue location can also have an impact on the interactions between actors (Cairney, 2006; Nohrstedt, 2011).

The theories offer a wide variety, or lack thereof, of descriptions on how actors interact. Details are often scarce on the interactions themselves and how they happen specifically. However, they do provide valuable insights into what affects these interactions. The interactions between the actors can be seen as forms of lobbying, influence argumentation or coordination (Huitema and Meijerink, 2010; Sabatier and Weible, 2007; Pierce, 2011). These efforts are often mentioned only within the context of policy entrepreneurs and with the subtext that not all actors engage in such behaviours. This is prominent within the MSF. The common language allows interactions for all actor roles and lets researchers decide.

The ACF seems to ignore interactions between actors explicitly. It does specify pathways that can lead to policy change, with most of them relating to a change in the actors' beliefs (Sabatier and Weible, 2007). This can implicitly be tied to interactions of some sort. Moreover, though the ACF does not necessarily give a name to these interactions, it provides an outline of their impact. This happens mostly through the study of the coalition, coalition coordination and the social networks that are at the origin of these coalitions (Calanni et al., 2014; Matti and Sandström, 2011; Weible et al., 2018). The goal of the common language is to provide a way to emulate such impact. By influencing each other on their beliefs, actors can emulate the policy learning effect that is described as a pathway to policy change. The interaction with the environment is also a way to include the external shocks that are described as a second pathway to policy change. And the common language also allows for the inclusion of coalition within policy process models, opening the door to a large swath on the literature.

The ACF provides an additional dimension to actor interactions: resources (Nohrstedt, 2011). Interactions between actors are not free of costs, may that cost be monetary or time-related. It takes time for actors to influence one another and, as is shown throughout the literature, it can also cost much money. The ACF introduces the concept of resources. As mentioned before, these resources include formal legal authority to make policy decisions, public opinion, information, mobilisable troops, financial resources and skilful leadership. These resources can be used to limit the number of interactions that actors can perform.

3.2 Hybrid simulation

A policy process model cannot be simulated without its context. To account for this, after the construction of a policy process model, researchers will have to select an environment model, to provide the policy process context. To connect the two models together, a hybrid modelling approach is proposed. This approach is one that is systematic as it can be used for a wide variety of models. It relies on obtaining a set of inputs and outputs from the policy process model that is detailed later on.

3.2.1 Hybrid simulation in the literature

Hybrid simulation models are simulations that consist of at least two different modelling methods. In most cases, these methods are either agent-based, system dynamics or discrete event models (Brailsford et al., 2010; Onggo, 2014; Schieritz and Grobler, 2003). Such hybrid models always require three parts: the modules each usually containing a different modelling method, the module interface used to connect them, and the updating rule that helps define how one model impacts the other. This is shown in Onggo (2014) with the example of a blood supply chain in the health sector.

Hybrid models can be classified in different categories (Nava Guerrero et al., 2016; Swinerd and McNaught, 2012; Vincenot et al., 2011). These classifications can vary in the literature depending on the modelling methods considered. One classification often cited is the one given by Swinerd and McNaught (2012) which differentiates between three main classes of hybrid architectures:

1. The sequential class where each module is used as the input for the next module. The modules do not run at the same time.
2. The interfaced class where there is a non-sequential combination of modules.
3. The integrated class where modules and model outcomes provide feedback to one another.

This classification mainly focused on the hybridisation of agent-based and system dynamics models (Kim and Juhn, 1997). Lättilä et al. (2010) goes further by comparing the two modelling methods and arguing why they should be used more together in hybrid models. A similar classification can be found using discrete event simulation and system dynamics in Chahal and Eldabi (2008). Vincenot et al. (2011) provides a slightly different classification containing four different categories, also for agent-based and system dynamics models. The overlap between these classifications is significant.

Hybrid modelling is widely used in a range of disciplines. For example, in Mayer et al. (2006), it was used to simulate agriculture land use. This study contains an agent-based model coupled to a landscape and climate model through the use of an interaction module. This approach allows for easy modifications in both models. In Barton et al. (2010), a stochastic model simulating land use is coupled to an agent-based model to simulate long-term dynamics of Socio-Ecological Systems (SES). Such an approach is relatively common for SES as many custom-made modelling approaches were built to study these systems. To study their interactions with actors, agent-based models are often used, leading to the need for hybrid models (Axtell et al., 2002).

Hybrid modelling is also used for health-related topics (Brailsford et al., 2010; Chahal and Eldabi, 2008; Onggo, 2014). In most of these cases, they are combinations of system dynamics models and discrete modelling simulations. The discrete event model describes the hospital processes which look into the processing of the patients. The system dynamics model focuses on how people transmit diseases to each other on the broader world using a combination of feedback loops.

Martinez-Moyano et al. (2007) used hybrid modelling to simulate financial stability with a combination of an agent-based and a system dynamics model. These two modelling methods were interfaced through a data manager built in Java. Within this hybrid model, the system dynamics model was controlling the overall model. The system dynamics model simulated different financial layers. Akkermans (2001) looked into emergent supply networks. He presented an exploratory study of supply chain networks using a combination of agent-based and system dynamics models for the supply chain.

Despite the numerous models outlined in this section, hybrid modelling remains an exception in the field of the modelling and not the norm. Most studies performed on socio-technical system use one modelling approach, may it be system dynamics, agent-based or any other method. Furthermore, in all cases, to the author's knowledge, these hybrid models are tailored to a specific system. The present work attempts to break this trend by presenting a policy emergence model that can be used with an extensive range of socio-technical models.

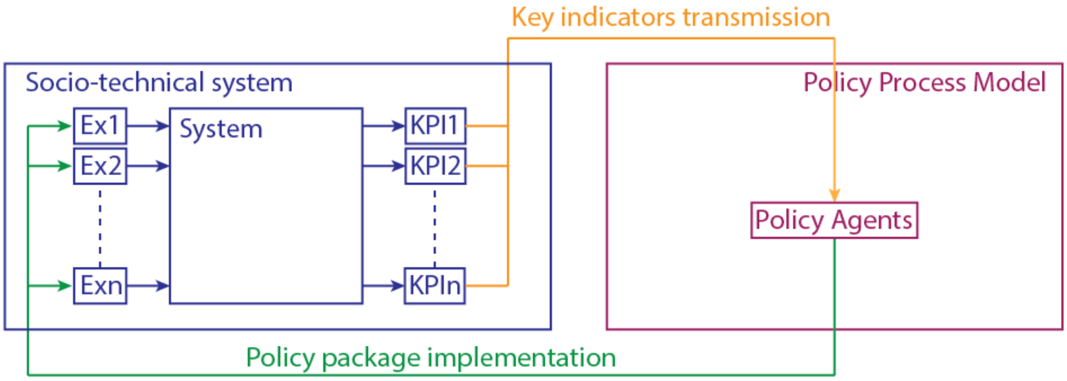

Figure 3.1 – Architecture of the hybrid model.

3.2.2 Endogenising the policy process

The approach presented is one where an agent-based model of the policy process is built in such a way that it can be used in a sequential hybrid model. The environment model - often a socio-technical system model - considered can be either an agent-based model or a system dynamics model as long as it can communicate the necessary data to the policy process model built in the open-source coding language Python. Environment models should also have the ability to use data from the policy process as input throughout the simulation.

The sequential hybrid model is drawn in Figure 3.1. With this hybrid model, the policy process model requires inputs from the socio-technical system to inform the policy agents. These inputs inform the belief system of the agents, updating their actual beliefs with what is happening in the system every time the policy process begins. The output from the policy process model is a policy instrument or policy package. The agents decide this output based on their deliberations. For the socio-technical system, the inputs correspond to changes in the exogenous parameters. These are based on the policy that has been selected by the policy agents. Its outputs are key performance indicators that are selected specially to inform the belief system of the agents.

Using this sequential approach, each of the models runs independently of the other and only communicate when needed. This is important as the policy process is not simulated for every time step of the socio-technical system. Depending on the experiment being run, the policy process might be triggered once every few months or years. This approach allows for such a sequence of steps. In the model demonstrations presented later in this dissertation, the policy process model is only run a fraction of time compared to the number of steps that the socio-technical systems go through.

This approach was also selected because it is the most modular and versatile possible. It provides the most significant access to the policy process model by making it simple to integrate into almost any socio-technical system model. All that is needed is a model that can export data and read data throughout its simulation. If that is possible, then theoretically, the policy process model can be added to perform additional studies that relate to policymaking and the impact of dynamic policies on the socio-technical system. This will be demonstrated in the next two parts of the dissertation with different examples of socio-technical systems.

3.3 Discussion

The approach presented in this chapter suffers from a number of limitations. This is both the case at the level of the common language and the level of the modelling and simulation. The common language presented is limited by design. It is composed of the minimum set of building blocks needed to simulate the policy process. This means that many concepts from the theories, even prominent concepts, have not been included. This limited approach is intentional and plays into the goal of the common language: providing a set of minimum requirements for the building of policy process models. The common language is meant to be a stepping stone for modellers that have only limited knowledge of the literature studying the policy process theories. The limitations related to the modelling approach are broader and are limitations that are often present in models. All are discussed in this section.

3.3.1 The policy steps

The use of simulation forces the common language, and hence the policy process to be approached linearly with a set of steps. The main reasons were described previously. This is an assumption that is made solely for simulation purposes. In reality, the policy process is often not linear (Jann and Wegrich, 2007). It is often much more complex with actors interacting at the same time in different locations. Often, the steps that are specified in the theory of the policy cycle are not followed. Sometimes, there will be no need for an agenda. Other times, the agenda will be created as an excuse to bring up a policy for a vote. It is argued here that, although we cannot reproduce such events, in most cases a linear process will not detract from the outcome.

Furthermore, even in qualitative studies, the chaos of the policy process often needs to be sorted through for researchers to gain an understanding. By looking at the events in a linear way, one can much more easily understand what has happened and why. The common language forces this analysis lens onto the researchers and the modellers.

The selection of the steps and their number is left to the modellers.

3.3.2 The belief system

The common language has been mostly articulated around the ACF belief system. It is a lynchpin of the potential interactions as they are described, requiring actors to influence one another on their belief system. It is done because of the lack of a systematic and well-articulated equivalent alternative approach to the decision making process throughout the literature. Despite - or because of - the use of the ACF belief system architecture, many theories can be simulated. It includes the potential for simulating the MSF, which does not prescribe a system to model the actor decision-making process. Other theories, such as the narrative policy framework (McBeth et al., 2007), can also be emulated using the ACF belief system approach. For example, using that framework, a change of narrative can be seen as a change of understanding of the environment by the actors, something that can be modelled as a change of causal beliefs.

Despite the use of the ACF belief system, other approaches are not prescribed. There might be alternatives that have been overlooked considering the vast policy sciences literature. The primary threshold criterion that these approaches should meet is the ability to simulate the actors' decision-making process and they should be able to inform the actors based on information from the environment model, whichever form that might take. If that is possible, then there is no reason why such an approach could not be used.

Alternative considering either a utility based approach could be considered at a lower level (Bueno de Mesquita, 2011; Stokman et al., 2000). This could be complementary of the belief system considered, where the utility approach would be used when actors interact to determine which interactions they perform. It also has the potential to replace the entire belief system approach considered as long as utilities can be calculated for all of the parameters present in the environment context. Note that the way the policies are assessed and implemented is left open-ended in this common language, opening the way to different approaches, including ones based on the utility approach. Similar considerations apply to the VBN theory and the BDI paradigms. They could either be made complementary to the approach presented here or they could replace it (Henry and Dietz, 2012; Taillandier et al., 2012).

3.3.3 The policy entrepreneurs

The common language does not define roles for the actors. It only specifies the need for actors in general. The reason for this approach is to avoid pre-emptively selecting the roles that should be included in models built using the common language, except for the policymaker role of course. The common language being a tool, this should be the choice of the modeller and be dependent on the case being studied. Different theories consider different roles as more or less critical, and the addition of specific role might have limited the choices possibilities for modellers later on. That said, an architecture has been put into place for the actor decision-making process. Therefore, it is simple to introduce different roles with different strategies. This can include the role of policy entrepreneur for modellers interested in simulating the MSF, for example (Mintrom and Vergari, 1996; Zahariadis, 2003).

3.3.4 The streams

The concept of streams also seems to be entirely absent from the common language. This is not the case. By default, the politics and the problem streams are present within the common language (Zahariadis, 2003). The actors can be seen as illustrating the politics stream. The problem stream is inherently linked to the belief system. Each issue in the actors' belief system can be seen as a problem or potential problem preference-wise. The policy stream is missing from the common language, but it could be added in a model based on the common language if one were to study the MSF specifically. It could be done by allowing the actors to choose both a problem and a policy at specific points in time along the process, similarly to what the MSF conceptualises. These points in time would then be considered to be windows of opportunity. The common language does not forbid the use of streams and can facilitate it.

3.3.5 The coalitions

Coalitions are at the centre of the ACF and are often seen as facilitating policy change (Sabatier, 1987). Nevertheless, they are not included in the common language. Once again, this is because the common language only provides a set of minimal building blocks required for every model of the policy process. Coalitions are not a required element of the policy process.

However, and similarly to policy entrepreneurs and the streams for the MSF, coalitions can be added to any model built using the common language. For example, coalitions can be modelled as super-actors. Coalitions are made of actors present in the policy arena and have their networks and belief system. They can also have resources, pooled from their members. Coalitions can also use the same interaction rules like the one provided for the actors, or they can use more advanced ones. Finally, based on the literature, coalitions could be provided strategies such that they best advance their interests. This approached is further developed in Chapter 9.

3.3.6 The common language potential

Throughout this discussion, it has been outlined that models built with the four building blocks of the common language can be used to simulate a wide variety of theories or cases. The common language enables researchers to test a plethora of hypotheses through the modelling and simulation of a large variety of policy process models. And this can be done for different purposes, may it be for the exploration of the policy process theories or to better understand the effects of the policy process on an environment model.

For example, using a policy process model built using the common language, one could try to understand the impact of the electorate on the selection of policies by policy actors and ultimately their impact on policy change. A model that includes public opinion and its interactions with the policy actors would be needed, along with the potential for additional policy process steps depending on the jurisdiction considered. Based on hypotheses formulated by the modeller, such a model can take a variety of forms. Simulations and their results can then be used to assess whether these hypotheses are appropriate or whether they need to be changed.

The common language can also be used to simulate theories and test hypotheses and assumptions that are present in the literature. For example, by using the four building blocks plus the necessary theory-specific concepts, it is possible to build a model that can simulate the MSF. Such a model will need to introduce the three streams. This model might consider that two of the three streams can already be inferred from the building blocks of the common language as it has already been detailed. Only the policy stream would then be required additional to the elements present in the common language. A model of the MSF would also need to put an emphasis on the policy entrepreneur role and specify in detail their attributes and abilities.

The same can be said if the common language is used to build a model to simulate the ACF. The modeller will have to focus on the introduction of coalitions and clearly define the assumptions that structure their decision-making process. This is something that is addressed in the third part of this dissertation.

3.3.7 Model complexity

Some limitations arise when considering modelling. They all relate, in one way or another, to the modeller's choice for the model complexity. A modeller has to choose the boundaries of the model carefully. The entire world cannot be modelled within one single model. At the same time, setting boundaries will introduce limitations into the model. In the present case, this includes a loss of crucial dynamics, a lack of true policy emergence and several cross-subsystem behaviours.

The simplicity of the common language, as it was described in a previous section, compounds with difficulties that researchers might have in fitting their specific cases to it. Because of its simplicity, fitting any case will require several simplifying assumptions. This can lead to a loss of crucial behaviours and results in a limited gain an understanding of the system. This will be the case with Switzerland in Chapter 7 and Chapter 12. If done improperly, this can limit the usefulness of the model in answering focused research questions. In such cases, there might be a need to further detail the policy process model before fitting the case and simulating it. Alternatively, in some cases, it might be better not to use the language proposed in this dissertation, and either use a different approach or forgo the use of modelling entirely.

Another point of contention that stems from modelling and the need for boundaries is the lack of true policy emergence. In contrast to many agent-based models that impose policy as an external variable, models built using the common language allow agents in the model to make their own decision about agenda setting, problem definition and choice of solutions. Their decisions are based on a unique and evolving understanding of reality, and on interaction with other agents who have their own understanding and preferences. However, the problems, the solutions (policies), the causal relations that dictate agents' preliminary understanding of their environment, and the agents' initial goals are all determined by the modeller. They are intended to reflect actual actors' belief structures at a certain point in time. Actors, within the confine of the model, cannot think up new policy solutions, introduce new issues or find new ways to define problems beyond the initial input and choices made by the modeller. Emergence is therefore restricted to the choices the agents make. It is based on their changing understanding of the environment and of their evolving preferences.

The use of a model, because of its boundaries, also prevents the study of cross-subsystem interactions. The ACF specifies that most processes happen within a specific policy subsystem (Sabatier, 1987). However, in reality, this does not have to be the case. The actions of actors in a policy subsystem can have an impact on the outcome of another policy subsystem. Such behaviours, though not taken into account within the theories, are often mentioned within qualitative cases as having had a role in the policy change. The common language prevents cross-subsystem interactions. One way to deal with such interactions would be through the use of an external event to which actors will react. Another would be to mention cross-subsystem behaviours within the discussion of the results of the model.

3.3.8 Data requirements

In addition to the data needed to simulate the environment model, the amount of data needed to simulate the policy process is likely to be a hurdle in the usability of the common language and the subsequent hybrid models created from it. The modellers will need to gather data on the beliefs of the actors, their policy network, and their resources. In most cases, this data can be obtained, though with difficulty. In some cases, the data might not be obtainable at all and will have to be estimated. This limits the usefulness of the common language for focused cases in the vein of current qualitative studies.

Data requirements are less of an issue when considering other applications of the common language. In cases where theory exploration is favoured, or only to gain a better understanding of the behaviours of actors is needed, for example, precise data would not be necessarily crucial. Instead, there is only a need to have data that is the right order of magnitude.

All these issues are further discussed later in this dissertation.

3.3.9 Model validation

Model validation is an issue for social systems. This does not mean that social system models cannot or should not be validated. Some form of validation for a model built using the elements of the common language is possible and needed. The validation, in this case, needs to be performed to check whether the model built is one that conforms with the theories. This includes the validation of the temporal narrative, the roles of the actors and the actors or/and environment interactions as these are part of the common language. It also includes validating elements that would be added for any specific model.

Different methods exist to validate a policy process model. The one favoured here is the validation by experts. For this, the common language, and models built using it, are shown to experts, often the same researchers that have developed the theories in the first place. They then assess whether the models presented corresponds to what they intended or understood the theories to be. Any feedback received is then incorporated into the model that is built. For focused cases, historical validation can also be used. This consists of checking the results from the model with outcomes in real life.

The validation can also be seen as an integral part of theory exploration. Theory exploration, by definition, requires testing elements introduced in the model that cannot yet be validated. The aim is then to show that such elements, and the resulting behaviours, are coherent with behaviours described in the literature or observed in real life.

3.4 Conclusions

This chapter has presented the methodological approach that is used throughout this dissertation. This included the presentation of a tool build to help researchers design policy process models: the common language. With the presentation of its four building blocks, I have shown how this common language was devised and hints at how it can be used to build policy process models. The outline of the hybrid modelling approach has further reinforced how researchers can simulate the policy process with environment models once a policy process has been formalised using the common language.

The rest of the dissertation focuses on demonstrating how the common language can be used through the presentation of two models of the policy process. Each of these models are then tested with different environment models either to outline the theory exploration possibilities or the understanding that can be gained through the endogenisation of the policy process in a complex socio-technical model.

The simplest implementation model Part II

4 Conceptualisation

The common language is first used to build a model referred to as the simplest implementation model. This implementation is used as the first common language use, and as the first demonstration of a policy process model. The simplest implementation model is meant to be the simplest model that can be built using the common language, and that can emulate the policy process. To construct this model, Occam's razor approach is used and all of the model's possible complexity is removed. This model is therefore not meant to represent a theory in particular. Instead, it is only used to explore the potential of a model of the policy process.

In this chapter, I present the conceptualisation of this simplest implementation model built using the common language.

4.1 Time

A policy process with one step is considered for the simplest implementation model. This step is the policy formulation step, and it consists of the selection of a policy instrument by the policy actors. They select the instrument based on its expected impact on their secondary issues. If a majority of actors have chosen the same instrument, the actors will be implementing this policy instrument in the environment.

A one step process is selected because it is the simplest possible option. Furthermore, the use of only one step does not detract from the fact that the model will still able to simulate the policy process where policies can be selected by actors, and these policies will then go on to influence an environment model. This limited complexity will also limit the results that can be drawn from this model.

4.2 The policy arena

The policy arena is limited to the environment model that is considered for each simulation. For this dissertation, as will be shown in subsequent chapters, this either consists of a predation system, the Swiss electricity market system or a flood safety system.

4.3 The actors

The only actor role considered for the simplest implementation model is the role of policymaker. Each policymaker can select a policy core issue, a secondary issue, and a policy instrument based on their belief system. Collectively, policymakers can decide to implement a policy instrument.

The choice of selecting only one actor role is motivated by the need to limit complexity. To simulate the policy process, as it was outlined in the common language, there is no need to have more than policymakers. They are needed because they decide which policy instrument should be selected and without them, no instrument could be selected. Other roles would introduce complexity such as new behaviours or new parameters that are not needed for simulate the policy process, though they would allow for different insights into the model. This is considered in the second model implementation in Part III.

Furthermore, and this will be illustrated later on in the simulation examples, the simplest implementation model also considers the electorate as a separate actor role in specific simulations. The electorate's purpose is to influence the preferred states of the policymakers such that the beliefs of the electorate are reflected in their representatives'. The introduction of the electorate is not meant to be representative of any specific jurisdiction. Instead, the electorate is used to outline the potential emergent behaviours that can be present when adding only one additional actor role.

4.4 The interactions

The interactions allowed are limited to interactions between policymakers and the environment and vice versa. No other interactions are considered to limit the complexity of the model to a minimum. These interactions consist of the policymakers influencing the environment by implementing policy instruments. It also includes the environment affecting the policymakers actual beliefs through changes in the environment itself.

The common language mentioned the possibility for the introduction of interactions between the actors. However, considering the goal of the present model, these interactions are not considered necessary. They would only introduce more complexity and are not required for the policymakers to select policy instruments and therefore the simulation of the policy process.

4.5 Discussion

As outlined before, this model is kept minimally complex. One of the reasons behind this choice is to allow for a thorough exploration of the model created. Though limited in complexity, a large number of concepts, stemming either from the theory or from assumptions made to fill the gap left by the literature, are appearing for the first time in a model. They need to be tested thoroughly, and the assumptions made in this conceptualisation need to be explored. Having a model with limited complexity makes this task easier to fulfil.

Despite the limited complexity, a range of scenarios and experiments are possible with the simplest implementation policy process model. It includes getting an understanding of the impact of the issues on the actors' policy instruments selection. There is also the possibility to test the effect of the electorate influences on the policymakers, and how this influence might differ based on different parameters. Finally, the impact that the preferred states have on policy change can also be investigated through a set of experiments where these preferred states are strategically changed.

Insights can also be gained on the hybrid model as a whole where the socio-technical model and the policy process model interact through feedback effects. This is an aspect that is always present regardless of the complexity of the policy process model. Overall, despite the limited complexity, this model allows for extensive simulation and testing. The results will also show how valuable even a simple model can be to gain new insights and generate emergent behaviours.

One of the more prominent assumptions made in this conceptualisation is the use of a one-step policy cycle consisting of only a policy formulation step. Initially, a two-step approach was used, including an agenda-setting step. However, it was found that because only policymakers are present in the model, and there are no possible interactions between the actors, the agenda-setting step was made irrelevant. Other aspects relating to the way the model is set up also have impacted this fact. It was therefore decided to remove the agenda setting step from this simplest implementation model. Future discussion will go into further details on why this step was not necessary.

4.6 Conclusions

This chapter has demonstrated, in a concise manner, the use of the common language to create a policy process model. This was done by following the requirements set by the four building blocks. The resulting model is simple with only a few moving parts. However, the conceptualisation of a model is not sufficient to simulate it. For this, more details are needed, operational aspects of the model. This includes a detailed set of steps along with the equations required for the computer that will run the simulations. This is presented in the next chapter which is the formalisation of the simplest implementation model.

5 Formalisation

Following the conceptualisation of the simplest implementation model, this chapter presents the formalisation of the model. It consists of taking the concepts considered in the four building blocks and filling in the details that are needed for the operationalisation and simulation of the agent-based model. This includes a detailed step-based process, the agents' characteristics and behaviours and all of the associated equations.

5.1 The process

In the simplest implementation model, I assume that the policy process followed is a one-step process with a policy formulation step. The initialisation of the model also needs to be considered. It does not include the environment simulation that is introduced each time the model is implemented with an environment model as will be shown in the three following chapters.

The steps used to model this approach are detailed below and illustrated in Figure 5.1.

1. Initialisation:
   (a) *Trigger of external events:* Any event that the modeller decides to implement are activated at this stage of the model cycle.
   (b) *Update of the truth agent:* Information from the environment is used to inform the truth agent's actual beliefs.
   (c) *Electorate actions on policymakers*
   (d) *Transmission of the actual beliefs:* The agents are informed about the environment from the truth agent.
2. Policy formulation step:

(a) *Preferences calculation (policy instruments):* All agents update their preferences for their secondary beliefs. They then all select a policy instrument that s/he will be advocating for.

(b) *Policy instrument implementation:* The policymakers vote to implement a policy instrument. If a majority is found, then a policy instrument is implemented, if not the status quo remains.

An additional step is provided for the simulation of the policy process: the initialisation. This is needed to include actions that are considered out of the formal policy process. This includes external events, the transfer of information from the environment model and the electorate related actions when an electorate is considered. The policy formulation step is also further detailed and includes two sub steps where the agents must first calculate their preferences before a policy instrument can be selected at the policy subsystem level.

As mentioned above, more steps would need to be considered if one were to consider the overall hybrid model. These steps, in chronological order, would be the implementation of the policy instrument into the environment model - resulting in a change of the exogenous parameters, the simulation of the environment model for a given number of steps, the update of the impact of the policy instruments on the environment model, and the transfer of the data from the environment model to the policy process model.

5.2 The agents

There are two main categories of agents in the model: the active agents that have a direct influence on the agenda and the passive agents that have an indirect impact on the agenda and policy change. These categories are needed for model architecture as a means to cluster different agents roles. They do not replace these roles as they were outlined in the conceptualisation. Ultimately, active agents will be the only agents that can perform interactions, justifying their need to be clustered separately.

5.2.1 Active agents

For the simplest implementation model, following the conceptualisation, the only active agent role considered is that of policymaker. Its attributes are given as follows:

1. The *active agent* is represented as an 3-tuple given by *agent = (ID, beliefHierar-*

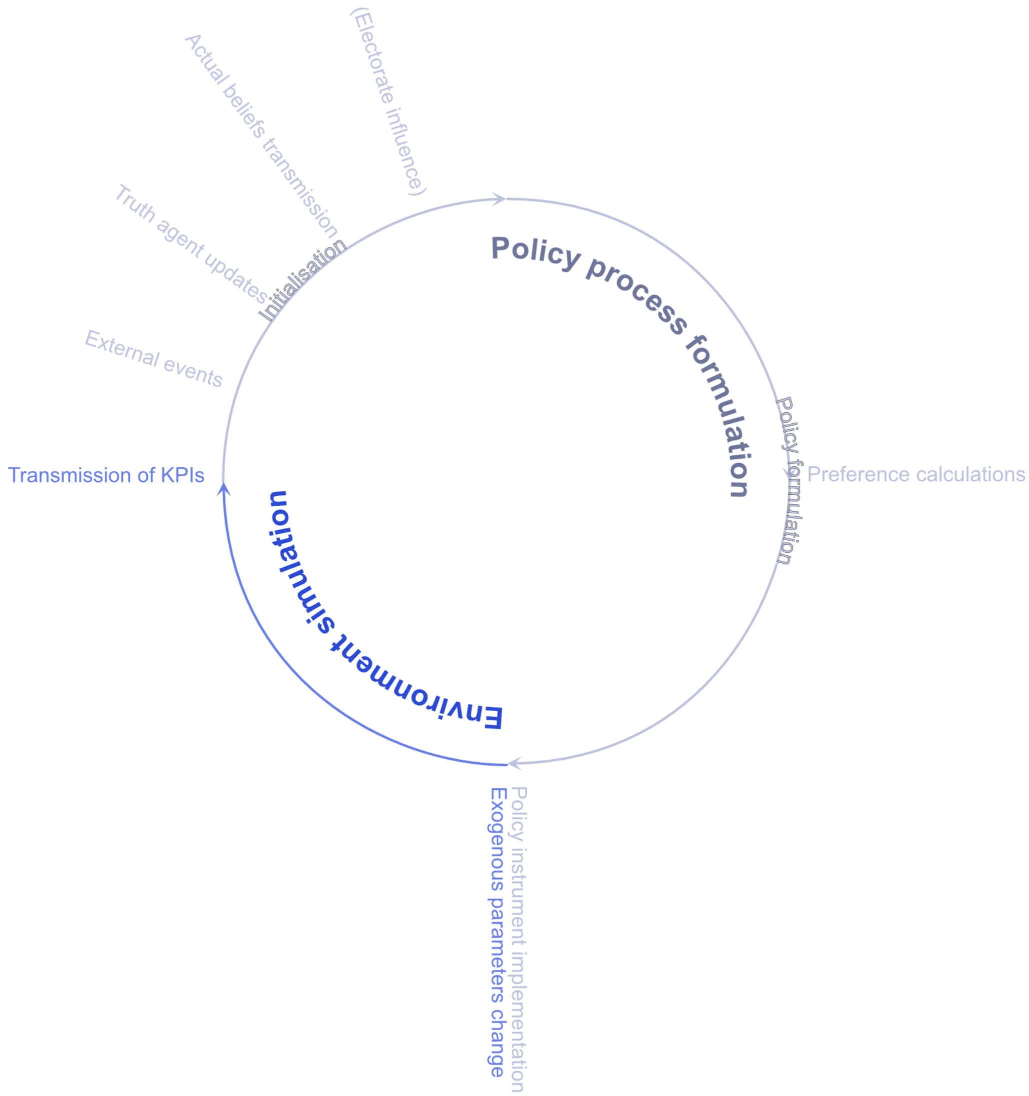

Figure 5.1 – Steps followed by the policy process for the simplest implementation model.

chy, advocacy) where *ID* is the agent unique ID, *beliefHierarchy* is the agent's personalised belief hierarchy, *advocacy* is the list of the issues the agent is supporting.

2. The *belief hierarchy* is made of the agent's own belief hierarchy structure and associated values.

3. The *advocacy* is represented as a 2-tuple *(issue_pf, policy_pf)* where *issue_pf* is the policy core issue preferred and *policy_pf* is the policy instrument selected in the policy formulation step.

5.2.2 Passive agents

The passive agents are the truth agent and the electorate.

The truth agent: The truth agent is an agent included only for programming purposes and has no correspondence to a real world actor. This agent provides a link between the environment and the policy agents. All of the environment's states are communicated to the truth agent. It uses these states to inform the actual beliefs of the active agents. The only attribute of the truth agent is the belief hierarchy composed only of actual beliefs for each of the issues.

The electorate: The electorate represents the voting public. The role of the electorate is to influence the preferred states of the policymakers. The following defines the attributes of the electorate:

The *electorate* can be given as a 2-tuple written as: *electorate = (ID, beliefHierarchy)* where *ID* is the electorate unique ID, and *beliefHierarchy* is the associated belief hierarchy of the electorate. The belief hierarchy of the electorate is similar in structure to the one of the truth agent consisting only of preferred states for all issues.

5.2.3 Belief hierarchy

The belief hierarchy follows the approach presented in the common language in Chapter 3. It is composed of two main parts: issues and causal beliefs. The issues are categorised in multiple layers: the deep core issues (the top layer - layer 1), the policy core issues (the middle layer - layer 2) and the secondary issues (the bottom layer - layer 3). The secondary issues are linked to the policy core issues through causal beliefs and similarly for the policy core issues and the deep core issues. An example of a belief hierarchy structure is shown in Figure 5.2.

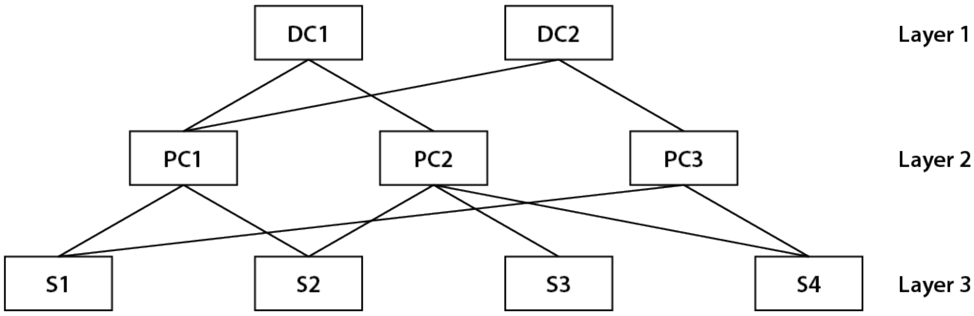

Figure 5.2 – Belief hierarchy architecture. The number of issues on each layer is provided as an example.

Each issue is composed of three parameters: the actual belief, the preferred state and the preference. The actual belief defines the view of an agent of a specific issue as it is in the environment. They can take values within the interval $[0, 1]$ meaning that all beliefs have to be normalised.

The preferred states show where an agent would like the issue to be in the future. They also take values within the normalised interval $[0, 1]$. This can be seen as the goal of the agent.

The preference, which is a calculated parameter, defines the urgency that agents place on each of the issues. It is obtained depending on the actual belief, the preferred state and the causal beliefs linked to the issue. The sum of all issue preferences on any single layer of the belief hierarchy has to be equal to 1. Its calculation is presented in a later section.

The belief hierarchy structure also contains causal beliefs. These can be understood as causal relations. They represent the understanding of an agent about how a change in one issue's actual belief will lead to a change in another on a different layer. The causal beliefs can take values within the interval $[-1, 1]$. A negative value means that an issue, when growing, will affect negatively another issue.

5.3 The actions and interactions

There are several actions and interactions that the agents can perform. I will first outline the actions that actors have to make for the model to proceed. This includes the calculation of the preferences. I then outline the policy instrument selection for the entire policy arena. Finally, I outline how I conceptualise that the electorate influences the policymakers, an action that is only considered for demonstration purposes.

5.3.1 Preference calculation - secondary issues

To calculate the preference, the gap between preferred state and actual belief for each issue is considered along with the impact of the causal belief on the gap of the issues on the above layers. The causal beliefs do not always help bridge the gap between the preferred state and the actual belief on the higher layer. If this is not the case, then the causal belief is not considered within the calculation as its effort is counter-productive within the agent's mind. The resulting equation that can be used to calculate the preference for the secondary issues is given by:

$$P_k = \frac{|G_k - B_k| + \sum_{j=1}^{n} |C_{j,k}\left(G_j - B_j\right)|}{\sum_{l=1}^{p}\left[|G_l - B_l| + \sum_{j=1}^{n}\left|C_{j,l}\left(G_{j,l} - B_{j,l}\right)\right|\right]} \tag{5.1}$$

The sums only include the second terms if C_j and $\left(G_j - B_j\right)$ have the same sign. If it is not the case, they are not considered.

And where G is the goal - or preferred state, B is the actual belief, p is defined as the number of secondary issues, k characterises the issue being selected for the calculation, j specifies the policy core issues, and C represents the causal beliefs.

Based on these preferences obtained, each agent will select one secondary issue to advocate for.

5.3.2 Policy instruments and preference calculation

The policy instruments are used to influence the environment model. Policymakers have to select one policy instrument during the policy formulation step. To assess the different policy instruments, the agents assess the impact of these instruments on the secondary issues in their belief hierarchy based on the calculated impact that they induce in the environment model. These impacts are translated into changes to the actual beliefs of the agents. Each policy instrument has the following parameters:

1. A *policy instrument* is represented as a 2-tuple *(ID, impact)* where *impact* is related to the impact of the policy on a specific secondary issue.
2. There are as many *impacts* as there are secondary issues. These impacts provide information to the agents on how much the secondary issue actual belief

parameter will change if the instrument is implemented. These are calculated in the environment model through a separate simulation.

The preference for a policy instrument is calculated using the following equation:

$$P_k = \frac{\sum_{p=0}^{n} |G_p - \left(B_p(1 + I_{k,p})\right)|}{\sum_{q=0}^{m} \sum_{p=0}^{n} |G_q - \left(B_q(1 + I_{q,p})\right)|} \tag{5.2}$$

And where k is the policy instrument for which the preference is calculated, n is the number of secondary issues, m is the number of policy instruments, I is the impact of the instrument on a specific secondary issue.

Once all the preferences have been calculated, each agent will select the instrument with the highest preference. This is the instrument they will advocate for when a decision is made on which policy should be implemented.

5.3.3 Policy instrument selection and implementation

The policy instrument selection is performed following a majority vote. If a majority of the policymakers have selected the same instrument, that instrument is implemented. If no majority is found, then the status quo is maintained, and no policy instrument is implemented.

5.3.4 Electorate passive action on policy makers

The simplest implementation model also includes the demonstration of the electorate - relating the impact the public opinion can have on the policy process. Their actions are approached following research from Laver and Sergenti (2011). It considers that the policymakers will tend to adapt their beliefs based on their electorate so as to maximise their chances to be reelected.

In the model, this translates into having the electorate passively influence the policymakers. This happens as a passive effect where the preferred states of the policymakers slowly progress towards the preferred states of the electorate. The equation to calculate the change in the preferred state of the policymaker is given as follows:

$$G_k := G_k + (G_{El} - G_k) \cdot C_i \tag{5.3}$$

where El stands for the electorate, k is a policymaker, C_i is the constant influence that allows variation in the speed of the change of the preferred states of the actors.

5.4 Conclusions

This chapter has provided all the details and equations required for the programming of the policy process model conceptualised in Chapter 4. This includes the process that the model should follow, the agents that should be included, and the actions and interactions that they can perform. Using this, the model can be coded in Python and simulated. But before any simulation can take place, an environment model needs to be considered. This is presented in the next three chapters for three different environment models.

6 The predation model

The first example of the simplest implementation model simulation is simulated using a wolf-sheep-grass predation model. The predation model, formalised below, was selected because of its simplicity. It is a well established and tested model may it be through agent-based simulation or using Lotka-Volterra differential equations (Hethcote et al., 2004; Wilensky and Reisman, 1998). Because the behaviours are known and understood, it makes it easier to test the policy process model and better understand the resulting hybrid model. With this in mind, the predation model is used to test several aspects of the policy process model, including understanding the impact of changing preferred states and changing causal beliefs on policy change. A test is also performed to understand how the electorate can impact the actors and the policymaking process.

6.1 The model

The predation model is taken directly from NetLogo[1] (Wilensky and Reisman, 1998, 2006). It is an agent-based model that has also been explored using differential equations (Hethcote et al., 2004). Throughout the literature, it has been used, extended and changed depending on the applications considered (Wang et al., 2011). Being a very simple model, it is also often used to teach agent-based modelling at university as well. It is a perfect model to illustrate emergent behaviours.

The predation model has a number of different variations. The version used here is one with three agents: wolves, sheep and grass patches. Wolves and sheep both have energy and can reproduce. They move randomly around a grid in search of food to maintain their energy without which they die. Wolves eat sheep when they

[1] https://ccl.northwestern.edu/netlogo/models/WolfSheepPredation

encounter them. Sheep eat grass patches. Every step, each sheep and wolf has a certain probability of reproducing and creating offsprings. Grass, on the other hand, can appear on every square of the grid. Once a grass patch has been eaten, it takes a few time steps for it to regrow. The results from the simulation of these agents are generally stable oscillations of the agents' population. The different parameters present throughout the model, such as the reproduction rate, the energy gained from eating or lost from moving can all have an impact on the outcome of the simulations.

Though the predation model can be seen as abstract and distant from real policy issues, it can also be linked to important policy topics. The model can relate to how we approach the dynamics of prey and predator populations. This is an important topic when considering conservation efforts or carnivore management (Boyce, 1992; Matti and Sandström, 2011; McCullough and Barrett, 1992). Through such models, one can begin to understand the dynamics that are at play at the most simple level between prey and predators. Such models can then be built upon with additional complexity so as to include environmental or geographical aspects as is done in when modelling wolf recovery for example (Boyce, 1992).

6.2 The hybrid integration

As outlined in Chapter 3, to simulate the policy process and predation models together, a hybrid framework needs to be set up. The policy process model being versatile, only a few parts of the model need to be defined as inputs so that both models can be connected. These inputs include the creation of the belief system following the structure provided in the formalisation and the definition of policy instruments based on the exogenous parameters of the predation model. These two parts are outlined below.

6.2.1 The belief system

The belief system outlined in the formalisation is one composed of three tiers of issues including at the top a deep core tier, followed by a policy core tier and a secondary tier. This structure stems from the ACF literature (Sabatier, 1987). In the present case, the deep core tier is disregarded as norms and values are not present directly or indirectly in the predation model. It would be difficult and highly speculative to include them in the belief system and as a result, add very little to the overall analysis. The belief system considered is one with two tiers: policy core secondary issues. These issues are all summarised below:

- Policy core (PC) issues:
  1. *Sheep* is defined as the number of sheep on the grid.
  2. *Wolf* is defined as the number of wolves on the grid.
  3. *Grass* is defined as the number of grass-filled patches on the grid.
- Secondary (S) issues:
  1. *Net sheep population* change is defined as the amount of sheep difference between the beginning and the end of a time step.
  2. *Net wolf population* change is defined as the amount of wolves difference between the beginning and the end of a time step.
  3. *Net grown grass patch* change is defined as the amount of grass patches difference between the beginning and the end of a time step.

The policy core and the secondary issues are closely related because of the simplicity, and therefore the limited amount of quantifiable parameters, of the predation model. This highlight one of the intricacies of adapting the policy process model to an already existing model as only parameters present in the environment model can be considered for the belief system. One essential requirement for the selection of the secondary issues is that a change in the policy instrument directly impact them. This is preferred because agents will assess the effectiveness of policy instruments based on the expected changes in secondary issues first. The impact on the policy core issues is indirect, viewed through the impact of the causal beliefs within the belief system.

6.2.2 The policy instruments

The selection of the policy instruments is linked to the available exogenous parameters in the predation model. There can only be as many policy instruments as there are parameters, excluding variations in the intensity of these instruments and the potential for grouping of instruments, also known as policy packages.

For the predation model, the main parameters retained are the parameters related to the reproduction rate of wolves and sheep, and grass patch regrowth. From this, six policy instruments are formulated plus the status quo. Each either increases or decreases the parameters by a certain magnitude, as shown below:

1. Change in sheep reproduction: ± 1% chance that a sheep will reproduce every step.

2. Change in wolf reproduction: ± 1% chance that a wolf will reproduce every step.

3. Change in grass regrowth time: ± 2 change in the number of time steps it takes grass to grow back to fully grown.

4. Status quo: Nothing is changed in this case.

The policy instruments in this example might seem far from practical policies that could be implemented in real life. The reason behind this is because they are expressed directly in changes of the exogenous parameters from the predation model. However, they can be associated to real policies. For example, changes in sheep reproduction can be implemented through culling program or the introduction of new sheep. Similarly for the wolves which have been reintroduced in the Yellowstone National Park for example (Boyce, 1992). Such culling or reintroduction would reduce or boost the reproduction rates of these animals. A similar analogy can be made for the grass regrowth time where areas can be fenced to boost the growth of grass.

An important aspect needs to be considered when policy instruments are selected: there should not be many more policy instruments than policy agents in the model. The reason for this is that if there are too many instruments, for example, several instruments with different changes of magnitude, this could split the policy instrument selection amongst policy agents. In the current way the policy process is implemented, this could results in no policy instruments ever be selected as there would not be a majority for any instruments at any point. Such a problem could be dealt with by considering a more complex algorithm, but that would further complexify the model as well. Instead, the approach taken is to limit the number of instruments to a reasonable amount of incremental instruments at the discretion of the modeller.

6.3 The research questions

Once the two models have been conceptually integrated, the scope of the simulation is known. It is possible to define some research questions of interests for this example. As stated earlier, the choice of the predation model drives the purpose of this example: it is used to explore the potential of the policy process model on a methodological level. Therefore, the research questions devised all relate to this methodological exploration goal. The questions are mostly focused on the inner workings of the policy process model. Future examples will focus on other parts of the hybrid simulation as they are more complex than the predation model and, in these cases, the interactions between the policy process and the socio-technical system will be of more interest.

The research questions are given as:

1. *What is the impact of a difference in the policy core preferred states of the agents?*
2. *What is the impact of a difference in the secondary preferred states of the agents?*
3. *What is the impact of a difference in the causal beliefs of the agents?*

For some simulation, as a demonstration, the policy process model also includes electorate agents. This exploration opportunity is used to gauge how the present implementation of an electorate impacts the rest of the model. The following research question is devised:

4. *What impact does the electorate have on the policies implemented?*

6.4 The scenarios

The scenarios are constructed based on the four research questions presented above, to try to answer them best.

In this simplest implementation model, the number of agents is irrelevant. There is the possibility to consider scenarios with more or fewer policymakers with different belief systems. However, because of the nature of the model, and the fact that agents cannot influence one another, it will always be the largest set of agents with similar beliefs that will be the deciding factor on the policy instruments that are implemented. It would therefore have no impact if more than one policymaker is considered. In effect, the model is so simple that it can be compared to an adaptive control policymaking model (Lempert et al., 2009), though the architecture of the model is much more advanced than that. Future chapters, including the ACF implementation model in the third part of the dissertation, will show how this architecture can be leveraged into more complex models.

The scenarios consist mostly of the initial parameters selection for the policy process model, including the agents' belief system. This consists of selecting the preferred states and the causal belief for each of the scenarios. The details of the beliefs selected are summarised in Table 6.1 and Table 6.2.

In total, five scenarios are considered. They are detailed below:

	Issues					
	PC1 Sheep	PC2 Wolves	PC3 Grass	S1 Sheep growth	S2 Wolves growth	S3 Grass growth
Agents						
Scenario 0	300	300	2000	100	50	200
Scenario 1	150	0	2000	100	50	200
Scenario 2	300	300	2000	50	-25	100
Scenario 3	300	300	2000	100	50	200
Scenario 4						
Policy maker	300	300	2000	100	50	200
Electorate	150	0	2000	100	50	200
Scenario 5						
Policy maker	300	300	2000	100	50	200
Electorate	300	300	2000	50	-25	100

Table 6.1 – Preferred states for the policy makers on a the interval [0,1] for all scenarios.

- Scenario 0 (Benchmark) - The benchmark scenario is used to obtain a baseline with which to compare the other scenarios. The beliefs are selected such that the outcomes of the model are average, with clear behaviour and populations that are in equilibrium.
- Scenario 1 & 2 - For this scenario, changes in the preferred states are investigating. These results in scenarios where either the policy core or the secondary issue preferred states are changed.
- Scenario 3 - The causal beliefs are changed to reflect a different understanding of the predation model from the part of the agents.
- Scenario 4 & 5 - These two scenarios are used to test the influence of the electorate by changing either the policy core or the secondary issue preferred states. Beyond this, the simulations are run for a varied electorate impact to understand how this might affect the system. Three different passive influence strengths of the electorate are considered: 2%, 20% or 50%.

Note from all of the values of Table 6.1 that the beliefs of the agent do not have to be consistent. The agent is perfectly allowed to wish for an equal number of sheep and wolves. This only illustrates the impossibility of such preferred states while illustrating the agent's lack of understanding. It is not an issue with the model itself. Similar inconsistencies in actor beliefs can also be found in reality.

Scenario 0/1/2/4/5				Scenario 3			
	PC1	PC2	PC3		PC1	PC2	PC3
-S1	1.00	0.75	-0.75	-S1	-0.50	-0.10	0.25
-S2	-0.75	1.00	0.25	-S2	0.05	-0.50	-0.25
-S3	0.50	0.75	1.00	-S3	-0.25	0.00	-0.50

Table 6.2 – Causal beliefs for the policy maker for all scenarios. These causal relations can be read as: an increase of 1 in S2 will lead to a decrease of 0.75 in PC1. They are all given on the interval [-1,1].

6.5 The simulations

For each of the scenarios mentioned above, the hybrid model is simulated fifty times. This is needed as each simulation is different taking into account the fact that the sheep and wolves move randomly, affecting the states of the predation model differently every time, and in turn the agents' actual beliefs. Fifty repetitions are considered sufficient to gain knowledge of the overall output distribution. Each simulation is run for 1600 time steps for the predation model. Every hundred-time steps, the policy process model is simulated, and a policy instrument is selected. This results in the selection of sixteen policies over the length of a single simulation.

The predation model is always initialised with the same values, contrarily to the policy process model. No scenarios are considered as this is a policy process exploration exercise. All parameters are selected such that they lead to a stable cyclical simulation when considering the predation model alone. Introducing the policy process might disrupt this through the adoption of policy instruments. The predation model starts with a hundred sheep and fifty wolves. The sheep have a four per cent probability of reproducing while the wolves have a five per cent probability. The grass regrowth time is thirty time-steps, and the total energy gained from food for sheep is four while for wolves, it is thirty. The grid is a square grid with a side length of twenty-five.

6.6 The results

The analysis of the simulations' results has for goal to answer the four research questions outlined previously. Throughout this section, two types of results are presented: the predation model agent populations and the policy instruments selections. The first focuses on the evolution of the wolf, sheep and grass populations over time. They are provided in percentiles to better see the results' distribution. The second graph category plots the sum of the policy instruments selected by the policymaker at each

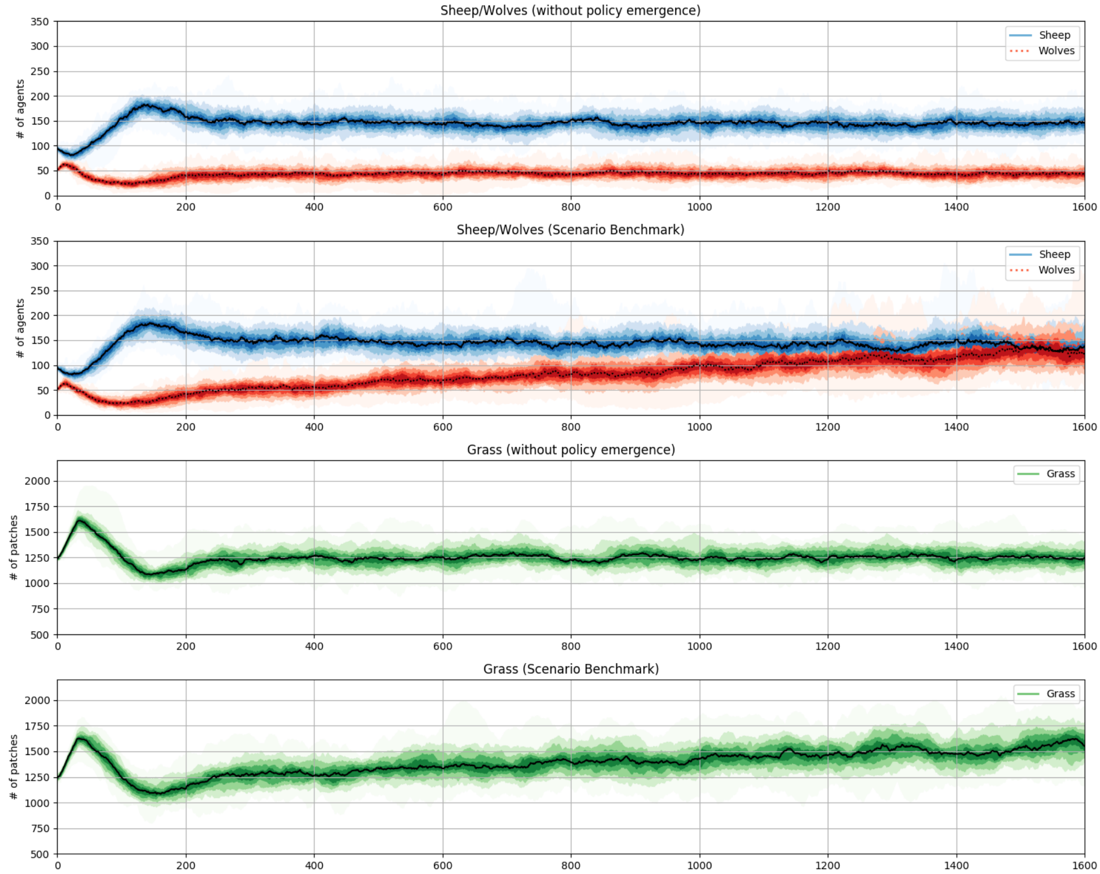

Figure 6.1 – Comparison of the population of sheep and wolves between a simulation of the predation model alone and the benchmark scenario (scenario 0) incorporating the policy emergence model.

time step. Because there are fifty simulations, the sum of all results can never be higher than fifty at each time step. This should give an idea of which policy instruments are favoured by the policymaker at what point in time.

For conciseness, not all graphs are presented in this chapter. For the complete graphs of all results for all scenarios, see Appendix C.

6.6.1 With and without policy emergence

Before answering the research questions, an observation is made on the difference that is present in the results between the benchmark scenario and a situation where there is no policy process model. The results are shown for the predation model populations in Figure 6.1. This comparison is used to understand the effect of the policy process endogenisation on the outcomes of the predation model.

Figure 6.1 shows that there is a definite impact on the predation model populations as a result of the introduction of policy instruments. Without the policy emergence model, the sheep and wolf populations stabilise around different levels. The wolf population stabilises around fifty wolves and the sheep population around 150 individuals. When the policy process is introduced, the behaviours are very different. The wolf population is increasing at a constant rate until it reaches the sheep population, at which point the simulations end. The sheep population remains stable, propped up by a large amount of grass continuously growing.

To understand these results, one must understand why and how the policy instruments are implemented. The policymaker attempts to achieve its preferred states through the policies it implements. These preferred state would like to achieve the impossible: sustaining a sheep population of three hundred while at the same time having three hundred wolves. The wolves and sheep growth preferred states are also very high and unrealistic for the dynamics at hand. A sheep growth of one hundred per time step cannot lead to a population of only three hundred sheep. It results in a policymaker that is never satisfied with the states of the predation model and will endlessly attempt to influence the model with the introduction of new policy instruments. This has the impact of making the model less stable, leading to the behaviours observed. Different beliefs would have affected the model differently.

6.6.2 Changing policy core preferred states

Continuing the exploration of the hybrid model, scenario 1 is used to observe what happens if the policy core issues' preferred states are changed. The expectation is that this will impact the overall sheep and wolf populations through a change in the policy instrument selected by the policy agent. The results are compared with the ones obtained for the benchmark scenario to assess this. They are plotted in Figure 6.2 and Figure 6.3.

The policy agent's sheep, wolves and grass count preferences have been adjusted to test these research questions. The new preferred states are highly unrealistic with a population of 150 sheep for 0 wolves and 2000 grass patches. At the same time, the secondary issues, that is, the growth parameters, are kept at the same levels. Such unrealistic preferred states are purposefully selected to test the model boundaries and to make it more likely to observe a change in the outcomes of the predation model.

The results show a limited change in the outcomes of the predation model. There is an overlap of the populations in both scenarios with the population of wolves growing while the sheep population remains somewhat stable. The slight difference is in the

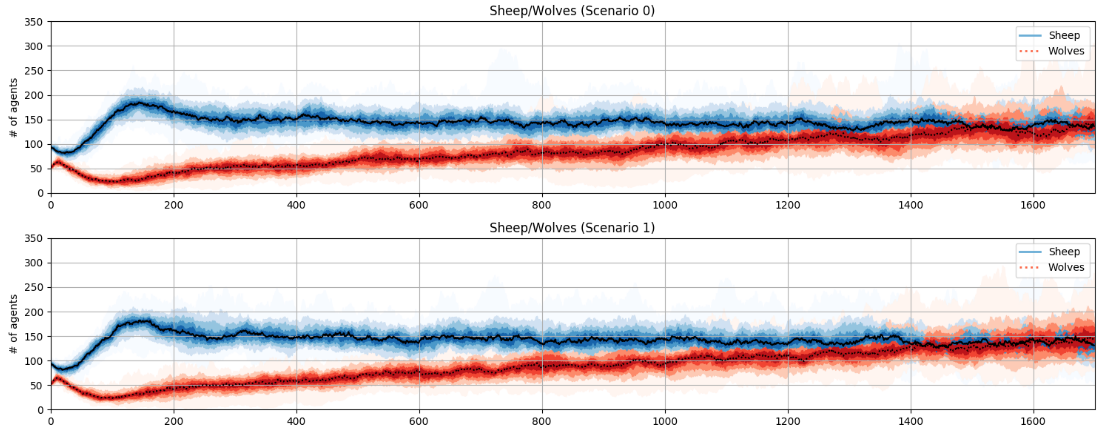

Figure 6.2 – Predation model populations for the benchmark scenario and scenario 1.

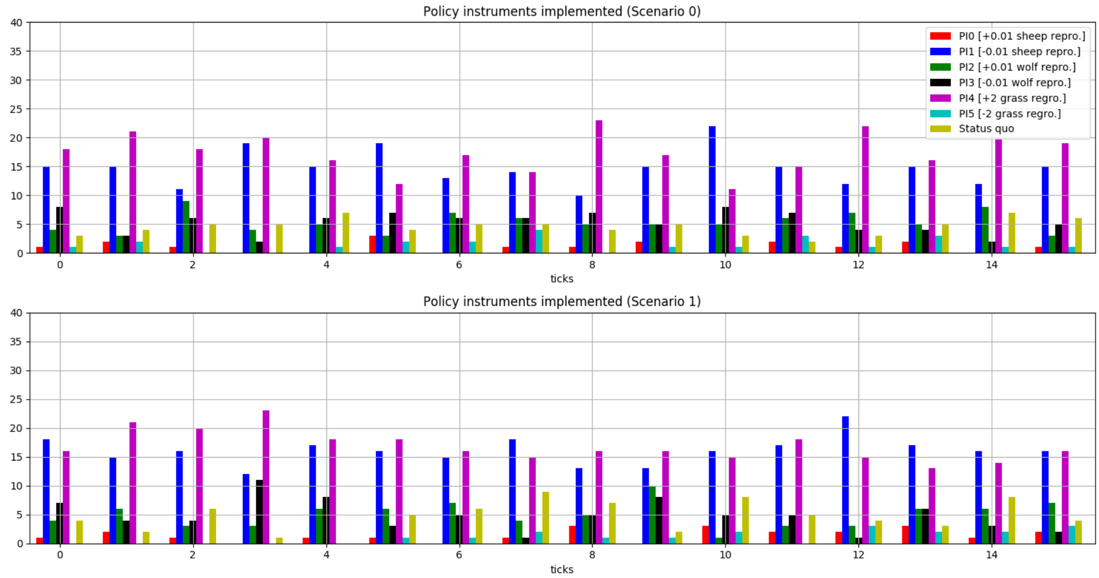

Figure 6.3 – Policy instrument selection for the benchmark scenario and scenario 1.

speed of the wolf population growth, which is slightly higher, happening hundred-time steps before, compared to scenario 1. This suggests only a limited impact on policy change as the policy core issues vary. The grass count is also similar for results in both scenarios.

These changes can be explained by the policies that are implemented by the agent, as shown in Figure 6.3. They are very similar. One can see that for scenario 0, it is policy instrument PI4 (+2 grass regrowth) that is the most often selected. For scenario 1, the selection is more balanced between the policy instrument PI1 and policy instrument PI4 (- 0.01 sheep reproduction and +2 grass growth). The differences are minimal.

Overall, a change in the policy core issue preferred states leads to only limited changes in the policy instruments selected and the predation model populations. These limited impacts can be explained by the fact that the policy core issues have no direct impact on the predation model. They are more akin to long-term guidance for the agent's policy instruments' selection. This is because they take time to affect the predation model. Simulating for a more extended period might have lead to more differences in the results between both scenarios.

6.6.3 Changing secondary preferred states

When changing the secondary issue preferred states, the expectation is that the impact will be more pronounced than for the previous scenario. In this case, it is because a change in the preferred states has the potential to directly affect the policy instruments' selection as these are chosen based on the expected impact on the secondary issues. Once again, the new secondary issue preferred states were set at levels which are unrealistic compared to reality, and to the policy core issue preferred states (which have been reset to the benchmark scenario levels).

The results are shown in Figure 6.4 for the predation model and Figure 6.5 for the policy instrument selection. The first thing that can be observed is a scatter in the population results of scenario 2. Several simulations see a crash in their wolf populations and a most likely associated skyrocketing of the sheep populations. Overall the population averages remain constant, with wolves around fifty and sheep around one hundred and fifty. It can be attributed to the preferred states being highly unrealistic with the goal of having a constant negative growth for the wolf population - though maintaining a population of three hundred wolves - for example. The impact of the policy instruments in this scenario is clear and observable.

This drastic change can be explained by the policy instruments preferred by the policy

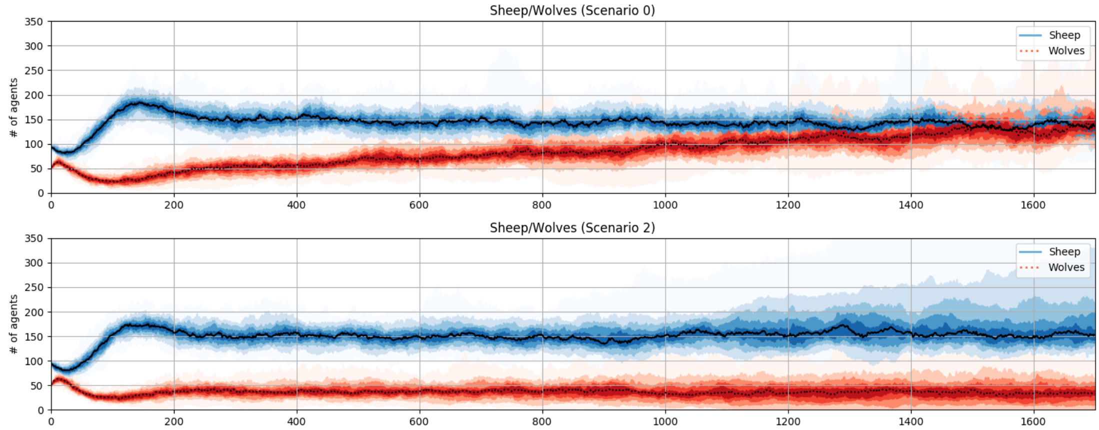

Figure 6.4 – Predation model populations for the benchmark scenario and scenario 2.

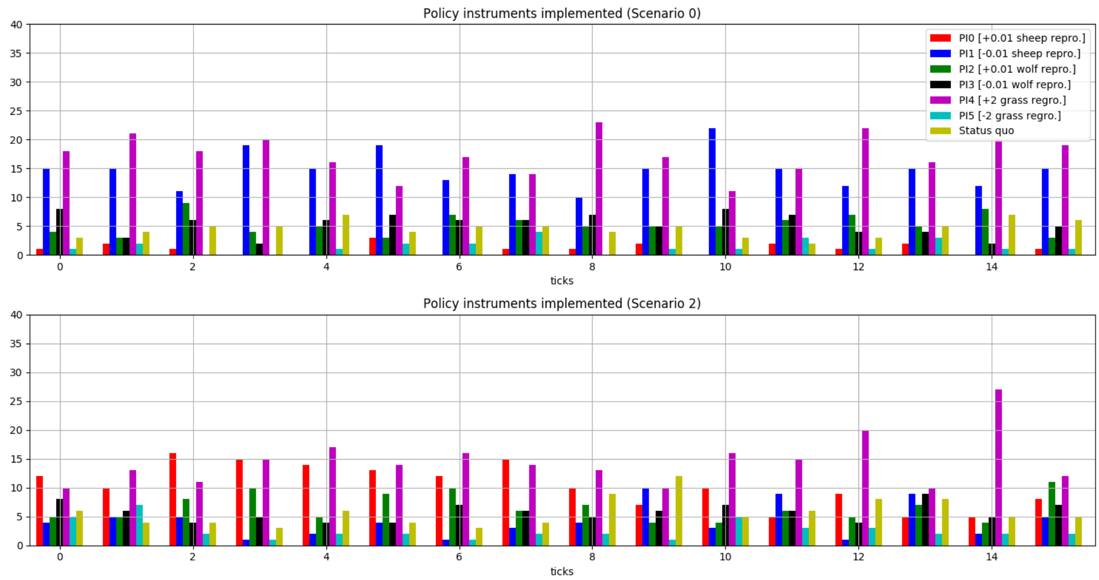

Figure 6.5 – Policy instrument selection for the benchmark scenario and scenario 2.

agent. The policy most selected in scenario 2 is initially policy instrument 0, which leads to an increase in the sheep reproduction rate. This is the first scenario for which this policy instrument is preferred as much. More importantly, throughout most of the simulations, policy instrument PI1 is completely neglected compared to the benchmark scenario where it was often selected. Policy instrument PI4 remains one that is most often favoured, as the policy agent sees nothing wrong in increasing the speed at which grass grows back. Together, the grass regrowth rate and the sheep reproduction rate can help sustain one another and help explain the explosion of the sheep population in several simulations while the amount of grass patches always remains in check. This, combined with a limited selection of policy instrument PI3, which reduces the wolves reproduction rate, leads to the populations observed in Figure 6.4.

Overall, one can see that changing the secondary issues' preferred states has a direct impact on the policies selected by the policy agent as expected. The outcome within the predation model, however, was not necessarily expected. In many cases, the policy agent ends up running the predation system into the ground because of its beliefs. Though this is unfortunate, it is not necessarily unrealistic.

6.6.4 Changing understanding of the world

The last parameter that is changed within the simplest implementation model is the causal belief parameter. These beliefs set the agent's understanding of the predation model by having them link changes in the secondary issues to changes in the policy core issues. In scenario 3, the expectation is only a limited impact on the predation model because of limited change of the policy instruments selected. This is because a change in the causal belief only has an indirect and marginal effect on the agent's decision-making process, similarly to the policy core preferred states. In scenario 3, the causal beliefs of the agent are reversed.

The results are shown in Figure 6.6 and Figure 6.7. The populations' distribution seems to follow similar trends to the one for the benchmark scenario. The main difference is that it is slower. In scenario 3, the wolf population slowly grows while the sheep population remains mostly at the same level. This is surprising considering a total upending of the causal beliefs of the agent. This can be explained by the way the belief system is built, and the preferences are calculated. More important changes can be observed in the policy instruments selected by the agents. Instead of dominating the selection with policy instruments PI1 and PI4, several policy instruments seem to be often selected, with policy instruments PI1 and PI4 still selected often.

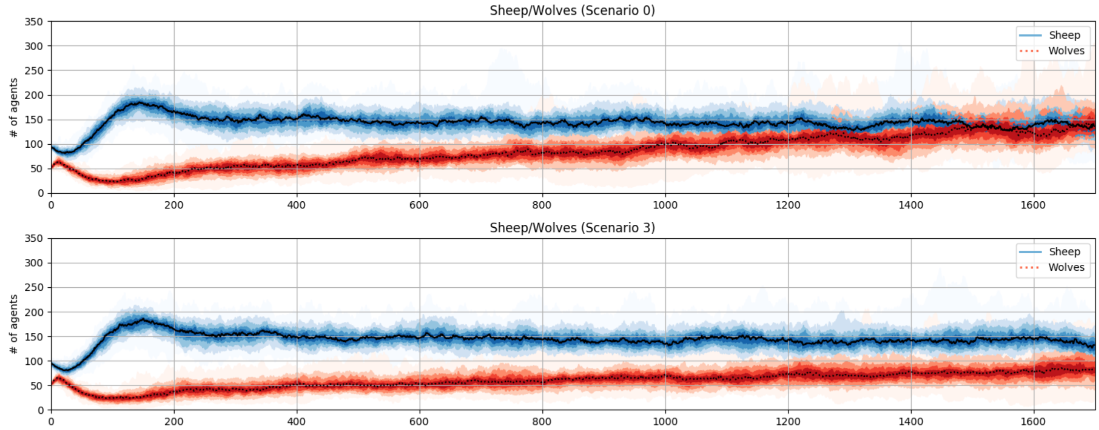

Figure 6.6 – Predation model populations for the benchmark scenario and scenario 3.

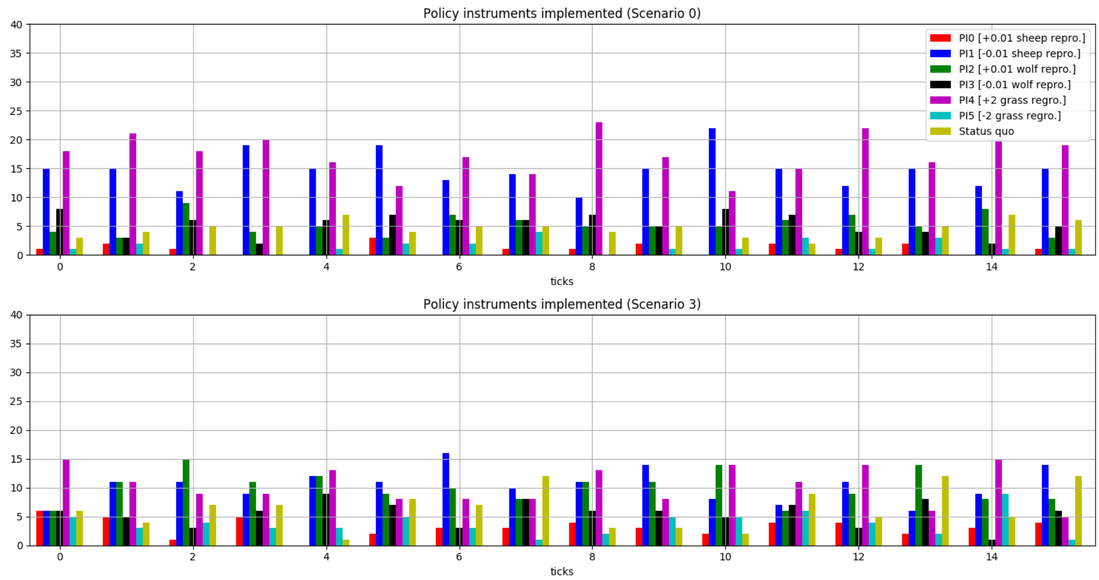

Figure 6.7 – Policy instrument selection for the benchmark scenario and scenario 3.

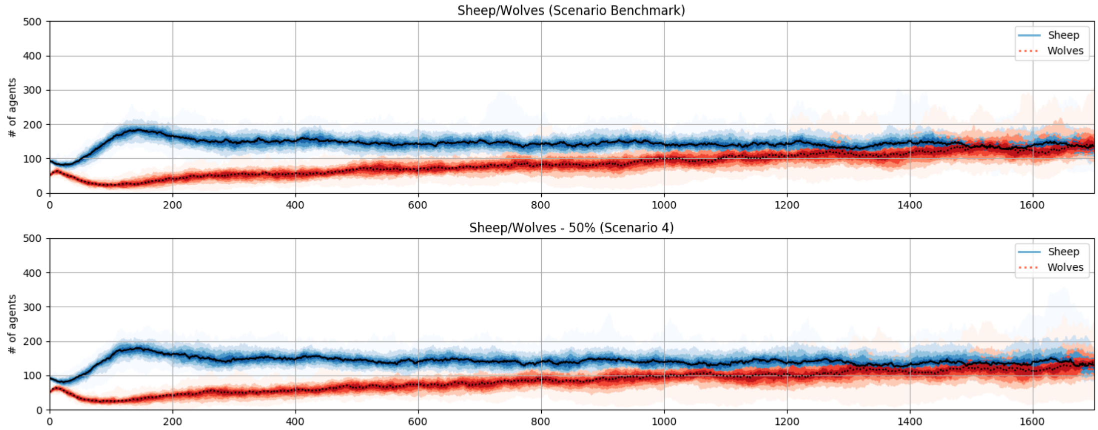

Figure 6.8 – Predation model populations for the benchmark scenario, scenario 2 and scenario 4 (50% influence only).

The architecture of the model can explain these results. A change in the causal relations affects, to a limited extent, the secondary issue preferences. It has a remote effect on the selection of the policy instruments. Overall, this leads to limited changes in the policy instruments being selected - besides having policy instrument PI4 selected less than for the benchmark scenario. Beyond this, the results seem to suggest an almost random selection as the agent cannot make sense of the expected impact of policy instruments and how they might affect the policy core issues.

The results from the predation system are more difficult to understand. One way to explain these results is that the system is inherently stable and that through the introduction of just about any policy instruments. The effects balance out. This means that the agent will never meet its goals, whatever policy instruments it implements. Note that policy instruments PI1 and PI4 are still implemented quite a lot and therefore probably driving the behaviour of the model.

6.6.5 The impact of the electorate

The electorate introduction is expected to result in similar outcomes as for the first two scenarios but with impacts over time. This is because the electorate influences the preferred states of the policy agent. This is what was changed, at once and not overtime, in scenario 1 and 2. In the present scenarios, the influence of the electorate is tested for three different speeds where, on the one hand, the electorate have minimal influence over time while on the opposite side, their influence is considerable.

Some of the results are shown in Figure 6.8, Figure 6.9 and Figure 6.10. These results

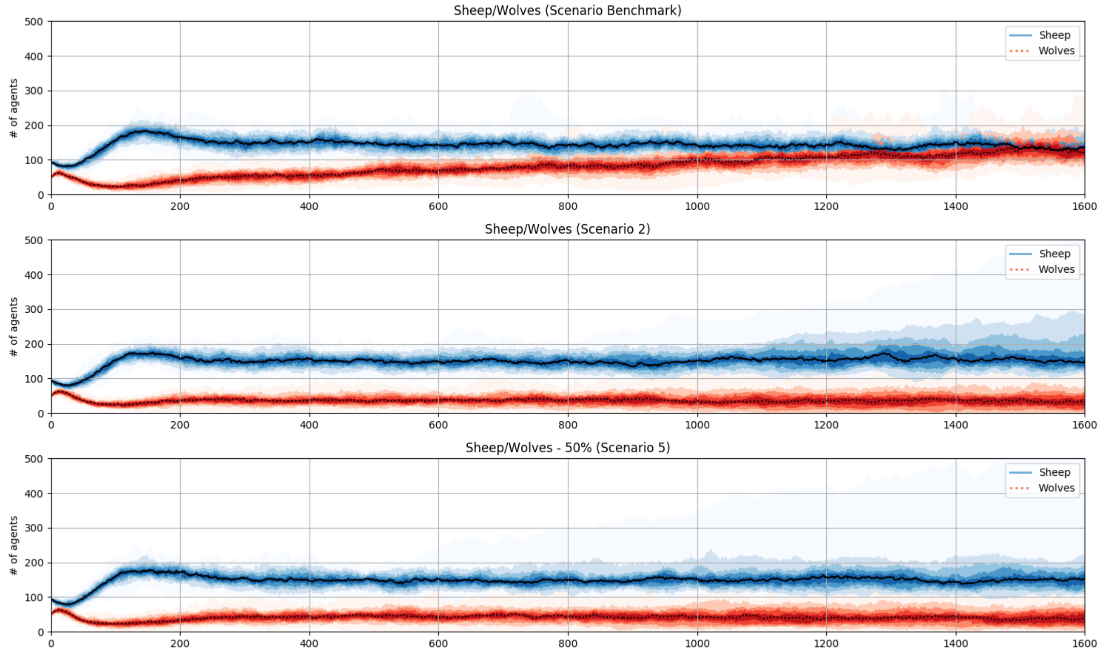

Figure 6.9 – Predation model populations for the benchmark scenario, scenario 2 and scenario 5 (50% influence only).

are segregated depending on the strength of the influence of the electorate. The results show that in the case of scenario 4, the predation population hardly changes. This can be expected as this scenario is very much a parallel of scenario 1 but with less strength as the electorate is slow to influence the policymaker than an outright change initially. They make little difference.

Similarly, the results from scenario 5 show some changes that are in line with the results of scenario 2. In some cases, this translates in the crashing of the wolf population. This is more prevalent in the simulations where the impact of the electorate is strong as the policy core preferred states of the policy agent is more quickly affected by the electorate. However, this remains hard to spot from the results.

This trend of parallels with scenarios 1 and 2 is confirmed when looking at the policy instruments being selected. There, the importance of the strength of the influence is also made very clear with changes in the results occurring only when the electorate influence parameter passes 20%. At 2% influence, the presence of the electorate makes very little difference within the time scale of the simulation considered.

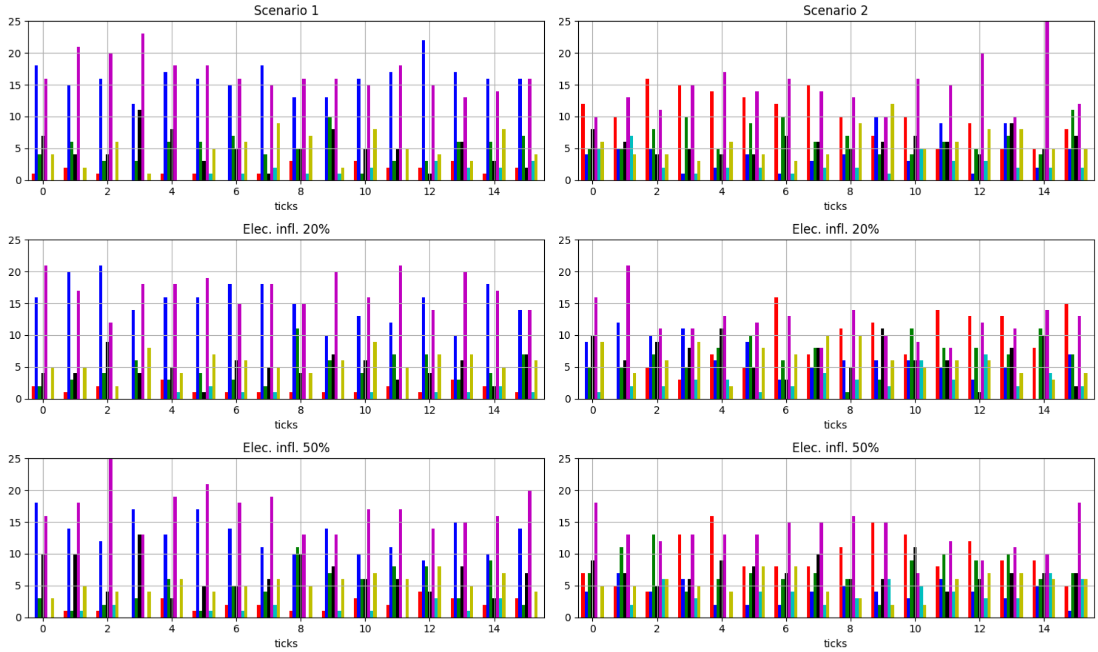

Figure 6.10 – Policy instrument selection for scenario 1, 2, and scenarios 4 and 5 (20% and 50% influence only).

6.7 Discussion

This simplest implementation model endogenisation, in the present case with a predation model, has shown the kind of results that can be obtained when simulating the policy process. Through these results, some limitations have emerged, demonstrating that simulation is not a silver bullet solution and that it also raises additional questions regarding the policy process.

6.7.1 Hybrid integration

The hybrid integration, though short and at the beginning of the study, is a crucial aspect of the policy process endogenisation. First, it requires an in-depth knowledge of the environment model. It is needed because one cannot create either the belief system or the policy instruments without this knowledge. It forces the modeller to understand all the dynamics possible in the model used.

Second, the design of the belief system is crucial. In the present case, the predation model was limited in complexity. Hence it limited the possibilities in the belief system creation. However, when considering a more complex model, the modeller will have to choose which issues are significant enough that they need to be considered in the belief system. Then the modeller will have to assess whether the issue corresponds to a secondary, a policy core or a deep core issue. This selection should also consider a balanced set of issues, with a preference for more secondary issues than policy core issues. The same applies for the policy instruments. As outlined before, there should be a limited set of instruments with potentially incremental changes in the exogenous parameters of the socio-technical system. All of this requires an in-depth knowledge of the model being studied.

Finally, an aspect that was not shown here is the calibration of the belief system. To calculate the agents' actual beliefs, informed by the predation model, a mapping is used such that the actual beliefs are always contained within an interval $[0,1]$. To do this, one has to assume a minimum and maximum value for all parameters. Choosing inappropriate values can affect the results obtained, and, most importantly, will have an impact on the decision-making process of the agents. It is crucial to choose these values appropriately.

Further remarks will be made on this topic in subsequent chapters. For the predation model, the tricky part was to select a minimum and maximum values for the agent growths such that all possible scenarios were taken into account. The populations had a minimum by default and a maximum that could easily be bounded.

6.7.2 The policy process simulation potential

The policy process endogenisation presented here has shown the potential of the approach proposed. This simulation with a predation model has shown that one can gain insights may it relate to the policy sciences or modelling and simulation. Several insights can be gained from the policy process model itself. Beyond the construction of the model which requires many assumptions that can be adapted depending on the theories followed or modeller goals, insights can also be gained by looking at the policy instruments implemented by the policy agents. In the present model, it is possible to track how a particular policy instrument made it through to implementation and how the beliefs of the agent have influenced these decisions. Beyond observing, it is possible to explain policy change within the simulation's context.

Through the scenarios, it has been shown that it is also possible to explore how changes in the agents' beliefs can lead to policy change. It includes understanding how similar changes in the different tiers of the belief system have different impacts on the ultimate policy instruments selected. These behaviours, again, are the results of the assumptions used to build the policy process model. This is an important point because if the results of the models were to not match reality, it could be due to incorrect assumptions that would then have to be changed. This is something that can be validated. Modelling allows for such changes to be done quickly and that tested almost instantaneously. This makes the model particularly useful for the theory exploration of the policy process theories.

Researchers can also use this model to understand better how a socio-technical system will be affected by policy instruments. This was shown with the predation model, where one can observe how the populations of agents varied over time, depending on the policy process. Because this is a long term simulation, it is also possible to gather information on the policy change path. Depending on the initial states of the predation and policy process models, a set of policy instruments can be identified as most likely to be implemented. Pathways can be identified from these. For the predation model, depending on the beliefs of the policy agents, it was possible to show a progression in the selection of the policy instruments, or in some scenarios, the constant selection of the same policy instruments. This can be valuable information for policy analysis and to advise policymakers. As outlined before, beyond illustrating the pathways, the model can also be used to understand why specific pathways have been followed instead of others, something that is much harder to do in a qualitative case study.

This approach also provides an avenue to study emergent phenomena that can emerge from the interaction of a socio-technical system and its associated policy process. Irrational behaviours were illustrated by agents having contradictory beliefs and choosing policies that had little impact on the outcomes of the model or that were contradictory to their own goals. Provided that the modellers have the initial data for the agents in the policy process, insights can be gained on such irrational behaviours, which happen in the model as a result of rational decision making. To best illustrate this phenomenon, all that has to be done is to provide inaccurate and inconsistent causal beliefs to the agents. They will then go on to randomly select policy instruments, confused by their own beliefs. The use of incoherent preferred states can also have the same effect, preferences that have been observed in real cases.

When looking solely at the predation model, and for any socio-technical system model, the results have also helped discern in which conditions the wolf and sheep agent populations are more likely to crash. After specific combinations of policy instruments being implemented, the results have shown that the model can crash, often with one of the population plummeting to zero. These population crashes only occurred in particularly extreme scenarios such as when the beliefs of the agents are made incoherent (a preference for negative wolf growth while keeping a population positive). By extrapolating such results to the different socio-technical system, one can understand the potential of endogenising the policy process. This will be further demonstrated in subsequent examples in this dissertation.

Despite all of the insights that can be gained from the simplest implementation model, one criticism is expected to emerge relating to the policy process theories: the simplest implementation model does not represent a single policy process theory. Instead, it is composed of some cherry-picked concepts from different policy theories. Because of this, the validity of these concepts considered outside of their respective theory can be questioned. This is a valid criticism, but it misses the point. First, the goal of this simplest implementation model is only to show that it is possible to simulate the policy process, something that is currently hard to find in the literature. Second, as was argued and shown before, no theory is sufficient to represent the entire policy process. It is therefore not possible to use concepts from a single theory or framework to construct a complete and simulatable model of the policy process. This does not mean that the approach presented here is the only one possible. It is one of a multitude of potential approach. Finally, although the simplest implementation model does not represent a single theory, the goal of the model presented in the third part of this dissertation will show such an example with the ACF.

6.7.3 The agent distribution

During the creation of the hybrid model, it was argued that the number of policy-makers did not make a difference for the present experiment. This is a consequential assumption that has limited the potential studies of this simplest implementation model. With this assumption in place, it becomes impossible to test more involved scenarios that would explore the power dynamics between agents. Such scenarios could see a variety of agents with varying beliefs, may that be causal beliefs or preferred states. The scenarios could then be built such that the agents' beliefs are changed at critical points along with the simulation. The goal would be to analyse how such changes might affect the policies implemented and, in turn, affect the predation model.

Without manually changing the beliefs of the agents throughout the simulation, it makes no difference whether to include one or more agents, even with varying beliefs. If two agents are present, and then if they do not agree, the status quo will always remain as no agreement will be found. If three agents are considered, then it will always be the combination of agents that have similar preferences that will decide the policy instrument. Because the preferred states and the causal beliefs of the agents do not vary over time, this fact would not change, rendering the third agent useless. For this to change, the policy process model should allow mechanics to influence the beliefs of the agents throughout the simulation. This is shown in the third part of this dissertation.

6.7.4 The inherent stability of the policy process model

Some aspects of the results relating to the policy process were more limited. For example, it could have been expected to see a more significant impact from a change in the preferred states or the causal beliefs of the policy agents. When a causal belief is inverted, that means the understanding of the world of the agents is pushed upside down. It would have been possible to expect that their policy instrument selection would be drastically different and in turn, the predation model would be upended. As the results have shown, this was not the case. And, though this might have been unexpected, it was perfectly explainable.

Considering the structure of the model, a change in the causal beliefs or the preferred states of an agent is one part of the long chain that can lead to policy change. Such a change does not necessarily lead to a change in the agents' preferences which are calculated based on the actual beliefs and the other issues in the same tier. A change in a preferred state of one policy core issue, for example, is one step removed from the calculation of the preferences of the secondary issues and two steps removed from the policy instrument selection. So if the policy core preferred states change, that change needs to be large enough that it will affect both the secondary preference choice and, subsequently, the policy instrument selection. This is unlikely. Only if that happens can a different policy be implemented, and then the predation model can be affected.

Change is further limited due to the incremental nature of the policy instruments. As was shown previously, only a limited set of policy instruments are proposed to avoid agents from fragmenting their choice with the result of never reaching an agreement. Note that in the present case, there was only one agent. This has lead to the selection of small and incremental policy instruments. An instrument only implemented once will have a limited impact on the behaviour of the predation model. Only the recurrent implementation of the same instrument can lead to significant change. However, in making their decisions, the agents do not consider prior policy selection. They only base their choice on the current states of the model and their beliefs. This limit the chance that they select the same instrument several times strategically.

Finally, the stability of the overall model can also be linked to the predation model. The predation model is a simple model with built-in behaviours. For example, it would be impossible to have significantly more wolves than sheep sustainably. It would also be impossible to have a large amount of sheep with no grass patches. Therefore, the changes that can be observed in the results can only be minor in the population of wolves or sheep. The most extreme cases are when the populations collapse because of instability in the parameters. However, this only rarely happens, and as the results have shown, it is only the wolf populations that crash while the sheep populations mostly maintain themselves because the agents tend to prefer maintaining the grass growth rate or even increasing it, in turn sustaining the sheep.

6.7.5 The causal beliefs and policy instruments assessment

The causal beliefs are used as a proxy for the agents' understanding of the environment. They represent the way agents understand how one issue influence another. This is done using the lowest level complexity possible and within the constraints of the belief system structure that is provided by the ACF: unidirectional linear causalities. The agents' representation of the world is therefore limited and does not allow to consider feedback effects between issues or even causal links between issues on the same tier of the belief system. This simplification is an assumption for the belief system.

However, other approaches could be used. Using the assumption made, one could push the representation of the belief system to its natural most complex conclusion: a nested simplified system dynamics representation of the environment tailored to each policy agent. Using a simple system dynamics model for each of the tiers of the belief system, agents would add the possibility to have a more sophisticated understanding of the inner workings of the environment - and the possibility of better capturing the agents' decision-making process. Feedbacks and causal relations between issues on the same tier would be possible. Each agent would, of course, have different values for their causal beliefs making their understanding of the environment different and tailored.

However, replacing the current somewhat simplified belief system by one that effectively contains a sub-model in every policy agent might not affect policy change. Other limitations, such as the incremental nature of the policy instruments, could prevent any significant impact, blunting the added complexity and bloating the entire policy process model. There might, however, be value in using such an advanced belief system submodel in cases where the study specifically focuses on understanding how the agent's understanding of the environment might affect their decision-making process.

The current belief system approach also has repercussions on the way the policy instruments are assessed. Currently, the impact of the policy instrument is assessed using a perfect information approach. The policy instrument is introduced into a copy of the environmental model. That copy of the model is simulated for a given time. The results are then compared to the initial states. The differences are used to define the impact of the policy instrument. This is done for all policy instruments. This is a perfect information approach because each agent's assessment of the policy instrument is not based on their understanding of the environment but instead on the model itself. It guarantees that the assessment is correct. It also means that all agents have the same assessment for all policy instruments, an aspect that is not necessarily realistic.

Using the complex belief system approach presented above, this could be resolved. Each agent would use their belief system to assess the impact of each policy instruments themselves. The impacts would depend on the understanding of the environment of the agents. Once again, such details might not be needed for an overall study of the hybrid model but could be helpful for particular research questions focusing on this point explicitly.

Finally, the current policy instrument assessment approach can provide some depth to the model. It allows for short-term and long-term thinking agents. The copy of the model used to assess the policy instruments is simulated for a given amount of time. In the present simulation that was set to the time between policy processes. This can be considered short term. It is possible to change this simulation time to several years or more, in which case, this would account for longer-term policy agents that are looking far ahead at the impact of the instruments. Of course, such a long-term perspective comes with its downsides. It is likely to be less accurate as uncertainty increases the further the simulation runs - and it does not take into account subsequent implementation of policies.

6.8 Conclusions

Overall, this chapter has shown the usefulness of the policy process model with a simple environment model. The next step is to use this model with a realistic full-blown environment model. This will help deal with the limitations in complexity that relate to the predation model. The use of a more complex environment model will help exhibit additional and different emergent behaviours.

In the next chapter, this is done using a model of the Swiss electricity market. For that example, the hybrid integration is more complex. The results obtained are more challenging to understand as they stem from complex feedback loops within the socio-technical system itself and with the policy process model. Some of the limitations highlighted here are dealt with as well.

7 The electricity market model

In this second example of the simplest implementation model, the policy process model is simulated within an electricity market model. This is the first implementation with a fully-fledged complex socio-technical model. The goal of this chapter is different than for the predation model. It is no longer to study the policy process model and its potential, but it is to demonstrate how the model developed can be used to gain deeper insights from a hypothetical realistic complex socio-technical system. For this, a Swiss electricity market model is used. A different socio-technical model could have be used for this task. The model was chosen because of its availability to the research within a tight time scope, the model's validity and the limited efforts required to link it to the policy process model developed in this thesis.

In this chapter, I look at the hypothetical impact that an electorate might have on an electricity market - considering the way it has been implemented, and the potential for a transition within a context of increasing electrification of the overall energy system. This current context is translated into the model through a range of electricity demand growth scenarios. This chapter first outlines the model broadly, then discusses the research questions and experiments that can be addressed using such a complex socio-technical system, before presenting the results and discussing them.

The work presented in this chapter is part of a book chapter written in cooperation with Prof. Matthias Finger of the EPFL (Klein and Finger, 2020).

7.1 The model

The electricity market model used for this chapter is one of Switzerland that was initially developed by van Baal (2016). This model was built following a system dynamics approach. In subsequent years, it was further developed and adapted into a hybrid system dynamics - agent-based model for research on a Swiss strategic energy reserve (van Baal, 2019). Overall, this made the model somewhat unwieldy and challenging to connect to the agent-based policy process model. For this reason, compounded with proprietary software issues, the model was entirely rebuilt for this work using an agent-based approach only and in the open-source programming language Python.

The electricity market model is a model that can be split into two main parts: the spot market and the investor modules. The spot market is responsible for matching supply and demand. This is done on an hourly basis. The demand is composed of the inflexible Swiss domestic electricity consumption and the elastic demand for foreign countries or assets that can store electricity. The supply is provided by assets, power plants, that represent the Swiss electric production park. The different technologies considered are solar, wind, combined-cycle gas turbines (CCGT), waste incineration, run of river, nuclear, hydro and hydro-pumping. Different costs are incurred for each of these technologies, affecting the price at which they bid on the spot market. Imports and exports from neighbouring countries are also considered for the spot market.

The investor module deals with investment in new capacity. These new investments are limited to CCGT, wind and solar assets only. Investor agents perform them based on a profitability index that includes a net present value calculation. The investors also have decision making power on what happens at the end of the life of an asset. They have the choice between extending the lifetime of the asset, mothballing it or decommissioning it.

The detailed formalisation of the model can be found in Appendix A.

7.2 The hybrid integration

Similarly to what was done for the predation model, the policy process and electricity market models need to be integrated into a single hybrid model. This consists of defining a belief system based on the parameters and dynamics within the electricity market model and defining a set of policy instruments that the policy agents will be able to select. This is detailed in this section in the same fashion as was done previously.

7.2.1 The belief system

The belief system is specific to the electricity market model as the key performance indicators inform it from that model. Beyond this, the issues selected are also informed by Markard et al. (2016) due to the fact that this is a model that was built specifically for Switzerland. In the paper, Markard et al. present the belief system of the actors in the Swiss electricity sector. They identified several issues (they referred to them as beliefs) that are specific to the Swiss context. However, they were limited to the deep core and policy core levels. Secondary issues were not presented and are informed only by the model and its dynamics in the present application.

The difficulty in the creation of the belief system is to associate the right indicators to the right issues. The first step is not to consider the deep core issues. These are normative issues that are beyond the boundaries of the model and, therefore, out of the scope. They are not a crucial aspect of the process which is focused on policy core issues. The next step is to select the secondary issues. These are extracted directly within the Swiss electricity market agent-based model. Not all indicators are made into secondary issues as most indicators are not important enough. Only five are ultimately retained, accounting for the most critical aspects of the model. They can provide a good understanding of the electricity model to the policy agents and allow for a sufficient set of policy instruments. Finally, there is the selection of the policy core issues. These are, in general, aggregates of the secondary issues. They are selected based on what was shown in Markard et al. (2016), as is detailed below. They do not stem directly from the model but are instead combinations of several beliefs.

For the policy core issues, there is an additional aspect that needs to be considered. In work performed by Markard et al., policy core issues within the Swiss electricity market policy subsystem were identified. These were: the seriousness of the problem, the role of the state, the environment, the economy and society. Several of these are not within the scope of the model at hand. The environment and the economy are the only two that can be calculated or deducted from other parameters within the model. They are selected as the policy core issues. Markard et al. (2016) also identified four secondary issues. They are not considered suitable for the model as they are formulated as questions rather than issues. Ultimately, the policy core issues are calculated using linear equations, including the indicators also used for secondary issues.

The belief system is given as follows:

- Policy core (PC) issues:
    1. Economy
    2. Environment
- Secondary (S) issues:

1. Renewable energy production
2. Electricity prices
3. Renewable energy investments
4. Domestic level emissions
5. Imported emissions

The economy takes into account elements related to the profits of firms along with the Swiss security of supply. The environment takes into account aspects such as emissions, the amount of renewable energy and the amount of imported emissions. The equations that are used to obtain the values of the actual beliefs for all issues are provided in Appendix B.

7.2.2 The policy instruments

The selection of the policy instruments is linked to the secondary issues that were chosen for the belief system. Similarly to the predation model, no policy packages are considered. Each policy instrument has one impact on one exogenous parameter. In total, eleven policy instruments are introduced:

1. Solar subsidies: ± 0.04 CHF/MWh for solar investments.
2. Wind turbine subsidies: ± 0.04 CHF/MWh for wind turbine investments.
3. Agent's hurdle rate: ± 0.02% hurdle rate change for the investors.
4. Carbon tax on domestic emissions: ± 10 CHF/MWh for costs of domestic emissions.
5. Carbon tax on imported emissions: ± 10 CHF/MWh for costs of imported emissions.
6. Status quo: Nothing changes.

Note that the same considerations had to be made as for the predation model. The policy instruments had to be limited to incremental changes. Additionally, only a limited set of instruments, representing what could be introduced as policies in real life, were considered. The hurdle rate is the only instrument that was added that would be more difficult to represent in real life though it can be translated into an incentive by policymakers for investments, or its opposite, through tax cuts or other tools.

	Issues				
	S1 RES [%]	S2 Price [CHF/MWh]	S3 REI [%]	S4 Dom. em. [tons]	S5 Imp. em. [tons]
Agents					
Policy makers	60	50	70	4 000 000	60 000
Electorate	100	75	100	0	5 000

Table 7.1 – Preferred states for the policy makers and the electorate on a the interval [0,1] for all scenarios.

7.3 The research question

Through the use of a complex socio-technical system, such as this model of the Swiss electricity market, the goal is to study the introduction of a policy process model to better account for deep uncertainty in the electricity market. With that in mind, one research question is formulated as follows: *What additional insights can be gained from the endogenisation of a policy process model that includes a simple electorate influence as part of the electricity market model?*

Note that results to this questions are not meant to be used in the real Swiss context but should only be used as an example of how deep uncertainty can be better understood using the present model. The way the electorate is conceptualised and included within the policy process model used in this chapter is not representative of the Swiss situation which is much more complicated.

7.4 The scenarios

The scenarios are constructed with the sole goal of answering the research question. Several different scenarios can be considered for the electricity model while at the same time considering scenarios for the policy process in the same fashion as it was done for the predation model. For the electricity model, the demand growth is varied to reflect the future potential for the electrification of the energy sector. This includes electricity growth scenarios of 0%, 1.5 % and 3%. For the policy process model, the influence of the electorate is varied to represent increasingly vocal voters or an increasingly competitive political field better. This is the same as was done for the last two scenarios of the predation model, where the electorate impact was tested as well. The impact of the electorate is varied from 0% to 5% and 50%.

	PC1	PC2
-S1	0.00	0.25
-S2	0.75	0.00
-S3	0.25	0.25
-S4	0.00	0.25
-S5	0.00	0.25

Table 7.2 – Causal beliefs for the policy makers for all scenarios. They are all given on the interval [-1,1].

Within this initial model, it is assumed that the majority coalition, as described in Markard et al. (2016), is one that is pro-economy. The same assumption is made as in the predation model: only one policymaker is considered because of the model simplicity. The agent's preferred states are defined by a pro-economy stance. It is also assumed that the policymaker has reasonable causal beliefs, which means that it reflects, to the greatest extent, what is happening in the electricity market model. For the electorate, because of the goal to show a transition, the assumption is that the electorate is pro-ecology. Its preferred states are defined such that it is pro-renewable energy production and against emissions. This is all reflected in the tables with the preferred states in Table 7.1 and the causal beliefs in Table 7.2.

7.5 The simulations

For the simulations, and like with the predation model, the two models are not run the same number of steps. This is because the electricity model is run on an hourly basis, and the policymaker is not realistically expected to select policies every hour. Instead, it is assumed that it decides on the implementation of new policies every three years. This results in a hybrid model where the electricity model is simulated for three years, then the policymaker decides on a policy. Once that is done, the electricity model continues for another three years and so on. The impact of policies is estimated for the policy agents based on what might happen over the following three years if that policy were to be implemented. The hybrid simulation is run for 30 years starting in 2016 with historical data for the Swiss electricity market model used until 2018. Furthermore, each combination of scenarios is run fifty times for repetitions purposes. This provides more representative results that are analysed in the next section.

7.6 The results

The results are presented using three main figures only. In Figure 7.1, the monthly averaged electricity prices for all of the scenario combinations considered is plotted. The plots are grouped according to electorate influence with each of the three demand growth scenario plotted for each of the electorate influence parameters. Figure 7.2 presents the electricity supplied by the energy sources that can be invested in - solar, wind and CCGT. The results are split depending on the demand growth scenarios. The demand is also plotted as a reference. Finally, the third figure presents the sum of policies implemented across 50 simulations for the three different electorate influence rates for a demand growth of 3%. This is presented in Figure 7.3. Note that all figures are plotted in this section. All figures regarding this model can be found in Appendix D.

Figure 7.1 shows that prices are bound to go up over time, regardless of the scenario, though this increase is less pronounced for smaller demand growth scenarios. Additionally, it shows that the price rise is correlated to demand growth and not so much to electorate influence. The highest growth in prices is related to the 3% demand growth scenario, as could be expected. In this scenario, demand would go from roughly 10GW to 20GW, an additional demand that would have to be fulfilled by a combination of new investments and additional imports from neighbouring countries as shown in Figure 7.2. One can conclude from this that prices of electricity across the country will depend heavily on the extent and speed of electrification and, in some respect, on the speed of efficiency increases as well.

Figure 7.2 outlines the supply from solar, wind and CCGT, and therefore by proxy, the investments in those technologies over time. One can see a drastic difference that is driven by the demand growth, and not so much by the electorate influence. The electorate influence has so little impact on the results that it is difficult to differentiate the different electorate influence scenarios in the graphs. For the scenario with 0% growth, most of the investments are limited to wind and solar. Note that, within the electricity model, solar and wind each has a cap of installed capacity based on numbers from the Swiss Federal Office for Energy (SFOE), for wind at 2282MW and for solar at 19702MW. There is some investment in CCGT for this scenario, but overall gas plants are not considered financially viable and are therefore avoided. The electricity prices are too low. Things change in the case of 1.5% and 3% demand growth scenarios. In the first scenario, there is a clear investment in CCGT capacity around the time where solar and wind are maxed out. This is justified by increasing demand and rising prices. For the 3% scenario, the rise in prices is so quick that investment in CCGT occurs early on, even before solar and wind have been maxed out. CCGT is viable early on and is seen as the only possible technology able to bridge the demand gap,

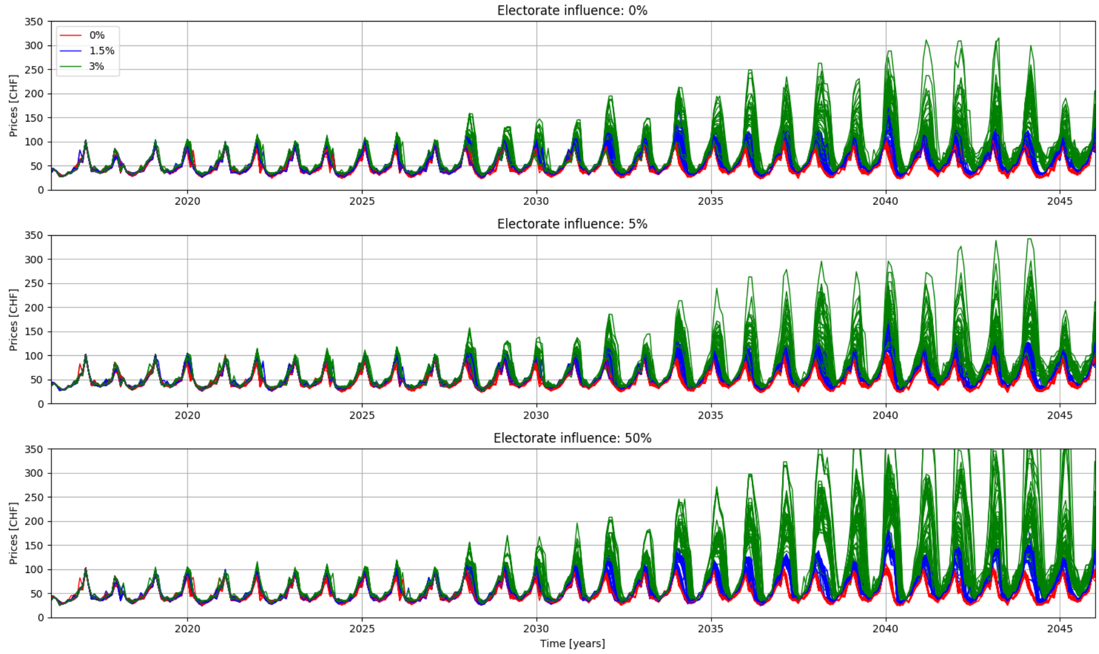

Figure 7.1 – Monthly averaged electricity prices, split according to the three electorate influence rates and for the three demand growth scenarios.

considering that only a limited increase in the imports and exports is available.

Finally, Figure 7.3 is the only figure that describes the outcome of the policymaking process model. The figure was made by summing, for each time step, the policy instruments. With this, one can see which policy instrument is favoured by the policymaker and at which time step over the fifty simulations for each of the scenarios. Only the 3% demand growth scenario results are presented here because these results are the most striking and provide the most insights. The results for the other scenarios are similar. Within these results, one can observe that the agent overall prefers the use of the carbon tax policy, may it be domestic or foreign. This is the case for all scenario growth with one difference between them: the timing at which the taxes are imposed. In the case where the electorate has a quick influence on the policymakers, the taxes are implemented very early on. It is the only scenario where the tax is imposed on the second step in 2019. For the other scenarios, where the electorate has less or no influence on the policymaker, it takes more time for the tax to be considered, with the policy actor implementing a reduction of the already present carbon tax in the first steps. Note that the introduction of the carbon tax is the main driver for the electricity prices increase, and this is reflected in Figure 7.1. All other policies are implemented at a much slower rate throughout the simulations, and when they implement a subsidy or a change in the hurdle rate, it is often balanced out later by another change in that

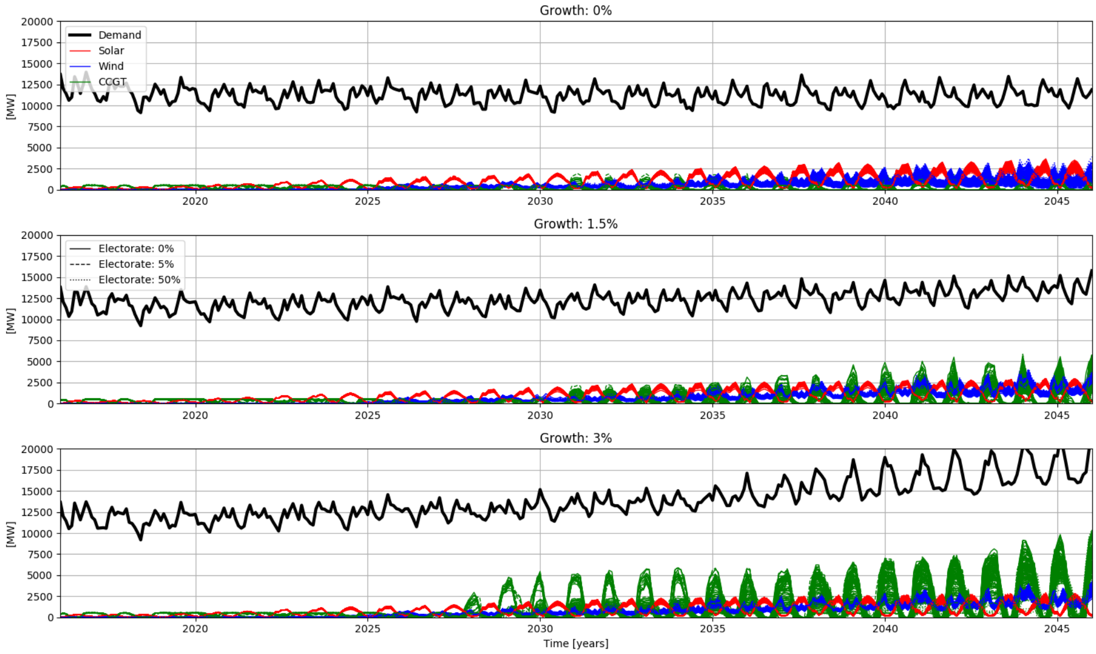

Figure 7.2 – Overall demand with the electricity supplied by solar, wind and CCGT sources for all scenarios.

policy, having little long term impact on the overall simulation. Only the carbon tax is implemented and mostly not removed for all simulations considered.

7.7 Discussion

The research focused on the methodological insights that can be gained from the endogenisation of the policy process in an electricity market model. In the present case, this further related to the additional inclusion of a simple electorate influence element in the policy process model and it hints at the impact of policy-related deep uncertainty in a complex socio-technical system. The results have shown that the a number of methodological insights can be gained from the endogenisation of a simple one-step policy process model.

Before going over these insights, it is important to address one of the major limitation of the present policy process model: it does not resemble the Swiss policy process, especially when including the electorate. The policy process in Switzerland is complex and relates to the country being a direct democracy. As a result, the electorate can brought into the process not only to vote in its parliamentary members but also through referendums on items that have passed through parliament. The present model does not address that, it simplifies drastically this process by reducing the

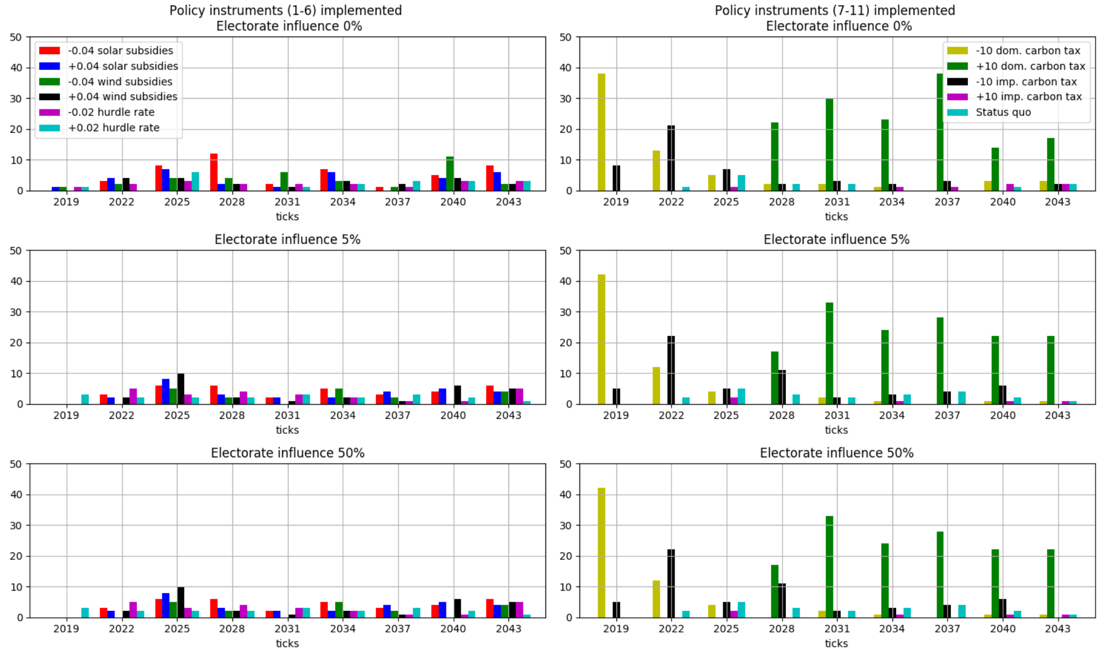

Figure 7.3 – Cumulative policy instrument selection for a 3% electricity demand growth scenario for all fifty repetitions for each simulation.

electorate to influencers. Though this is not necessarily wrong, it also means that the model cannot and should not be considered to have any predictive powers for the Swiss case. Instead, the focus should be placed on the trends and behaviours that the results exhibit, an approach used to discuss the results onwards. This can be done as the electricity market model was validated in previous work (van Baal, 2016).

First, the results have shown that the electorate can have a significant impact on the policymaker. However, its impact on the electricity system as a whole is much more questionable. As shown throughout the results, demand growth has a much higher impact on the outcome of the electricity mix than the electorate. It seems that the only impact the electorate has is on the prices. The electorate influences the policymaker to be more environmentally friendly, which in turn implies the sustained implementation of a carbon tax. However, because of the electricity model's assumptions - a cap on the amount of solar and wind and a maximum net transfer capacity - this cannot result in an increase in solar and wind power plants. It results instead in the additional construction of CCGT plants regardless of carbon price because of the lack of other possible options. Therefore, an enormous growth in emissions can be observed throughout all growth scenarios.

These results should be put into the broader context of the simplistic approach used for the electricity model. The electricity model is limited to a conventional approach of the electricity market. It does not consider the technological innovation that would accompany the electrification assumed for the high demand growth scenario. This includes the introduction of batteries, demand-side management and other innovations that will come as a result of an increased digitalisation. When considered, such innovation could significantly affect the demand curve, even if the demand growth were to reach 3% annually as is postulated in the most extreme scenario in this chapter.

On the other hand, the results can also be considered optimistic in certain respects. They show a significant adoption of wind turbines, on the scale that the SFOE has predicted. The current situation has shown that wind turbines currently face strong local opposition in Switzerland and that the goals are not currently met. The results therefore show a situation that would be optimal for the adoption of wind power, but that remains unlikely without more social acceptance, accompanied by significant regulatory and legal change.

Besides the optimistic increase of wind turbines, the model suggests that a large number of CCGT power plants would be constructed in Switzerland. These results are purely based on the economics of the model and do not reflect the politics, or acceptance, of this technology in Switzerland. In fact, proposals for the construction of a CCGT power plant would likely face the same amount of rejection as a wind turbine, if not more. Highlighting this issue, and especially in a 1.5% or 3% demand growth scenario, it is essential to note that the production stemming from CCGT power plants would need to be filled by other electricity sources. From the results of this model, it is unclear what those could be. A more detailed model could be used to assess this.

The results also depict a trend where Switzerland could become even more seasonally dependent on its neighbours. The decrease in nuclear power production, accompanied by more solar power and the potential of demand growth means that Switzerland would be even less able to supply its entire electricity in the winter period and would further depend on its neighbours for its security of supply. This is a politically sensitive issue, considering the current negotiations between Switzerland and the European Union (van Baal and Finger, 2019). Nonetheless, it is an issue that will become more important as time passes as the results suggest.

Methodological elements related to the policymaking process model have also affected the results. Within the policymaking process model, the policymaker is only able to select one policy instrument at a time. This is an assumption that is made based on insights gained from the punctuated equilibrium theory (Baumgartner et al., 2014). A different implementation could have seen policymakers negotiating to introduce a policy package that would include several different policy instruments. This might affect the results as it could include the introduction of subsidies for solar power plants and an increase of carbon taxes by policymakers. However, considering the net difference between the number of times the carbon tax policy instruments are selected and the other instruments, it is unlikely that considering such a policy package would make a significant difference.

Beyond the potential for policy packages, the policymaking process model used within this chapter is inherently simple. Most of the model is replaced by one agent whose preferred states change over time and who tests policy instruments, selecting the best and implementing it. One of the reasons for the use of this model was to present a novel approach for the study of socio-technical system simulations. The other reason was the research question. The policymaking model used is sufficient to answer the question.

For other research questions and more policy-focused research questions, the policymaking model can be made to incorporate more complexity beyond being a better representation of the Swiss policy process. For example, Markard et al. (2016) have shown that coalitions play a vital role in the shaping of the electricity market policies. To investigate the impact of coalitions on the electricity system, the policymaking process model would need to incorporate aspects such as coalitions, policy entrepreneurs, agent-on-agent interactions, and potential additional elements such as imperfect knowledge transfer. The approach presented here allows for that when needed. However, this complexity increase would bring limitations of its own but would allow for answering many more research questions.

To answer the research question directly, a number of insights could be gained from the endogenisation of the policy process, and especially because the electorate was included. These insights go from understanding the impact that the electorate can have, while considering how the electorate was included, to placing a focus on trends that are likely to develop in the future as a result or despite the introduction of new policy instruments. This is the case with the increased emissions trend, for example, despite the expectation that additional taxes would help reduce emissions. These insights are one way to deal with policy-related deep uncertainty. Though not predictive, they can nonetheless be used to inform policymakers on the workings of the overall system and the trends that can be expected. This could help them inform their decision making process.

Overall, the study presented here was the first time that such an approach was utilised. This allowed gaining a better understanding of the feedback effects between the policymaking process and the electricity system, despite the simplicity of both models. It has also shown that there is room for improvement, mostly through an increase in complexity in either model. Increasing complexity in the electricity model would allow for more insightful results and could help guide the design of new regulations. A more complex policymaking process might grant more insights on the interactions between the policy process and the electricity market.

7.8 Conclusions

The goal of this chapter was to show that additional insights can be gained when considering a more complex environment model. This has been shown here with an electricity market model. The insights ranged from additional information on the electricity market model itself to insights on the interaction between the two models.

Similarly to its implementation with the predation model, a number of limitations were raised relating to the simplest implementation policy process model, most relating to its lack of complexity. It's one step approach and lack of agent interactions means a lot of the dynamics of the process are left out. One way to remedy to this problem would be to simulate a policy process theory that consider these aspects. It would mean a policy process model with more complexity, provide a second example of the use of the common language, and present a policy process model specifically designed to emulate a specific theory. This is done in the subsequent part with an implementation of the advocacy coalition framework. But before this can be done, the next chapter presents a short third environment model to outline the versatility of the approach developed in this thesis. The environment model is a system dynamics model and it helps show that the policy process agent-based model developed here is versatile enough to be be used with different modelling paradigms.

8 The flood safety model

The final example for the simplest implementation model of the policy process is one that is designed to demonstrate that the model can also be used with socio-technical systems built in a different modelling paradigm. The goal behind this example is to show the versatility of the approach developed in this dissertation. In this chapter, the policy process is introduced within a flood safety system dynamics model built for a previous project. Agents within the policy process have a choice between four different policy packages that prioritise different approaches to flood safety. With this example, I look at how different beliefs on flood safety can lead to different flood safety protection outcomes.

The work presented in this chapter is part of a paper that is in the process of being published in collaboration with Dr. Jill Slinger of the TU Delft.

8.1 The model

The system dynamics model simulates the complexity of a flood defence system for a ring dyke. A ring dike encircles a plot of land to protect it from floods. This approach is used in flood protection in The Netherlands. The model presented in this paper is similar to the one presented in Klein et al. (2016). It uses a co-flow approach for determining the food safety standard. The system diagram of the model is shown in Figure 8.1 with the exogenous factors in green, the endogenous factors in grey and the policy packages in magenta (these are detailed in a later section).

The model is composed of three main parts: the levee co-flow which tracks the levee height and the quality of the levees, a part which establishes the safety standards and a part that tracks the safety level with the perceived and official safety. Floods are introduced into the system as external inputs. The Causal Loop Diagram (CLD)

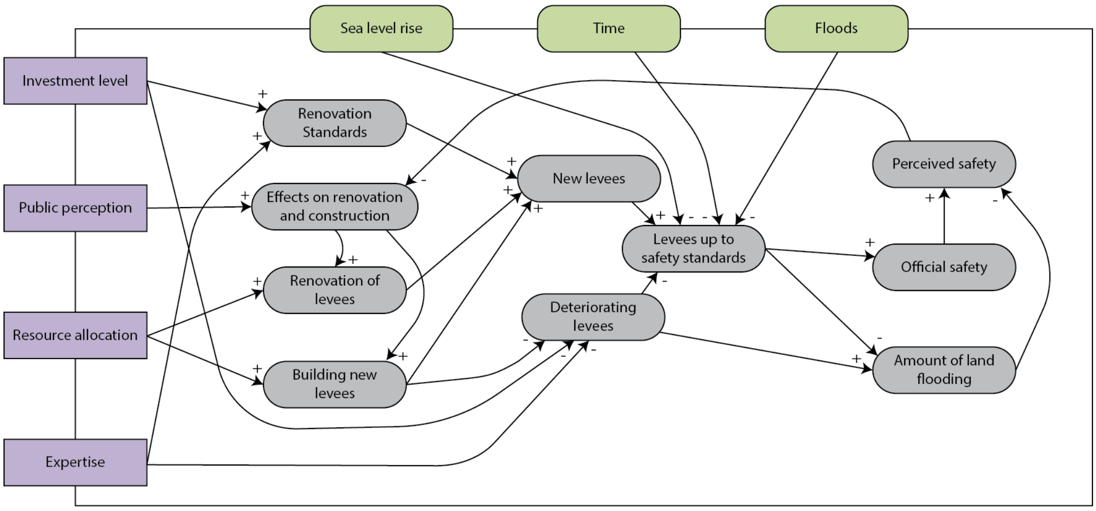

Figure 8.1 – System diagram of the system dynamics model where the non-policy dependent exogenous factors are presented in green, the endogenous factors in grey and the policy dependent exogenous factors are in magenta (adapted from (Klein et al., 2016)).

explaining this model is shown in Figure 8.2. The CLD is composed of four main loops which represent the maintenance repair loop, two sufficient safety loops which act in opposite ways and the levee obsolescence loop. The maintenance loop describes the relationship between the number of levees and the renovation of levees. The sufficient safety loops are feedback loops where an increase of levees leads to an increase in perceived safety and effects on renovation and construction of levees. Finally, the levee obsolescence loop is a negative delayed loop where an increase in the levees that should be maintained leads to an increase in levee obsolescence.

One of the main components of the model is the co-flow tracking the levee ageing. First, new levees are designed, they are then built into standard levees which deteriorate over time. These deteriorated, or old, levees will then either disappear or be renovated and turn back into standard levees (meeting the continually increasing safety standards). The co-flow is used to track the levee heights, which is a proxy for levee quality based on the safety norms established.

Linked to these levees is the calculation of both the official current safety and the perceived safety. The official current safety is a function of the length of levees and their quality (height of the levees). The perceived safety is a function of the official current safety, but it is a measure that reflects the confidence of inhabitants and their memory of the damage of previous floods. When destructive floods occur, this perceived safety will plummet regardless of the official current safety of the ring dyke

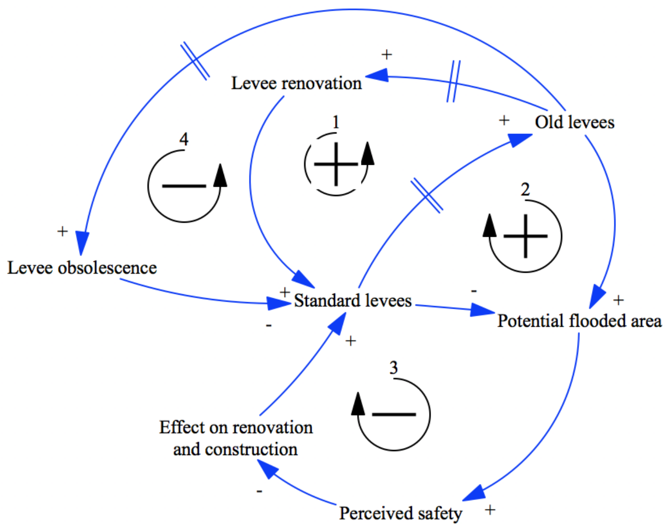

Figure 8.2 – Causal loop diagram representing the system dynamics model conceptualisation. (1) is the maintenance repair loop, (2) and (3) are sufficient safety loops acting in opposite directions and (4) is the broken levee loop.

system but depending on the amount of land flooded.

The required safety standards are primarily determined based on the number of levees, their quality and the current sea level. Throughout the simulation, the sea level rises, leading to an increase in the required safety standard. It is also affected by several external factors, such as the planning horizon. These factors can differ from country to country.

The final part of the model allows for the creation of different types of floods. These floods can be produced as pulses but also as one-off events. They can vary in duration and happen anytime throughout the simulation.

8.2 The hybrid integration

Similarly to both previous models, the flood safety model needs to be integrated into the policy process model. This consists of the creation of a belief system and the selection of a set of policy instruments. Additionally, because the flood safety model is a system dynamics model, I also look at the practical integration of the two models and what it entails to connect a system dynamics model to an agent-based model.

8.2.1 The belief system

The creation of the belief system is similar to the previous models. The fact that the model is a system dynamics model does not induce any changes. Parameters still have to be found and selected as issues within the three tiers of the belief hierarchy. Once again, the deep core issues are not considered for the present model. They can be seen as norms and values, and these are not present within the flood safety models. Such concepts are too high level to be part of such a model. Therefore, only policy core and secondary issues are considered for the belief system.

The policy core issues selected are chosen because they are of paramount interest to the policy agents in the model. These two issues are investment priority and safety. The investment priority issue relates to the priority that policy agents have regarding renovating old levees and constructing new levees. Agents can have a different view of the approach that should be undertaken. Note that financially, a focus on renovation might be more economical and could suggest an underlying priority for another strategic approach which might be more related to evacuation planning instead of levee construction. This element is not present in this model. The second issue is safety. All agents are interested in remaining safe from flooding. This safety measure is obtained as a mix of the technical safety and the perceived safety by the inhabitants of the land protected by the ring dyke in the system dynamics model. Because neither of these two policy core issues can be found directly within the model, they are calculated using a linear combination of parameters present in the model.

The secondary issues are the more basic issues within the system dynamics model. The issues considered are mostly performance indicators that can be obtained directly from the model. They can be mapped directly into actual beliefs for the policy agents. They are the following:

1. Standard levee safety
2. Old levee safety
3. Standard levee length
4. Old levee length

8.2.2 The policy packages

For the past two model examples, policy instruments were preferred. This was due to the limited amount of overlapping parameters present in each of these models that

	Policy packages							
	0	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Exogenous parameters	Expertise		Public perception		Resource allocation		Investment level	
Ageing time	+1	-1			+1	- 1		
Obsolescence time	+2	-2			+2	-2		
Design time	-0.25	+0.25			-0.25	+0.25		
Flood perc. time			+0.05	-0.05				
Renovation time					-0.5	+0.5		
Adjustment time					-5	+5		
Planning horizon							+10	-10
Reno. standard							+0.05	-0.05
Construction time							-0.5	+0.5

Table 8.1 – Table presenting the different policy packages.

made policy instruments more logical than policy packages. In the present case, the number of exogenous parameters is significant. This allows the creation of policy packages that deal with larger parts of the model. These are detailed here. Before deciding which packages the policymaker will be able to decide on, the exogenous parameters within the model are listed. This is done in the first column of Table 8.1.

Following the system diagram in Figure 8.1, four main avenues are explored for the policy packages: an expertise policy package, a public perception policy package, a resource allocation policy package and an investment level policy package. Each of these packages acts on a different set of exogenous factors. This is also outlined in Table 8.1. Policy packages are incremental, similarly to the previous two models. Opposite sides of the same packages are also provided as a choice to the policy agents to choose from. They can introduce a policy or remove it from the system. Overall, this means that eight policy packages are considered plus the status quo package where no changes are introduced. This would be policy package 8. As an example, the introduction of the policy package that increases the expertise - policy package 0 - translates to an increase in ageing time, an increase in obsolescence time and a decrease in design time.

Other methods to generate policy packages have been demonstrated in the past (Taeihagh et al., 2009). They could ultimately be introduced in this model in the future but are currently beyond the scope.

8.2.3 Practical integration

The practical integration of the flood safety model is more complicated than for the two previous models. This is because the flood safety model is a system dynamics model, and it was constructed in Vensim. Previous models were agent-based models that were built in Python in each case.

There are different approaches to integrating a Vensim system dynamics model and a Python agent-based model. The most common approaches are to either build a wrapper in Python that will run the Vensim program from Python itself or to convert the Vensim model into Python and run everything as one Python program. For this paper, the second approach was used. The Vensim program was converted into a Python code using the PySD package[1]. This approach could be used because the Vensim model was simple enough to be converted into Python without significant issues.

Then, similarly to the previous approaches, the two models were considered as functions, each with a set of inputs and outputs. It then becomes an issue of building a script that will connect both models. This was done while the policy process model script remained the same as the one used for both the predation and the Swiss electricity market models. This was maintained for reproducibility and to show that the policy process architecture used is versatile. This script was built to include the initialisation of both models, the calculation of the states and actual beliefs from the technical model, and the policy packages implementation. It also includes a small part used to run a cycle of experiments for different inputs, to record all the data produced by both models and to experiment with different external events. The full model is accessible publicly on Github [2].

8.3 The research questions

The flood safety model is used primarily to demonstrate that the policy model can be simulated with agent-based models and system dynamics model. The primary goal is to show that the model developed for this dissertation is versatile and was built in such a way that it is easily reusable. This can be shown by the simple simulation of the flood safety model with the policy process model.

The present chapter goes further and exploits two parts of the models. First, it looks into the impact of the forecast for the policy agents on their decision making. This is

[1] https://github.com/JamesPHoughton/pysd
[2] https://github.com/kleinrap/policyemergencev4_SM_v2/tree/master/4_LeveeModel

	Issues					
	PC1 IP [-]	PC2 Safety [-]	S1 SLS [m]	S2 OLS [m]	S3 SLL [km]	S4 OLL [km]
Agents						
Policy maker	12	0.85	60 000	60 000	10 560	10 560

Table 8.2 – Preferred states for the policy maker on a the interval [0,1].

done by varying the time horizons the agents consider when they are estimating the impact of the policy packages. The question that this chapter aims to answer for this is the following: *What impact does a different foresight have on policy change?*

The second aspect that is investigated is directly related to the flood safety model. The model allows for the creation of different types of flooding as external events. It can be used to test the policy agents' reaction to different flooding events. The research question is formulated as follows: *Do different flooding events lead to a change in the policy packages that are implemented by the policy agents?*

8.4 The scenarios

Following the research questions formulated above, two sets of experiments were designed to test the hybrid model, each relating to the research questions. First, a benchmark scenario is simulated. This is done to establish a baseline for the simulations with no floods, and it is run without a policy agent.

The following scenarios take into account both the policymaker foresight and the flood events in combination. The policymaker foresight is varied from short term to long term. For this, the policymaker is given knowledge of the impact of the policy packages either five years ahead or, in the other case, fifteen years. The second parameter varied is the flood event in the system dynamics model. For this, either a single significant flood event is introduced or a series of smaller flood events every three years over a period of twenty years are considered.

For all simulations, the policy agent is provided with a set of preferred beliefs that are outlined in Table 8.2. Its causal beliefs, helping them understand the inner working of the flood safety model, are provided in Table 8.3. Once again, it is assumed that there is only one policymaker considering the policy process model used.

	PC1	PC2
-S1	0.00	0.50
-S2	0.00	0.50
-S3	0.75	0.00
-S4	-0.75	0.00

Table 8.3 – Causal beliefs for the policy maker. They are all given on the interval [-1,1].

8.5 The simulations

The hybrid model was once for each combination of a flood event and foresight value for the policymaker. It results in 4 simulations in total. Because all parameters are set - there is no randomness in the flood safety model and a lack of agent interactions in the simplest implementation model, no repetitions are considered. All experiments were run for twenty years, where the system dynamics model was run at a time step of 0.0078125 years. The policy process model was run with a one year time step starting at the end of the first year. This allows for a maximum of 19 policy packages to be implemented throughout the entire simulation. The benchmark simulations, which do not include the policy process, were also run for the same amount of time as a reference.

8.6 Results

Figure 8.3 displays a comparison between the simulations for both the single flood (left) and the pulse flood (right). At the top, it shows the levee length for all three levee stocks (designed, standard and old). The second line of graphs presents the safety measure for the levees themselves. This is one of the metrics that the policymaker has in its belief system and that it is attempting to change through the implementation of policy packages. In the third line of graphs, the two safety parameters, official safety and perceived safety are plotted. Official safety is calculated technically while the perceived safety depends on how safe inhabitants feel, which is a function of the amount of land flooded after each flood. Finally, the lower graphs show the choice of policy package by the policymaker. All of these results are in turn shown for both the five years and the 15 years foresight parameters.

First, the selection of policy packages is looked at. The different packages were presented in Table 8.1. One can compare the selection of policy packages between the two different flood events for the five years foresight case. In the single flood event, first, a range of packages is selected with policy agents tending to select contradicting

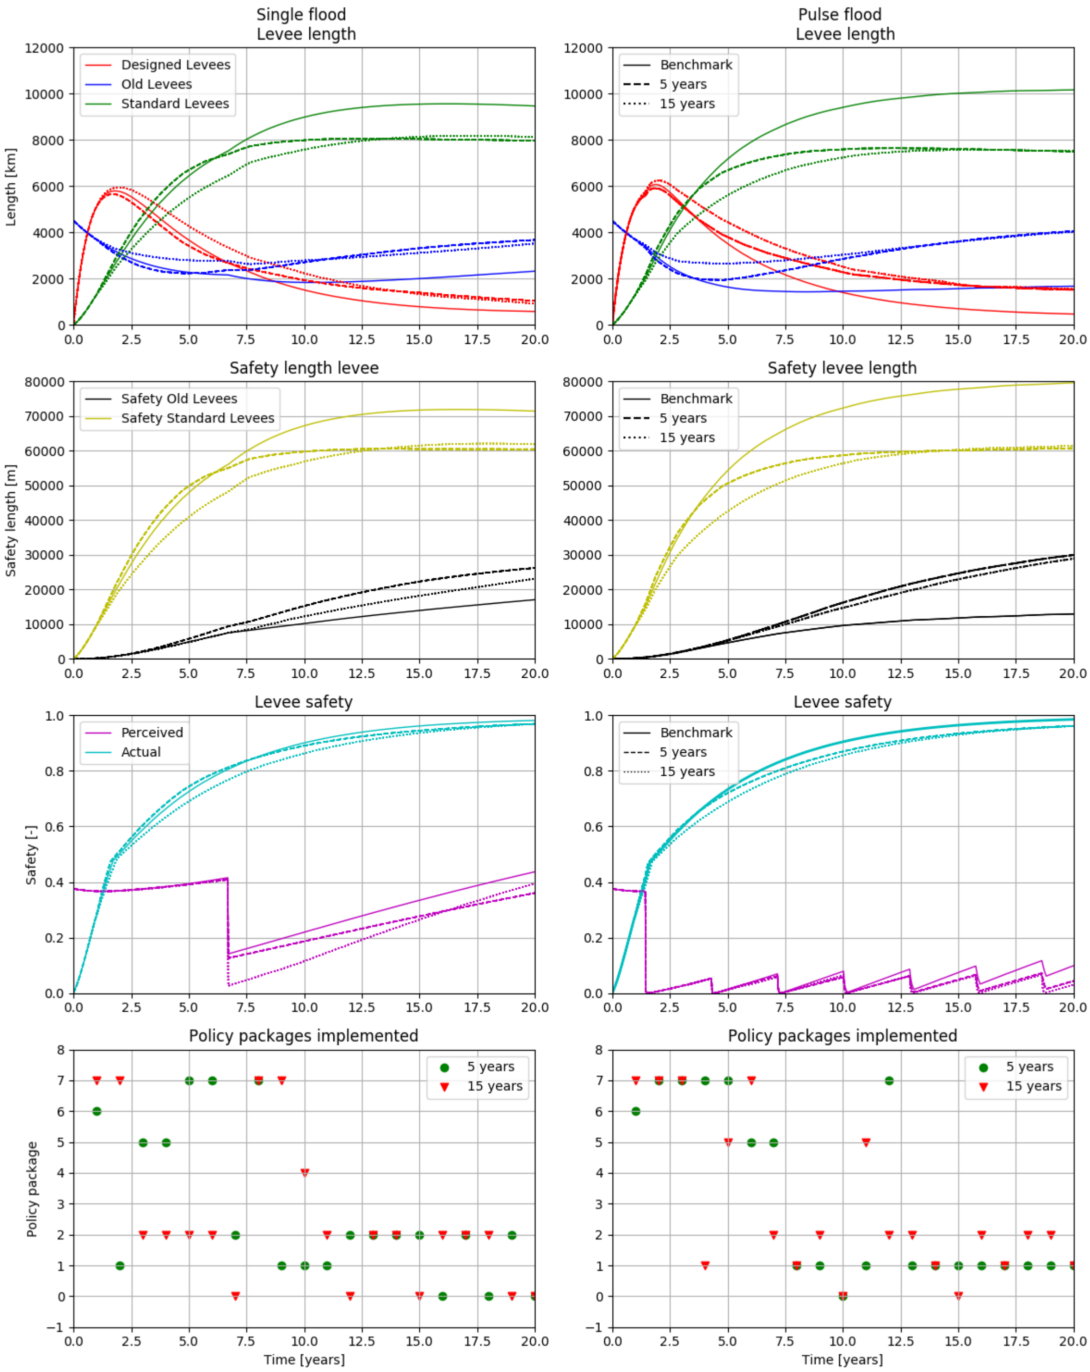

Figure 8.3 – Outputs from the flood safety hybrid model showing the levee length, the safety parameters and the policy choices for both flood scenarios.

packages in the form of package 6 and 7 (one increasing and the other reducing investment levels). When the floods strike, a change happens in both event situation. For the single flood, policy package 2 (increase in public perception) is highly favoured along with policy package 1 (decrease in expertise). For the pulse flood, it is policy package 1 that is solely selected. This means an increase in ageing and obsolescence time while reducing the design time. Ultimately, one can see that the outcomes are very close when it comes to levee length and actual safety, perceived safety still being affected by the pulsing flood.

The main differences between the different foresight levels can be seen not only in the policy packages chosen - they are slightly different, but mostly in the trends for the levee lengths. In the case of a short foresight, the results initially track closer the reference curves before diverging towards their end states. Ultimately though, the end states for the two foresight are mostly the same. This can be linked to the preferred states of the agent. Overall, the trend differences between the two foresight values can be attributed to the fact that the agent, acting based on a 15 years foresight, will select packages that achieve the goals the fastest. Note that though the agent has foresight, it cannot account for potential flows. Though this hardly makes a difference in the present results as the floods have little impact on the overall levee lengths and perceived safety. Additionally, one can see through these results that the number of times the agent can decide on policies also has an impact on the outcomes of the model. Differences would be observed if the policymaker would only be able to decide on policies every two or even five years instead of the current yearly process.

One aspect that can be observed from the results is the impact of the preferred states from the agent on the system. The agent has contradicting goals, wanting to have a very high amount of both old and standard levees, about 12 000 km each, and for the safety levee length where the agent wants to reach, for both old and standard, a value of 60 000 m. This has a direct impact on the system dynamics with a reduction of the standard levees and a net increase of the number of old levees. Similarly, for the safety length, a significant decrease in the standard levees and an increase for the old levees is observed, all in line with the preferred states of the policymaker. One can see the difference with the reference curve (full line in the graphs). Overall, this shows that the policymaker is attempting to reach its preferred states, though it cannot reach them all at once.

8.7 Discussion

One of the reasons why the policy agents cannot reach their preferred states is the fact that these are not realistic. The other relates to their causal beliefs. The causal beliefs of the agents are built with the assumption that there is a linear one-way causal relation between the secondary issue and the policy core issues. This is a simplification of the model and is not valid in the system dynamics model. This also has an impact on the policy packages that the agents chose as they select the ones that are more likely to have them achieve their policy core preferred states. This was already mentioned for the predation model.

Finally, note that this demonstration has one major limitation: a lack of repetitions. This was avoided because of the stated goal of this chapter. In a complete study, one would have to use a Monte Carlo approach with the flood safety model, varying the initial parameters such that a more representative sample of results is obtained. It would better account for random ranges that might happen in the system. This was done for the study of the flood safety model alone in a prior study (Klein et al., 2016).

8.8 Conclusions

Overall, this chapter has shown that it is possible to use the policy process model presented in this dissertation with a system dynamics-based flood safety model. This was just the first step to demonstrate the versatility of the policy process model. Further work focused on flooding, and its effect could be performed using this very model. More complexity could also be introduced in the policy process model to reflect the political composition of a specific country and how the policy agents might realistically react to the developments in the flood safety model.

This chapter also completes the presentation and demonstration of the simplest implementation model. This consisted of the development of this policy process model and its use with three different environment models showing its use for theory exploration, to deal with policy-related uncertainty and its versatility. However, the model suffered from limitations as a result of its lack of complexity. The third part of this thesis attempts to deal with these limitations by presenting a more complex policy process model based on the advocacy coalition framework.

The ACF implementation model Part III

9 Conceptualisation

The second implementation of the common language is made to demonstrate that it can be used to simulate a specific theory. In the present case, this is done for the advocacy coalition framework or ACF. Several new concepts need to be introduced to the model for the ACF, such as policy learning or coalitions. This third part of the dissertation is structured in the same way as the second part. First, a conceptualisation of the model is presented in this chapter. The next chapter presents its formalisation. Finally, the two chapters after that focus on the simulation of the ACF implementation model with two example environment models - the predation and the Swiss electricity market model.

This chapter presents the first step: the conceptualisation of the ACF implementation model. The difference between the simplest and the ACF implementation models is found mostly within their respective complexity. The ACF model is built using the simplest implementation model as its foundation. Being a complex framework, the ACF includes several additional concepts. To deal with this complexity without being overwhelmed by it, this implementation is approached in stages. Starting from the simplest implementation model, four layers of complexity are introduced based on specific concepts present in the framework. It consists first of the introduction of actor-on-actor interaction mechanisms to emulate the concept of policy learning. It is followed by the introduction of coalitions and their associated interaction mechanisms. Then two add-ons are considered, changing the fully rational actor assumption to the assumption that actors are bounded rationally. This includes the introduction of partial knowledge on the one hand, and imperfect relation on the other, making the actors entirely rationally bounded when both add-ons are considered at the same time.

The work presented in this chapter is part of a paper that is in the process of being published.

9.1 The simplest implementation model

Before introducing the new concepts for the ACF implementation model, this is a quick reminder of the simplest implementation conceptualisation following the buildings block provided by the common language. For the time building block, a one-stage policymaking process was considered, with a policy formulation stage where the actors decide on a policy instrument to be implemented. The policy arena corresponded to the environment within which the policy process is endogenised. This is a case-specific aspect of the model.

The only actor role that was considered at this stage was the policymaker role. The policymakers decided the policy instruments that were to be implemented. The actors were also all provided with a belief system. This belief system informed their decision making to select policy instruments.

Finally, the interactions introduced in this implementation were limited to the interactions between environment and actors. The environment informed the actual beliefs of the actors while the actors could implement a policy instrument that would influence the simulation of the environment.

9.2 Emulating policy learning

Policy learning is chosen as the first level of complexity as it is considered to be the most basic addition that can be found within the advocacy coalition framework. This means that it requires the less amount of added complexity to be integrated within the model. This will become clear with the elements added to the model in this section.

Policy learning is presented in the framework as one of the pathways to policy change (Sabatier and Weible, 2007). It is defined as a "relatively enduring alternations of thought or behavioural intentions that results from experience and/or new information and that are concerned with the attainment or revision of policy objectives" (Jenkins-Smith and Sabatier, 1999). In the model, this will approached as a change in the actors' beliefs as a result of interactions with the environment or with other actors. This can include changes in their actual beliefs, their causal beliefs, or their preferred states.

To emulate policy learning, modifications and additions to the simplest implementation model are made for the common language's time, actors and interactions building blocks. As policy learning relates to a change in the actors' beliefs, the main element missing from the simplest implementation model is the interactions between the actors. To allow for these interactions, three new concepts are introduced: the policy entrepreneur role, actor-on-actor interactions and actor resources.

Before these concepts are added, the time building block is completed. A more conventional two stages policy process is considered with the addition of an agenda-setting step before the policy formulation. In this step, the policy actors focus on the creation of an agenda that will narrow down the issues that can be discussed in the second step, the policy formulation. The discussions in the agenda setting step are limited to the policy core issues while for the policy formulation, they are restricted to the secondary issues.

The agenda created is composed of only one policy core issue that is selected through a majority vote by the actors. If no majority can be established, then the model continues without the consideration of the policy formulation step and therefore without the implementation of a new policy instrument. This is justified by the fact that the actors cannot agree on what issue to discuss. If an agenda had been created, in the policy formulation, the actors will only be able to discuss issues that they consider to be directly related to the policy core issue. On the basis of this discussion, they will then have the possibility to select a policy instrument and implement it, similarly to the simplest implementation model.

This second step is added to increase the complexity of the model such that more insights can be gained on the issues being selected. This can include knowledge on what issues are considered polarising for the actors and in which cases they are more likely to be present. The selection of an agenda setting step is based entirely on the multiple streams framework (MSF) despite considering this additional step for an ACF implementation (Kingdon, 2014). It also helps better include the concept of actor limited attention (Durant and Diehl, 1989). It was preferred to use an existing set of concepts, even if from a different framework, rather than using unvalidated concepts.

When it comes to the actor building block, besides the policymaker role, the policy entrepreneur role is introduced with the goal of emulating policy learning. The role of the policy entrepreneur is one purely of influence, exercised through actor-on-actor interactions - detailed later on. Policy entrepreneurs have their own interests, based on their belief system. This belief system follows the same architecture as the one described in the common language. Throughout the policy process, policy entrepreneurs attempt to influence policymakers and other policy entrepreneurs such that their interests are addressed, and their goals are met. They do this by pushing for specific policy instruments, specific issues or by helping shape an agenda that aligns with their interests. The assumptions is made that policy entrepreneurs have a say in the creation of the agenda, along with the policymakers. However, they have no say on the selection of policy instruments where only the policymakers have this power.

The last concepts that are introduced relate to the actor interactions building block. They include actor-on-actor interactions accompanied by the concepts of conflict level and resources. Actor-on-actor interactions are introduced as the main way to facilitate policy learning. Actors can influence one another by targeting other's preferred states and causal beliefs. In effect, this means that the actors can influence each other's goals or their understanding of how the environment works. Actors cannot influence each other's actual beliefs because, at this level of complexity, I assume that the actual beliefs are linked directly to the state of the environment under the assumption that actor have perfect information. Actors can only perform one influencing action at a time because of their limited attention (Durant and Diehl, 1989).

To decide which action to perform, actors consider all the actions they can perform. They then select the action that is likely to have the most impact, an assessment that is affected mostly by the average conflict level between actors amongst other things - actors with an average level of conflict are more likely to interact than when they have a high or low level of conflict, according to the literature (Henry, 2011). In reality, actors would be unlikely to consider all of their possible actions, they would not have time to do so. But, baring any specific literature that outlines exactly what actions actors consider, it was decided that for the present model it is easier if all actions be considered.

To limit the number of interactions and to weigh them depending on the actors' prominence in the policy subsystem, the concept of resources is introduced. Resources are spent by the actors when they want to influence other actors. After an actor has chosen the action it wants to perform, this action is implemented using a part of the actor's resources. Actions are performed until the actor runs out of resources. If an actor has run out of resources, it is unable to continue influencing other actors. These resources are replenished every round of the policymaking process. Only one type of resource is considered, contrary to the literature (Sabatier and Weible, 2007). I assume that all resources mentioned in the literature, from political resources to public opinion, are aggregated into one resource type. This is a quantitative resource that, if higher, will lead to a stronger ability to influence others.

In the different stages, the actors' influences are limited to issues in specific tiers of their belief system. For the agenda setting, this influence is focused on the policy core issue tier, while for the policy formulation, the influence is limited to the secondary issue tier. Furthermore, I assume that the actors have a limited attention span and, therefore, can only focus on one issue at a time for each stage. This is an assumption that comes from the common language and follows from the punctuated equilibrium theory (Baumgartner et al., 2014). The issue chosen is the one that they consider the most urgent - the one they assign the highest preference to. It is the issue where the difference between actual belief and preferred state is largest. To further their interests, actors will attempt to convince other actors that the issue they have chosen is an issue of importance.

The philosophy used to introduce policy learning is the same as the one followed for the common language and the conceptualisation of the simplest implementation model. For modelling and demonstration purposes, the goal is to build a model that is as simple as possible while satisfying the main requirement: simulating policy learning within the ACF context. In this thesis, policy learning is defined as a gradual change in the actors' belief system parameter. The goal is then to do this by adding the least amount of concepts and parameters. This is needed to retain a model that is manageable and that can be tested. It also provided a second demonstration of the use of the common language.

This approach does not preclude future researchers from diving deeper into the theories and to include more details in their models. For certain elements of the ACF, many more details can be found in the literature. Relating to policy learning, the literature has attempted to understand exactly what policy learning is and to differentiate different types of learning (Dunlop and Radaelli, 2018). The present model does not exclude the use of a more in-depth approach where parameters such as problem tractability and certification of actors are introduced into the model to try to extract behaviours that exhibit different types of learning as is hypothesised in Dunlop and Radaelli (2018).

Regarding the actors, the approach followed was to introduce only one new role for similar reasons. The role of the policy entrepreneur can be found throughout the literature but with varying degrees of importance depending on the theory. Furthermore, it can also be linked to the policy entrepreneurship literature (Cohen, 2016; Mintrom and Vergari, 1996; Mintrom and Norman, 2009; Mintrom and Luetjens, 2017; Oborn et al., 2011). In this literature, the role of the policy entrepreneur is further detailed. The present model attempts to reconcile these details with the belief system architecture present in the common language into a coherent, yet minimalist, model. The result is a set of actions that can be linked to actions mentioned in the literature. This includes for example the assumption that policy entrepreneurs display social acuity, define problems, build teams and lead by example (Mintrom and Norman, 2009). Although this is not explicitly spelled out in the present models through specific actions of the actors, these are underlying assumptions to the actions that the actors can perform.

Finally, the concept of resources is brought into the model with the explicit requirement to provide a limit on the number of actions that the actors can perform. To obtain a clear model, this requires that the resources mentioned in the literature be aggregated into one (Sabatier and Weible, 2007). Future work, focused on the study of the actor resources, could differentiate between different types of resources but this is not the subject of the present model. Indirectly, the inclusion of the concept of resources also brings in elements relating to actor power. Through their resources, certain actors are able to perform more actions than others, in effect providing them with more power over other actors. The introduction of the resources therefore also opens the door to the study of power concepts but this is out of the scope of the present model and could be the topic of future models.

9.3 Adding the coalitions

The second level of complexity comes in the form of coalitions. This is one of the key concepts of the ACF required when attempting to emulate the framework. Their introduction has an impact on the common language's actors and interactions building blocks. Throughout the literature coalitions are prominent (Calanni et al., 2014; Matti and Sandström, 2011; Weible et al., 2018; Zafonte and Sabatier, 2004). However, their functioning is not detailed enough making them hard to operationalise within the context of a model. This is especially true from a quantitative point of view.

To deal with these specificities, the approach taken here is one of simplicity similarly to the previous models. The goal is to find the simplest way possible to include coalitions within the policy process model while considering the already present architecture including the actor roles and the belief system. This means reconciling the theories with the parameters already present in the policy learning model while introducing the least amount of new parameters to emulate coalitions. This approach is only one of several approaches that can be considered. Like for the previous models, different researchers, in different contexts, either political, institutional or geographical, or with different goals for their studies, might interpret the literature and the different context-specific studies differently, leading to different choices and assumptions on how to model and simulate coalitions. This might include the introduction of more parameters, different operationalisation choices (the way the thresholds are quantified), or follow more closely specific key papers in their more detailed assumptions.

For the model, I assume that coalitions are made of actors that coalesce based on a shared preferred state for a salient policy core issue and that pool their resources and coordinate their actions with the goal to advance their interests. This assumption follows the literature which states that actors within a coalition should have "perceived belief correspondence" and "coordinate their actions to achieve political goals" (Matti and Sandström, 2011). When they assemble, I assume that coalitions effectively act as super-actors and more specifically super-policy entrepreneurs. This follows from their need to achieve political goals (Matti and Sandström, 2011). It also helps limits the addition of new parameters while reconciling the already present model with the theories. To further limit the model complexity, their influencing actions are the same as the policy entrepreneurs as described for the policy learning model. The main difference with other actors is the additional resources at their disposal.

Such influence is aimed at furthering the coalition's interests - their political goals. For example, this can be achieving their preferred state for the issues around which the coalition has been assembled. This is done through the implementation of policies supported by the coalition. With this in mind, I assume that each coalition can influence either actors within the coalition itself to increase its cohesion or, actors that are outside the coalition to implement policies aligned with the coalition's interests. These coalition cohesion actions are referenced throughout the literature. Ingold et al. (2017) refers to it as within-coalition coordination. In Jenkins-Smith et al. (2014), it is explained how a change in the cohesion of the coalition might affect the effectiveness of a coalition. Further, I assume that depending on the coalition cohesion - that is how close in beliefs the members are from one another, the coalition will assign a different amount of resources for either cohesion building influence or outside influence actions.

Using the architecture of the model presented for the policy learning complexity increase, the coalitions' influence can be exercised on the preferred states of other actors or their causal beliefs. This is similar to the policy entrepreneurs. At this point, the issues of implementation relate to the operationalisation of these actions. This is due to the fact that the literature does not detail who exercises this influence within the coalitions and how the actions are selected. When they do, it is often context dependent. Similarly to the other actors, and in line with the common language architecture, I assume that each coalition has a belief system. This belief system is populated with values for the actual beliefs, the causal beliefs and the preferred states that are calculated based on the averages of all of the coalition's members values. The preferences are calculated from the resulting values to define which issues is the focus for the coalition on each tier of the belief system.

The coalitions' influence actions are based on their belief systems, the same way it is for all actors. The actions are rated first by considering the conflict level, and the actions are performed using coalition resources. For cohesion actions, the actors considered are the members of the coalition. For external influence actions, the actors considered are all of the actors in the policy subsystem that are not coalition members. This assumption means that there is no need for a leader or key figure in the coalition. The complication with having a coalition leader - an alternative approach - would be to define how a leader can be systematically selected based on information already present in the model, such as the belief system or the actors' resources. By using this averaged belief system approach, the need for additional assumptions and parameters can be avoided, making the model less complex.

The goal with this approach for the coalitions and their actions is to emulate the effort of coordination described in the literature (Calanni et al., 2014; Weible et al., 2018). As mentioned before, the literature goes in a lot of details regarding the coalitions but these are often case specific details which cannot always be used in a general model such as the one described here. For example, Weible et al. (2018) explores coordination in the context of high-intensity policy conflicts in New York, Colorado and Texas. The study shows that different parameters are important in such situations for collaboration and coordination between the actors compared to the ones used in low-intensity conflicts. All of these details could be included in a model to study high-intensity policy conflicts but this is not the goal of the present model. The goal here is to show that it is possible to include coalitions within a model and to simulate them.

Furthermore, even if one were to go into the details of the literature, it would still lack specific details that can be used for the operationalisation of the coordination of coalitions. Staying with the example of Weible et al. (2018), the study does not tell us at what specific threshold actors need to attain to coordinate as the study remains qualitative. Qualitative assessments are difficult to operationalise and, in the present case, operationalisation requires hard numbers to determine when actors do what and based on what parameters. Finally, there is also a need to reconcile the theory with the belief system architecture that is used in the present model. In the literature, the term of beliefs is used at large when looking at coordination without specifying or defining exactly what a belief is. A precise and clear definition is required for modelling.

Elements such as the policy network also play a crucial role within the literature (Henry, 2011; Park and Rethemeyer, 2012). The shape of the network is often influenced by the trust between the actors and can influence the creation of coalitions, along with the resulting interactions throughout the policy subsystem. The collaboration between actors can also relate to the professional competence of actors or their need to acquire more resources (Weible et al., 2018). However, for the sake of simplicity, and because an explicit network is not needed to emulate coalitions, such network was not explicitly included in the model. Despite this omission, the model still contains an emergent policy network based only on the interactions that agents have. This network can be tracked throughout the simulation. This network is however not based on trust as the literature suggests but solely on beliefs. Future models could include the policy network more explicitly if it was the subject of the study for example.

9.4 Including partial knowledge

Up to now, the introduction of concepts to emulate policy learning and the addition of coalition have assumed that actors have perfect knowledge under the assumption that the actors are fully rational. This means that they know perfectly what the other actors' beliefs are. Their actions are informed by such knowledge and, therefore, are perfectly accurate. However, in real life, such situations are unlikely and actors are more likely to be partially rational - if not entirely irrational. There always remains a degree of uncertainty on another's beliefs. Ingold et al. (2017) mentions how what they call "mutual knowledge" can influence understandings and can help actors reach agreements.

Partial knowledge can be introduced in the model to account for this uncertainty. This has an impact mostly on the common language's fourth building element: the interactions. Note that this additional level of complexity can be added to both policy learning or coalitions, while the coalitions cannot be considered without the policy learning introduction. Therefore, it is approached as a model add-on.

By considering partial knowledge, it is assumed that actors only have approximate knowledge of other's beliefs. The only way to gain more knowledge of other's beliefs is to interact with them. Introducing partial knowledge means that actors will plan their interactions with a certain amount of uncertainty, leading to potentially ineffective influencing interactions when their knowledge is inaccurate. For example, an actor might try to influence another actor's causal beliefs while both actors have the same causal beliefs. This results in a waste of resources for the first actor and a gain of knowledge of the target actor's beliefs, meaning the mistake is less likely to be repeated.

Note that throughout the literature, the concept of trust can be conflated with the concept of partial knowledge and how it is approached in this section. Trust is often the key concept used to define how actors interact throughout the policy subsystem and that helps defines the policy network (Berardo and Scholz, 2010; Leach and Sabatier, 2005). It can also sometimes encompass the concept of partial knowledge in the way it is defined here.

9.5 Considering imperfect information

Finally, it has so far been assumed that information flowing from the environment into the policymaking process model is perfect. This means that the actors' actual beliefs reflect precisely what is happening in the environment. For example, for a policy core issue of crime, if there are 300 crimes per year, this will be what actors consider when selecting policy instruments. However, in reality, actors often receive imperfect or distorted information, or not trust the information that is communicated to them (Henry and Dietz, 2011). This information might be coming from experts, from other actors within the policy subsystem or from the media (Adams, 2004; Arceneaux et al., 2016; Birkland and Lawrence, 2009; Dela Santa, 2013; Elmelund-Præstekær and Wien, 2008; Mutz and Martin, 2001; Weible, 2008). Using the previous example, this difference in information might lead actors to make their decision based on the number of 300 crimes per year while others, having obtained different information, might decide based on 200 or 400 crimes per year. This will impact the actors' decision-making process. Such aspect of imperfect information can be taken into account in the model.

This is done through changes in the common language's actor interactions building blocks. Note that similarly to the partial knowledge inclusion, considering imperfect information can happen with either policy learning, coalitions or even the partial knowledge having been introduced to the model. Therefore, it is also approached as an add-on to the model.

Conveying information from the environment to the policy actors in the policy process model is now done by actors. For this, a new actor role is introduced: the external party. External parties are actors that can obtain the real states of the environment and have for role to inform the policy actors of these states. One can think of external parties as actors representing experts, the media, international institutions, academic institutions or NGOs, for example (Adams, 2004; Arceneaux et al., 2016; Weible, 2008).. Transmitting the states is something they do, influenced by their own biases. They will, for example, transfer information differently if they believe that that information is correct or not. The effectiveness of this transfer will also be affected by the policy actors' political affiliation they inform. For example, trust might not be present between actors of different political affiliation, leading to a lack of information transmission if any. In turn, this imperfect information that the policy actors obtain will impact their decision making as their actual beliefs will have distorted or incomplete information.

This might result in changes to their preferences or policy instrument selections.

Beyond this information transformation, it is assumed that external parties have the same potential for actions as policy entrepreneurs. Throughout the agenda setting and the policy formulation, they can influence other actors on their preferred states and their causal beliefs. One difference is that such interactions are performed on all actors simultaneously, following the assumption that external parties have the means to influence all actors at once instead of focusing their efforts on one actor at a time. For example, consider an external party being a political TV show. The message being delivered on that show will affect all of its audience, albeit differently depending on the variation in the audience's political affiliation. This is reflected through this mechanism.

Similarly to the previous complexity extensions, in this approach of the inclusion of imperfect information, trade-offs had to be considered. There was the need to reconcile the already present architecture of the model with the details of the literature while trying to include as little new parameters as possible. For example, the goal is to have a context-neutral approach for this model. This therefore excludes the explicit presence of roles such as the one of expert for example as they might have been presented in the literature (Weible, 2008). The selection of an external party role was preferred in this case. This approach does not however exclude the formulation of these same experts through the parameters presented above. In this case, experts, media and NGO could only be differentiated based on their resources and their beliefs retaining similar actions.

One of the issues present with information transfer is that it is heavily integrated with a range of other concepts within the qualitative literature and can span a wide range of literatures (Elmelund-Præstekær and Wien, 2008). This includes, but is not limited to the concepts of trust, type of policy subsystem or policy network. Note that not all of these concepts are present in the model outlined here because of the desire to keep the model simple and context-neutral. Furthermore, for modelling purposes, any concept introduced in the model is introduced alone as much as possible to simplify its study. Though this helps simplifying the model and the analysis of the results, it can also lead to the loss of critical behaviours. For example an adversarial policy subsystem where coalitions are very competitive might be more impacted by distorted information than a unitary subsystem (Weible, 2008). This is a behaviour that cannot be exhibited with the present approach.

9.6 Discussion

The ACF model proposed in this chapter is much more complex than the first simplest implementation model. This complexity is required because the framework is more complex it itself. To simplify the model exploration, and because all of the complexity might not be needed in every case study, the model was approached through four successive complexity levels. The challenge for each level of complexity was to account for the concepts already present in the original model, while integrating new elements as close as possible to the literature, and while keeping the model as simple as possible. The need for simplicity stems from the goal of the present approach which is aimed to be a demonstration of the power of the common language while being context-neutral. Simplicity makes it easier to understand the model, and analyse its results. This three-way challenge has lead to a number of trade-offs and assumptions that were detailed throughout this chapter and which are unfortunately not always be based on validated literature.

The presence of different complexity levels can be seen as an asset for the testing of the model, and for theory exploration in general. It grants researchers the ability to limit their model to only the amount of complexity they require for their models and gives them an example on how to do so. For example, in a model where the goal is to understand the effect of policy learning, the model can be limited to its policy learning complexity level, helping limit the complexity and making the model more understandable. Researchers should also be aware that this will of course come with limitations as the model might me behaviours that might otherwise be present in reality. All of this is demonstrated in the following chapters with the implementation of this ACF model with the predation and the Swiss electricity market models.

Using a tiered complexity approach also allows for an extensive set of experiments relating to the different levels of complexity. One can think of the large number of model combinations that could be implemented and tested. Avenues are proposed to test assumptions related to actor interactions, actor collaborations, actor strategies, policy learning, power dynamics, resources sharing, coalition behaviours or coalition creation rules. Additional assumptions testing related bounded rationality, and to the impact of partial knowledge and imperfect information on policy change are also in the cards. This significantly extends the possibilities brought by the use of this modelling and simulation approach to study the policy process and the ACF in particular.

Note that imperfect information and partial knowledge are introduced within this ACF model while they are not explicitly mentioned as part of the ACF literature. The reason why they are introduced is that, though they are not present as core concepts of the ACF, these concepts relate to bounded rationality and are inherently present in the framework even if not explicitly mentioned. They can play a crucial role in how coalitions and actors behave and are influenced. By incorporating them to this model, something that was not possible for the simplest model implementation because of the non-inclusion of actor interactions, the goal is to show that the use of modelling can widen the research scope significantly and allow for more in-depth studies of cases. It can also be used to test concepts that might be considered tangential to the critical policy process theories and that are often relegated to the background, but that play a key role in policy change.

9.7 Conclusions

This chapter has presented an implementation of the advocacy coalition framework using the common language. To simplify the model testing and exploration, this was presented through four different levels complexity increase to include policy learning, coalitions, partial knowledge and imperfect information. The model presented has a lot of potential when it comes to theory exploration and the different examples provided throughout this chapters have helped demonstrate that.

This chapter was limited to the conceptualisation of this model. This is insufficient for simulation. For that, a set of equations needs to be defined. This is done in the next chapter with the formalisation. Only after that chapter can the model be coded, experiments can be designed and the model can be simulated.

10 Formalisation

Following the conceptualisation of the ACF implementation model in the previous chapter, this chapter presents the formalisation of the model. This formalisation follows the structure used in the conceptualisation. It does not include the formalisation of the simplest implementation model from Chapter 5. However, all concepts and equations mentioned in that chapter still apply here. It includes the formalisation of the belief system for the agents or their preferences calculations.

The ACF implementation model is formalised for each additional level of complexity considered. This includes first the addition of policy learning, followed by the introduction of coalitions, the inclusion of the imperfect knowledge add-on and, finally, the formalisation of the partial information add-on. For each addition, the process is outlined, followed by the agents added and, finally, the new interactions.

10.1 Policy learning extension (+PL)

The introduction of policy learning (+PL) leads to the addition of policy entrepreneurs and several interactions between the agents. This is detailed within this section, starting with the new process that is considered for the model. The introduction of policy learning sees the addition of interactions that allow agents to influence one another on their beliefs.

10.1.1 The process

The policy process is now a two-step process. Each of the steps considered now includes several new interactions between the agents. First, the agenda-setting step is performed, then, if an agenda has been agreed on, comes the policy formulation step.

This does not include the environment simulation. The process is given below and illustrated in Figure 10.1.

1. Initialisation:

   (a) *Trigger of external events:* Any event that the modeller decides to implement are activated at this stage.

   (b) *Update of the truth agent:* Information from the environment is used to inform the truth agent's actual beliefs.

   (c) *Transmission of the actual beliefs:* The agents are informed about the environment from the truth agent. Their actual beliefs are updated.

2. Agenda-setting step:

   (a) *Resources distribution:* The resources needed for the actions are distributed to all of the agents.

   (b) *Preferences calculation:* Each agent calculates the preference for its deep core and policy core issues. Additionally, each agent selects a policy core issue that it will advocate for based on its preferences.

   (c) *Agent interactions:* Agents influence one another to advance their interests relating to the policy core issues of their choosing, spending all of their resources in the process.

   (d) *Preferences update:* Each agent updates its preferences before the selection of the agenda to take into account changes related to the agent interactions.

   (e) *Agenda selection:* The agents vote for a policy core issue to be placed on the agenda. If a majority is found, then a new agenda is formed, and the process can proceed to the policy formulation if not the policy formulation step is skipped and the status quo remains.

3. Policy formulation step:

   (a) *Resources distribution:* The resources needed for the actions are distributed to all of the agents.

   (b) *Preferences calculation:* Each agent updates its preferences for its secondary issues and policy instruments based on the issue on the agenda. Each agent then selects a policy instrument and a secondary issue that it will be advocating for.

   (c) *Agent interactions:* The agents influence one another to advance their interests relating to the secondary issues, spending all of their resources.

(d) *Preferences update (policymakers):* Policymakers update their preferences before the selection of the policy instrument to take into account the agent interactions.

(e) *Policy instrument implementation:* The policymakers vote to implement a policy instrument. If a majority is found, then a policy instrument is implemented, if not, the status quo remains.

Similarly to the simplest implementation model, an initialisation step is required. This step is used to link the environment model to the policy process model, to introduce potential external events that can relate to scenarios and to update the agents with information from the environment. This step is a computational step that would not be found in real life or in the literature.

10.1.2 The agents

One new active agent role is added to the +PL model as per the conceptualisation: the policy entrepreneur. It has the same attributes as the policymaker with the main difference being that policy entrepreneurs have no decision making power and therefore, can not be part of the policy instrument selection.

There are changes made to the active agents in general:

1. The *active agent* is now represented as an 5-tuple given by *agent = (ID, type, beliefHierarchy, advocacy, resources)*.

2. The *type* corresponds to the agent type: in this case, it can only be a policymaker or a policy entrepreneur. This will go on to define the actions possible for the agent.

3. The *advocacy* is now represented as a 3-tuple *(issue_as, issue_pf, policy_pf)* where *issue_as* is the policy core issue chosen during the agenda-setting step, while *issue_pf* is the secondary issue chosen and *policy_pf* is the policy instrument selected in the policy formulation step.

4. The *resources* is the currency that is used to perform actions. Resources are provided at the beginning of each step of the policy process.

The preference calculation for the policy core issues is the same as the one presented in Equation 5.1. Analogously, in the agenda-setting step, the agents will select one policy core issue to advocate for.

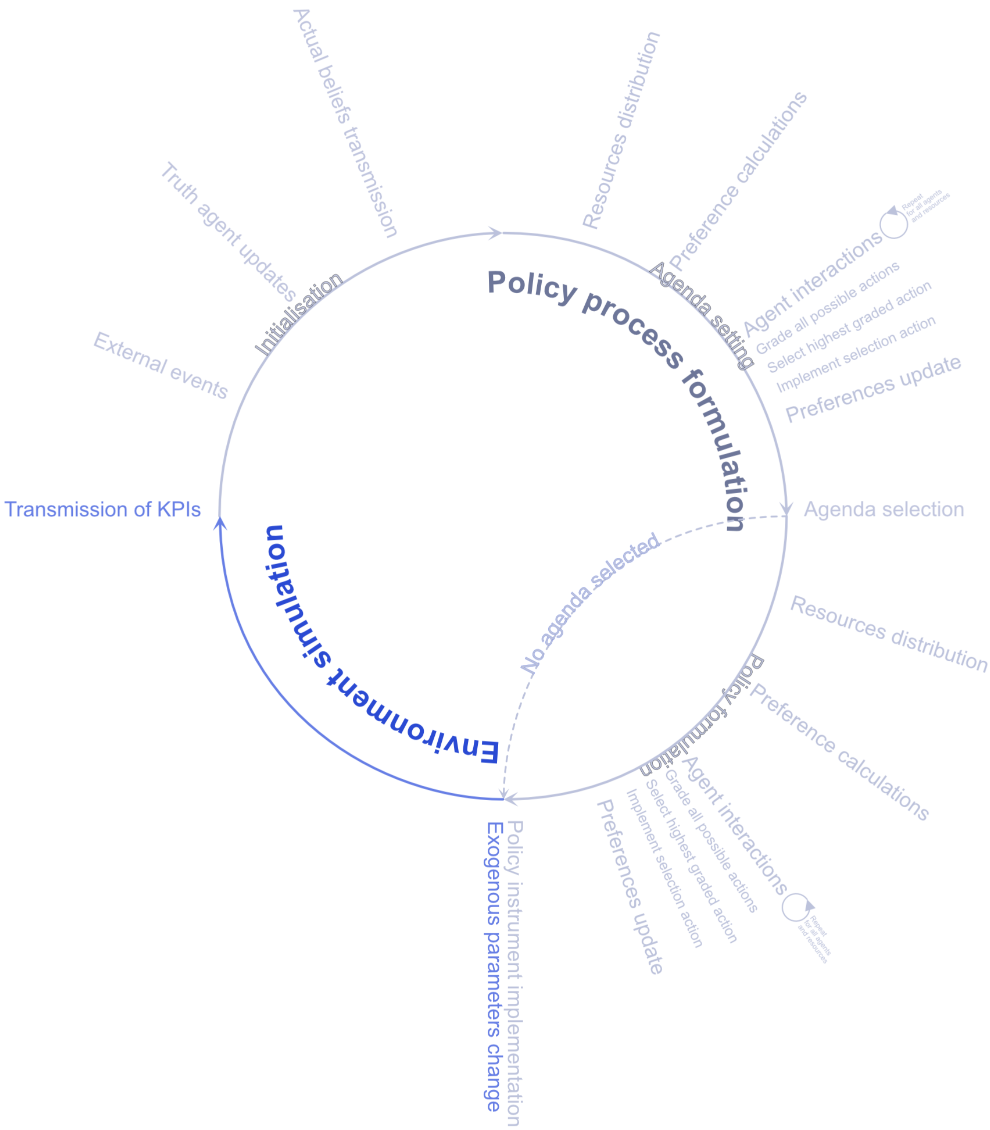

Figure 10.1 – Steps followed by the policy process for the ACF model with the first level of complexity +PL (with details for the agent interactions).

The architectures of the policymakers and the policy entrepreneurs agents are the same, only the actions they can perform are different.

10.1.3 The agent interactions

One type of interaction is added with the policy learning introduction: agent-on-agent interactions. These interactions aim to emulate policy learning, a key concept of the ACF. The rest of the interactions remain the same as the ones that were introduced in the simplest implementation model.

Throughout both the agenda-setting and the policy formulation steps, the agents can interact with one another. These interactions are limited to their causal beliefs and their preferred states. Agents only interact on the issues they are advocating for and their related causal beliefs. For the agenda-setting step, these interactions are limited to the policy core issue they prefer, and for the policy formulation, the interactions are on the secondary issue preferred.

To act, first, each agent considers all of the actions that it can perform, grading them. Once the actions have been graded, the agent will select the action that is expected to have the highest impact. It will then implement this action, affecting the beliefs of the agent targeted and spending a small number of resources in the process. This process repeats until the agent has spent all of its resources. In the present case, it is decided that agents are allowed to perform ten actions per step, meaning they spend 10% of their resources per action This is illustrated in Figure 10.1.

The first step is to grade all of the possible actions that an agent can perform. This includes actions on all agents that are present in the policy subsystem for each of the causal beliefs related to the preferred policy core issue, and on the preferred state of the preferred policy core issue. The grades are assigned based on the conflict level between agents for each issue. This conflict level is directly related to the difference in beliefs between the agents.

Following Henry (2011), it is assumed that agents with an average conflict level are most likely to interact, followed by a low and then high conflict level. These are established based on ranges of differences. For example, presently, a conflict level is defined as average when the beliefs of two agents are within 0.2 and 0.4 of one another. Low conflict level is then when the beliefs of the agents are within 0.2 of one another, and high conflict level is when there is a greater difference than 0.4. Practically, in the algorithm that is used to grade the actions, a low conflict level is assigned a grade of 0.50, for medium it is 0.75, and for high it is 0.25. To further differentiate the actions, this conflict level is then multiplied by the difference in beliefs between the agents.

An additional step is taken in the policy formulation step. As the policymakers are the only agents with decision making power in that step, it is assumed that agents will prefer influencing policymakers rather than other agents. A bonus to the grade of 10% is therefore provided for actions that influencing policymakers specifically, to favour such actions.

The equation that is used to calculate the grade is given as follows:

$$Gr_{j,p} = W_{CL_{j,p-n}} \cdot |\text{diff}| \cdot \text{Bo} \tag{10.1}$$

where Gr is the grade, j is the belief selected - preferred state or causal belief, p is the agent onto which the action would be performed, n is the agent performing the interaction, W_{CL} is the conflict level weight assigned based on the difference in beliefs, diff is the difference in beliefs and Bo is the policymaker bonus that, in the policy formulation step, is set to 1.1 when p is a policymaker and to 1 when that is not the case.

Once an action has been selected - the one with the highest grade, the influencing agent will know the specific belief and the agent that this action will be performed on. For a causal belief, this is done using the following equation:

$$C_{j,p} := C_{j,p} + (C_{j,n} - C_{j,p}) \cdot R \tag{10.2}$$

where C stands for causal belief, j is the causal belief selected, and R for the resources.

And similarly for a preferred state interaction:

$$G_{j,p} := G_{j,p} + (G_{j,n} - G_{j,p}) \cdot R \tag{10.3}$$

where G stands for the preferred state and j is the policy core issue concerned.

10.2 Coalition extension (+Co)

The second level of complexity increase for the ACF implementation model sees the addition of the coalitions. Such complexity increase can only be performed if the policy learning has already been introduced into the model. This is because coalitions can only impact policy learning through interactions of their own. The main element to be added when considering the coalitions is the coalitions themselves and all the interactions that surround them. This includes the way the coalitions are created, how they influence their members to keep them within the coalitions and how they influence agents that are outside of the coalition to advance their interests. This is detailed in this section.

10.2.1 The process

The introduction of coalitions adds a few sub-steps to the process described in the +PL model. This is detailed below and illustrated in Figure 10.2.

1. Initialisation:
    (a) *Trigger of external events*
    (b) *Update of the truth agent*
    (c) *Transmission of the actual beliefs*
    (d) *Coalitions creation:* The coalitions from the previous round are deleted and recreated based on the updated beliefs of the agents.
2. Agenda-setting step:
    (a) *Resources distribution*
    (b) *Preferences calculation*
    (c) *Coalition interactions:* The coalitions perform their coherence and external influence actions to advance their interest relating to their preferred policy core issue, spending all of their resources.
    (d) *Agent interactions*
    (e) *Preferences calculation*
    (f) *Agenda selection*
3. Policy formulation step:
    (a) *Resources distribution*
    (b) *Preferences calculation*
    (c) *Coalition interactions:* The coalitions perform their coherence and external influence actions one another to advance their interest relating to their preferred secondary issue, spending all of their resources.

(d) *Agent interactions*

(e) *Preferences calculation (policy makers)*

(f) *Policy instrument implementation*

10.2.2 The agents

The agents are the same as in the +PL model and include the previously added policy entrepreneurs. The only addition is that of the coalitions, which are considered as an active agent for this model. This can be justified by the fact that coalitions have mostly the same attributes as the other active agents in the model and they can also perform influencing actions.

The coalitions have the following attributes:

1. A *coalition* is represented as an 5-tuple given by *agent = (ID, members, beliefHierarchy, advocacy, resources)* where *ID* is the coalition unique ID, *members* is the list of agents that are members of the coalition, *beliefHierarchy* is the coalition's aggregate belief hierarchy build using its member's beliefs, *advocacy* is the list of the issues the coalition is supporting, *resources* corresponds to the number of resources the coalitions has.

2. The *belief hierarchy* is made based on the average of the member agents' belief hierarchies for causal beliefs, preferred states and actual beliefs. For all issues, the coalitions calculate their preferences based on their other beliefs, similar to all agents.

3. Similarly to the other agents, the *advocacy* is represented as a 3-tuple *(issue_as, issue_pf, policy_pf)*.

4. The *resources* are the resources pooled from the members. The modeller defines the amount the agents provide to the coalition. The resources tally is updated at the beginning of every step of the process.

10.2.3 The agent interactions

The introduction of coalitions leads to some additional interactions between coalitions and agents. This includes the systematic creation of coalitions and coalition interactions.

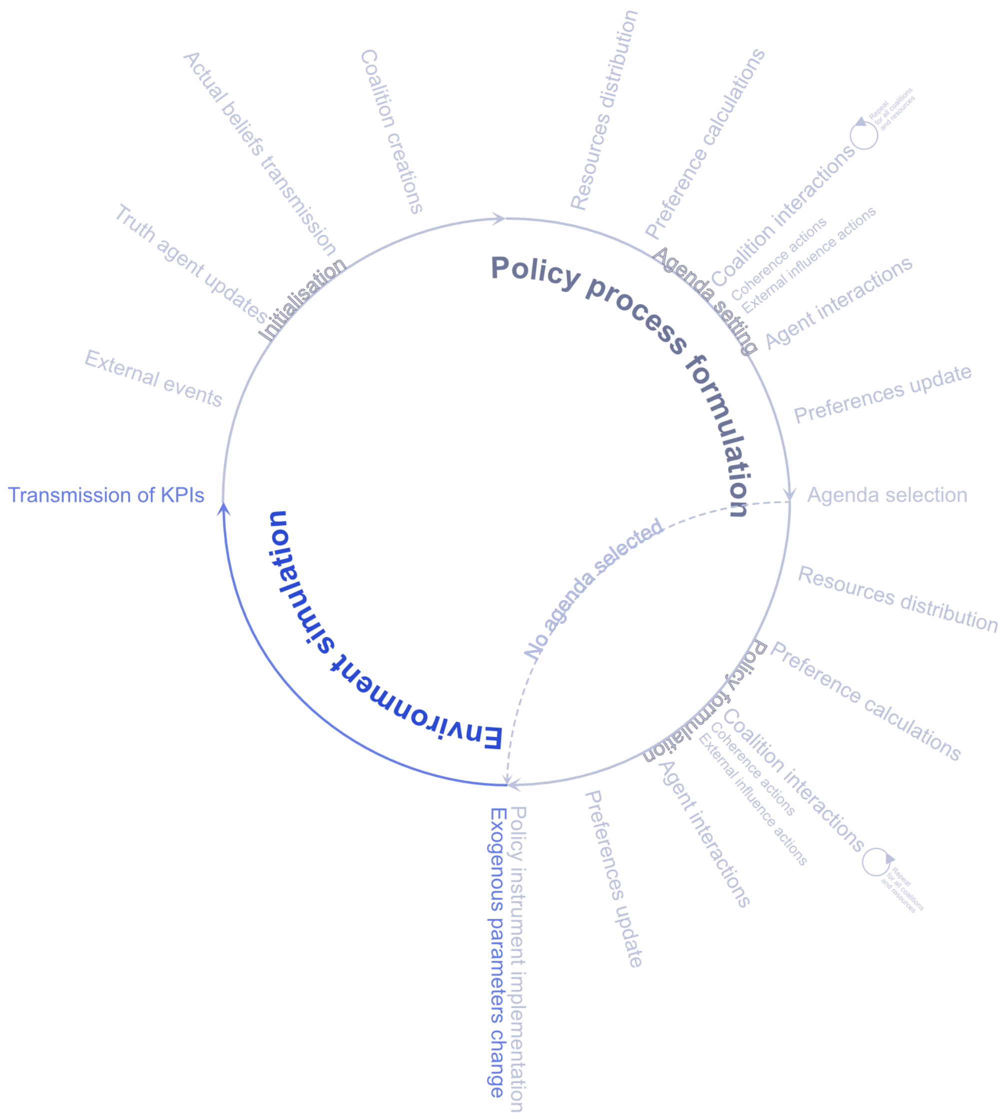

Figure 10.2 – Steps followed by the policy process for the ACF model with the second level of complexity +Co (with details for the coalitions).

Coalition creation

To be simulated, the coalitions need to be created systematically. For this, several criteria are considered to gauge which agents can assemble to form coalitions with the goals not to have a limited set of coalitions within one policy subsystem. The literature often cites three or fewer coalitions, and most often only two, present at any time in a subsystem. Furthermore, the coalitions should remain stable throughout most of the simulation.

Considering this, the assumption is made for coalitions to be created around policy core issues where agents have similar preferred states. Often one policy core issue is considered to be critical or salient within a subsystem, and this policy core issue does not change throughout the process. This issue is chosen before the simulation by the modeller. Agents are then grouped based on their preferred states for that salient policy core issue.

In the present formalisation, it is assumed that agents have similar beliefs if they are within 0.3 of one another - considering that the preferred states can have values between 0 and 1. Coalitions are created whenever several agents have their preferred states within 0.3 of each another. An algorithm is then designed to find the optimum points around which a maximum of two coalitions can be formed, under the assumption that any subsystem considered within this model will have a maximum of two coalitions. For cases where it is known that there are more than two coalitions, this requirement can be adjusted. Coalitions can only be created when more than three agents are grouped.

After a coalition has been created, its resources are calculated based on the number of resources that the agents have pooled together. Its belief system is established based on the average of its member's beliefs. Finally, note that coalitions are created for the entire process, and are therefore the same for the agenda-setting and policy formulations steps.

Coalition interactions

As was outlined previously, in the present formalisation, it is assumed that the coalitions behave like other agents, and therefore can perform the same interactions. The assumption is that coalitions, as a whole, will work to influence outsider agents to advance their interests. Besides this, coalitions can also influence their members as a way to maintain member cohesion. It is assumed that coalitions cannot influence entire other coalitions.

The coalitions spend a part of their resources to increase the belief coherence between their members. This consists of performing influencing actions their members that are considered to deviate too far from the coalition's goals. 'Too far' is defined as being higher than 0.2 from the beliefs advocated for by the coalition. This influence can be performed on issues' preferred states on both the policy core and the secondary levels. They can also influence their members on their causal beliefs, but this is considered less critical. The assumption is that coalitions will spend up to 30% of their resources working on increasing the coherence with their members. Such resources are spent in tranches of 5%, allowing up to six actions per step. If all the members are within the threshold of 0.2, then no resources are spent on such coherence actions.

The rest of the coalitions' resources are used to influence external agents in the same way that agents influence one another. Based on the conflict level, the coalitions grade potential interactions. They then choose the actions that they want to perform based on these grades and then influence other agents. Similarly to what was done for the +PL model, this includes actions on the causal beliefs and the preferred states only. Because of their considerable potential resources, the amount of resources coalitions can spend on any single action is limited to 5%.

10.3 Partial knowledge add-on (+PK)

The third level of complexity increase sees the introduction of partial knowledge into the model. This introduction has no impact on the process or the agents. It only impacts the agent interactions. This is explained below. Note that the +PK model can be combined with the +PL or +Co models. It can be used whenever agent influence interactions are considered. This is why it is termed as an add-on while being a complexity extension of the model.

An issue arises when selecting which knowledge the coalitions have. It is unclear how to transfer the partial knowledge of a coalition from round to round, as agents enter and exit coalitions and coalitions are updated every round. The assumption that is made to deal with this issue is to have the coalitions' partial knowledge be an average of all of its members' partial knowledge, following the approach used to calculate the coalitions' beliefs.

10.3.1 The agent interactions

The introduction of partial knowledge only leads to changes in the agent-on-agent interactions, changes that are detailed below.

Agent-on-agent interactions

Partial knowledge does not directly affect the actions between the agents. It affects how the actions are selected by changing their grades. As outlined previously, the conflict level is calculated based on the knowledge of the agents of other agents' beliefs. Previously it was assumed that agents had complete and perfect knowledge of each other's belief. When considering partial knowledge, this assumption is removed and agent grade the actions based on their knowledge of the beliefs of other agents.

The knowledge of other agent's belief is not static and can be erroneous. It changes as agents interact with one another. This is reflected in this model as well. Agents can update their knowledge of other agent's beliefs when they interact with them. This is a slow update that happens after each interaction and where both agents, the one influencing and the one being influenced, gain knowledge of the other's beliefs. The knowledge update is calculated using the following equation for a preferred state:

$$G_{j,p}^{pk} := G_{j,p}^{pk} + (G_n - G_p^{pk})W_{\text{gap}} \tag{10.4}$$

where G stands for the preferred belief, pk stands for partial knowledge, j is the issue selected, p is the agent who is updating his/her knowledge, n is the agent being gained knowledge from, W_{gap} is a weighing constant that can be used to scale the speed at which knowledge is acquired through interactions. This equation would be the same for a causal belief.

10.4 Imperfect information add-on (+PI)

The fourth and last level of complexity introduced is one where imperfect information is considered. This relates to the information that emanates from the environment and how it is relayed from the environment through the truth agent and to the policy agents. The consideration of imperfect information sees the introduction of a new agent role, changes in the process and the addition of new actions related to the new agent role. Note that this +PI model is an add-on that can be added to either model presented before, similarly to the +PK add-on. It can also be combined to the +PK add-on to consider a situation where full bounded rationality is present in the model.

10.4.1 The process

The main changes to the process after the introduction of the imperfect information add-on are changes that happen in the initialisation step of the model. This is mostly focused on the way the actual beliefs of the agents are updated from the environment. The rest of the process remains unchanged.

The steps for the initialisation are given below and illustrated in Figure 10.3.

1. *Trigger of external events*
2. *Update of the truth agent*
3. *Update of the external parties:* The external parties inform themselves with the truth agent actual beliefs depending on their bias.
4. *Transmission of the actual beliefs:* The policymakers and entrepreneurs have their actual beliefs informed from the external parties depending on their bias.

10.4.2 The agents

One additional active agent role is considered within the imperfect information add-on: the external party. External parties the same attributes as policy entrepreneurs but have additional actions that they can perform.

The introduction of imperfect information also requires the addition of the political *affiliation* parameter to all active agents. This is introduced to help define from whom policymakers and entrepreneurs are more likely to gather their actual beliefs. This attribute is set by the modeller for each agent and does not change throughout the simulation.

Finally, the external parties see one additional parameter attached to each of the issues that is present within their belief system. This parameter, named as *bias*, is used to determine the bias of the agents towards specific issues. These are provided as input by the modeller. This parameter is provided on a scale from 0 to 5. If the external party has a bias of 0, the agent has a bias for that issue. The agent will only trust its own actual belief. This effectively means that the agent will not be updating its actual beliefs from the environment altogether. At the other end of the scale, a bias of 5 will mean that the agent informs itself directly from the environment. The bias values in between means the agents follow more or less the information that is provided

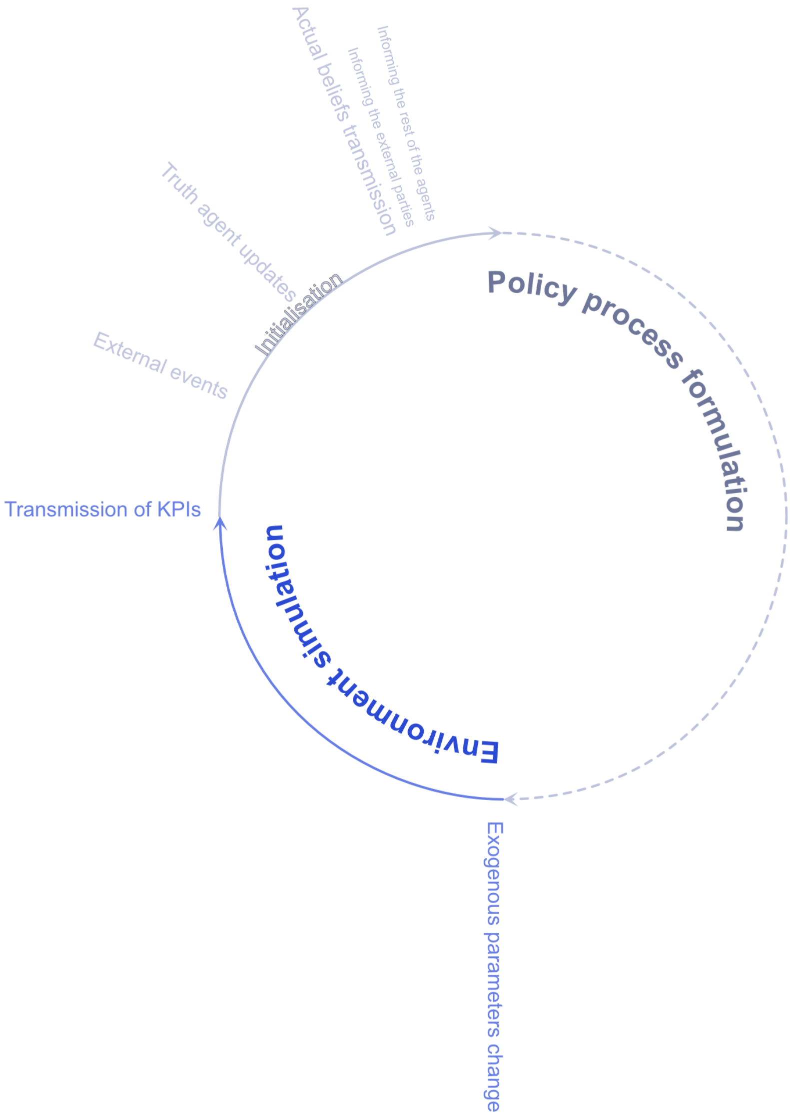

Figure 10.3 – Steps followed by the policy process for the ACF model for the initialisation with the imperfect information add-on.

by the truth agent from the environment. The equations used to operationalise this parameter are provided below.

10.4.3 The agent interactions

Three additional actions are considered with the introduction of a partial information system. Two of these interactions are related to informing of the external parties' actual beliefs and then those of the policymakers and entrepreneurs. The third is related to the interactions that the external parties can perform on other agents during the agenda-setting and the policy formulation when they behave similarly to policy entrepreneurs.

Informing the actual beliefs of the external parties

The informing of the external parties from the truth agent is based entirely on the issue-specific biased parameter. The equation that is used is given as:

$$B_{j,p} := B_{j,p} + (B_{j,\mathrm{TA}} - B_{j,p}) \cdot (0.2 \cdot W_{\mathrm{bias}}) \tag{10.5}$$

where B stands for the actual belief, j is the issue selected, W_{bias} is the bias value for issue j, p is the external party who is updating his/her knowledge and TA is the truth agent.

Informing the actual beliefs of the policymakers and entrepreneurs

Informing the policymakers and entrepreneurs from the actual beliefs of the external parties is linked to the agents' political affiliation. Agents sharing similar political affiliations will trust one another more and update their beliefs according to what the external parties give them. In the case where the affiliation is different, then the agents will be more sceptical of the information and take it with a grain of salt. In modelling term, this means they will partially change their actual beliefs based on what they observe in the concerned external parties. The speed of this update is a modeller input parameter defined per affiliation pairs.

The equation used to inform the actual beliefs of the policymakers and entrepreneurs is given as:

$$B_{j,p} := \frac{1}{n_{EP}} \sum_{k=0}^{n_{EP}} \left(B_{j,p} + (B_{j,k} - B_{j,p}) \cdot W_{\text{aff}}\right) \tag{10.6}$$

where W_{aff} is the affiliation parameter coefficient, p is the policymaker or entrepreneur who is updating his/her actual belief, k is the external party, and n is the number of external parties.

Blanket interactions

External parties can also influence other agents. It is assumed that their influence is performed on the entire policy subsystem. This is referred to as blanket interactions. These interactions are similar to the ones already known with the main difference being that all agents are impacted at once but to a smaller extent. This influence is more limited, as it is more diluted and less personal than other agents' influencing actions.

Similarly to the policy entrepreneurs, the actions are graded following the conflict of interests. The external parties then select the action with the highest likely impact to be performed. The equation to calculate the grades is the same as before but summed overall policy agents. It is given as follows:

$$Gr_{j,p} = \sum_{u} W_{CL_{j,p-n}} \cdot |\text{diff}| \cdot \text{Bo} \tag{10.7}$$

where Gr is the grade, j is the belief selected - preferred state or causal belief, p is the agent onto which the action would be performed, n is the agent performing the interaction, W_{CL} is the conflict level weight assigned based on the difference in beliefs, diff is the difference in beliefs, u consists of all policy agents, and Bo is the policymaker bonus that, in the policy formulation step, is set to 1.1 when p is a policymaker and to 1 when that is not the case.

The equations used to perform these blanket interactions are then given as follows for the causal relations:

$$C_{j,p} := C_{j,p} + \frac{1}{n_a}(C_{j,n} - C_{j,p}) \cdot R \tag{10.8}$$

where C stands for causal belief, R for resources, p for the agent being interacted upon, n for the agent performing the interaction, and n_a is the number of agents in the subsystem.

And similarly for a preferred state interaction:

$$G_{j,p} := G_{j,p} + \frac{1}{n_a}(G_{j,n} - G_{j,p}) \cdot R \tag{10.9}$$

where G stands for the preferred state.

And for the actual beliefs:

$$B_{j,p} := B_{j,p} + \frac{1}{n_a}(B_{j,n} - B_{j,p}) \cdot R \tag{10.10}$$

where B stands for the actual beliefs.

Actual belief agent-on-agent interactions

The last interaction that is added is one for the policymakers and entrepreneurs. Because actual beliefs are now subjective, agents can influence one another on these as well. Similarly to what was presented before, agents have a conflict level for these actual beliefs based on their differences in beliefs. This is used to select what interactions should be chosen between all three interactions (causal beliefs, actual beliefs and preferred states). If the preferred action selected is on an actual belief, the equation that is used to perform the influence is then given as follows:

$$B_{j,p} := B_{j,p} + (B_{j,n} - B_{j,p}) \cdot R \tag{10.11}$$

where B stands for actual belief, R for resources, j is the actual belief selected, p for the agent being interacted upon, and n for the agent performing the interaction.

10.5 Discussion

Following this formalisation, it is already possible to outline some elements that might require specific attention and issues that could arise during the simulation of the model based on the assumptions that have been made. This concerns each of the four levels of complexity that were added in this chapter. I discuss here what these elements are, why they might be problematic and potential alternatives.

10.5.1 The +PL model

Several operationalisation assumptions were made to introduce policy learning into the model. In several instances, different approaches could have been considered for these assumptions. This includes, for example, the way the interactions are performed. The current algorithm selects agents randomly to perform their actions until the agent's resources are exhausted. Once that is the case, a new agent is randomly selected, and this keeps going until all agents have spent all their resources. A different approach might have seen a process where, after each interaction, a new agent is selected randomly to perform one action, with this continuing until all agents have spent all their resources. Effectively this is a choice that was made based on computational efficiency considerations and not following the literature. It is unclear how this might impact the results of the model. Future work might try to understand better how a different algorithm might impact policy learning, if at all. This is an operationalisation problem for which no answer should be expected in the qualitative literature. On a related topic, agents do not update their preferences after every influencing action, again for computational efficiency reasons. In some rare cases, this might mean that agents are influencing others on an issue that they no longer consider to be most salient.

Another underlying assumption mentioned above argues that the agents spend all their resources each round. This assumption is once again one amongst a multitude of other possible assumptions. The literature does not clarify this point and does not deal with resources in the way they are considered here. There is also the possibility to implement another approach that would see agents preserve their resources for future rounds when a certain threshold has been reached - for example they are satisfied in the other agents' beliefs. This model, along with the common language approach, opens the door for different strategies that agents might follow, such as only spending resources in crucial moments, or withholding resources if they estimate that their interests are being fulfilled. A wide range of strategies can be considered. Such strategies would most likely be tailored to the context such that model exploration can be performed or because they have been observed in a specific case being studied. However, to implement such a strategy, the literature should first need to be expanded such that strategies observed and described in real cases be implemented in the model.

Resources have a more significant impact than outlined within this policy learning extension formalisation. The way the resources are tuned has a direct impact on the rate of agents' policy learning. This is important as this rate is described in the literature (Sabatier, 1987). It mentions that policy core issues only change over several years, while secondary issues can change more rapidly. As a result, when simulating a specific case, the way the resources are implemented and tuned should take into account in which stage of the process these resources are spent in. This could be done by changing the practical impact of the resources per process step. For cases where the evolution of the agents' beliefs is known, this could be used to calibrate the model, and in certain instances, validate it. Note that this does not relate to the relative difference in resources between agents. That difference should not be affected as it would affect the power dynamics at play in the policy subsystem.

In the present implementation, for exploration of this dissertation, it is assumed that policy learning should not exceed a change in belief of more than 0.1 for any agent from step to step. This is an arbitrary number that can be tweaked depending on the political and geographical context of cases studied. This value is affected both by the practical impact of resources but also by the number of agents in the model. If there are more agents, and they all focus their influence on one agent, that agent might see its belief change incredibly fast, something that would be unrealistic. Therefore the model should be calibrated such that such a scenario is made impossible.

10.5.2 The +Co model

Several operationalisation assumptions were made when introducing the coalitions into the model. These assumptions are necessary as details on how coalitions collaborate and coordinate at the operational level are scarce throughout the literature. The first assumption relates to the way the coalitions are created. In the proposed approach, it is assumed that coalitions are formed each time step anew, though the criteria used to form these coalitions make it such that the same coalitions are created over and over as the simulation progresses. This approach was taken for practicality reasons. Questions remain, however, on what a different approach would look like and how this might affect the composition of the resulting coalitions. In the literature, coalitions are very stable. They often see their memberships remain the same over the decades. If a different approach was used and coalitions were not remade each time step, what would then be the criteria to remove a member from a coalition? This is not clear in the literature but is an aspect that can be easily tested using this modelling approach.

Similarly, it was assumed here that agents coalesce around a single, modeller-defined, salient policy core issue and that this policy core issue does not change throughout the entire simulation. From the literature, it is not clear whether this assumption is valid or whether a different criterion should be used to form coalitions. There is the possibility, in specific policy subsystems, that agent be multi-dimensional and consider more than one policy core issue to come together into a coalition, complicating further the definition of thresholds for a simulation. There is also the possibility that this coalescing issue change over time, depending on several factors including, but not limited to, the polarisation of the political arena or dynamics that are external to the policy subsystem. In Pierce (2011), while studying coalition membership and its relation with the evolution of the policy core beliefs of the actors, Pierce found that "the belief systems of advocacy coalitions may converge to share multiple beliefs with a single advocacy coalition while maintaining their stability." Such studies can provide avenues on how to deal with coalition membership, though the teachings might remain case-specific, only pushing the problems of modelling further down the road.

Overall, the creation of this model raises more questions than it answers - which is coincidentally one of the uses of models. One aspect that is known is that there is no single answer to these questions. Each case, depending on geography and the political system, might have different mechanisms responsible for the creation of coalitions. Some cases will see the creation of more than three coalitions while others just one. In such cases, then several different approaches would have to be tested and explored to attempt to find the one that best emulates what is observed in reality.

These questions also extend to the resources that members are willing to contribute to their coalition. Can these contributions change over time, depending on the strategy developed by the coalition or by its members? Questions are also raised on the strategies that coalitions might adopt to influence other agents. In the present approach, the coalitions are considered rudderless, with no leader, and merely an average of their members. They pursue the actions perceived as most effective regardless of strategy. In reality, a leader is often observed in coalitions, and that actor will be the one organising a strategy for the coalition defining its future actions (Smith, 2000). However, how can this leader be systematically chosen within the confine of a model? Should it be based on the agents' beliefs, their resources, a combination of both, or maybe another parameter that has not been considered so far? Beyond this, it is also possible to ask following what beliefs the coalitions will perform its actions? These are only a few questions that need to be answered to get a better understanding of the functioning of coalitions. Once again, the answers probably vary depending on the context, and this model is an excellent opportunity to test and explore a variety of options.

Finally, and although this model raises many questions, it could also help explore a multitude of assumptions to advance the qualitative literature. Because it is simple to change these assumptions, test them and analyse the results, the use of a model might be the best tool to try to understand how systems such as coalition might function or at least build a model that approximates well the inner workings of coalitions.

10.5.3 The +PK add-on

Two assumptions were made with the introduction of partial knowledge that, if changed, could affect the results of the simulations. They are the partial knowledge initialisation and catch-up assumptions. To begin the simulation, the agents must be provided with initial values for their knowledge of other agents' beliefs, i.e., their partial knowledge. This consists of a large amount of data that is likely to be difficult to obtain. To collect this data, each actor would need to be asked about the beliefs of all other actors within the policy subsystem. Even if that were possible, discounting the time it would take, not all actors will have an opinion on every issue of every other actor, and therefore the data might not be available. This is not currently taken into account in the model. Furthermore, regardless of availability, in such cases, data will have to be assumed or guessed to simulate the model.

In a case where the model is simply tested or used for theory exploration, as it will be done for the predation model in the following chapter, the partial knowledge of the agents needs to be initialised as well. Here, depending on the assumptions made, the results could be significantly affected as the partial knowledge informs the agents on the actions they can perform, and in turn, affects the policy learning. Currently, for the simulation presented in this dissertation, partial knowledge is generated using random values. In the future, there might be merit in studying the impact of different initialisations methods on policy change to understand better whether a lack of data for a real case might interact with the results. If this is the case, then it would also be essential to determine in which case it might be better not to use the partial knowledge add-on so as not to draw incorrect conclusions from the results.

In the formalisation presented above, it is also assumed that agents gain knowledge based on interactions with other agents. This is an assumption that is made out of necessity and it is alluded to in the literature (Leach and Sabatier, 2005). Once again, this assumption can affect the results of the simulations. This knowledge gain is regulated by a constant that limits the speed at which agents update their partial knowledge, something that in turn affects their actions, policy learning and policy change. Furthermore, this way of dealing with the gain of knowledge assumes that the agents are truthful to one another or reveal a part of the truth when they interact.

It assumes no foul play or deceptive interactions between agents that might convey false information. This is something that can be true for most cases, but that can also be incorrect for cases where there is a high polarisation of the political arena. In such a case, a different approach would need to be considered.

10.5.4 The +PI add-on

The introduction of the imperfect information add-on is the part of the increase of complexity that counts the most assumptions and therefore, the most uncertainty. This directly relates to the fact that the literature hardly addresses this part of the policy process and means that the vast majority of these assumptions are not validated. They are introduced only to show the possibilities provided by the use of a model. The goal with the introduction of partial information is to ultimately deal with aspects such as political bias and fake news propagation, and their effects on the policy process.

The current approach provides ways to deal with bias and misinformation, but it remains a static approach. To illustrate this, let me take an example of simulation of fake news. One way to do so would be to initialise the model with an external party having fake news-like actual beliefs and a bias of 0 for that specific issue. This would prevent that external party from informing itself from the environment, therefore maintaining its beliefs, and then informing other agents about its incorrect actual belief. This approach can work but remains a static one. It does not address the potential for agents to change their strategies when it comes to informing other agents with inaccurate actual beliefs. Therefore, this approach is limited and should be further refined for future work focused on this topic.

The assumptions related to the communications between agents currently use the concept of political affiliation to transfer information. This affects whether the agents trust one another or not. This approach uses an embryonic political affiliation network. It only allows the trust element to vary from affiliation to affiliation. In reality, this is much more complex. Modelling wise, there is no need to mirror reality but introducing more complexity into this network and how it influences the transfer of information might allow for the better understanding of information propagation through a policy subsystem, may it be real or false information. For example, despite affiliations, agents might be more likely to listens to educational institutions or the judicial branch compared to media networks or other politicians.

Within the +PI add-on, the assumption that external parties can influence other agents through blanket influence is also made. This is a concept that is not present in the literature. Other assumptions could have been made such as one where external parties are either not able to influence other agents at all, besides their role of information purveyors, or one where external parties are only allowed to influence agents the same way policy entrepreneurs are. This is, once again, something that will be dependent on the case and its political and geographical context.

10.6 Conclusion

To conclude, the formalisation, along with its associated conceptualisation in the previous chapter, presented in this chapter provides one approach to the simulation of the advocacy coalition framework. This approach is filled with assumptions that remain untested or unexplored in the literature and that are therefore not validated. However, this is one approach amongst many other possible approaches. The goal of this dissertation is to show that the common language can be used to develop models of the policy process, and this implementation is an example of that.

In the discussion, I argued and showed that it is possible to use the common language to build alternative approaches to the one presented here. This might be needed to test specific context, address specific geographical specificities, or test new extensions of the current theories. The following chapter will present the simulation of this ACF implementation model with both the predation and the Swiss electricity market model. The goal with these demonstration is not predictive. Instead, it is to focus on what can be learnt from the simulation of these models with respect to the theories or the behaviours that can be exhibited from the simulations.

11 The predation model

The first simulation is once again with the predation model. This time it is to explore and test the ACF implementation model. The predation model is simple and convenient, making it the perfect candidate for theory exploration. This chapter follows the structure used in the previous chapter to outline the implementation of the ACF model. This is done in four steps, one step for each complexity level added. By testing the model in this way, it is easier to understand and explain the results that are obtained. Considering all complexity levels at the same time would make it difficult to extricate the different overlapping dynamic and interacting behaviours.

There are two goals to the testing of the ACF implementation model. The first one is to test the assumptions that were made throughout both the conceptualisation and then, in more detail, the implementation. However, the goal is not to experiment with alternative assumptions, as there would be no real case to compare the results to. This could be the subject of future research. This chapter is used to demonstrate the proposed approach.

The second goal is to understand better whether the added complexity is useful and insightful. Adding complexity makes the model more difficult to initialise and to analyse. Through these tests, I want to clarify whether it is indeed useful to increase the complexity or whether a simple model would generate similar behaviours and results that are sufficient to gain the insights needed into either the policy process or the socio-technical system.

Because the predation model was presented earlier in Chapter 6, it is not presented again here. Furthermore, the hybridisation of this model, that is the creation of the belief system, and the policy instruments are also the same for the ACF implementation model. They are also not repeated in this chapter. For each of the complexity increase, several steps are followed. First, one or more research questions are outlined to define the goals of each simulation. Scenarios are then constructed based on these research questions. These include both the initialisation parameters for the policy process model and the predation model. Finally, for each complexity level, the results are shown and explained on a per scenario basis. Not all figures are presented in this chapter as they would take too much space, and they do not necessarily all provide new insights for the reader. They are however all included in Appendix E. This chapter is then concluded with a discussion addressing insights from each of the four complexity additions.

11.1 Emulating policy learning (+PL)

The first complexity increase is the introduction of policy learning to the policy process model. This is achieved with the introduction of the role of the policy entrepreneur and with agent influence actions. The details of how the experiments were crafted are shown within this section, along with the results' analysis.

11.1.1 The research question

The main research question formulated to study the +PL model is the following: *What impact does the introduction of policy learning in the policy process model have on policy change?*

The goal is to understand the repercussions of policy learning onto the agents, their beliefs, their decision-making process, their policy instrument selections and, ultimately, the policies they implement. This is done within the context of the model designed here. This should also include an understanding of the impact of the resources on policy change and how the balance of power might affect the speed of policy learning. Throughout the ACF literature, policy learning is mentioned as one of the four ways leading to policy change (Sabatier, 1988). In this section, I look at whether this is also the case for the +PL model.

11.1.2 The scenarios

Four scenarios are considered to answer the question asked above. Compared to the simplest implementation model, several new parameters are present in this model. This includes the agent resources and the new agent distribution - considering the introduction of the policy entrepreneur role. The agent distribution is perhaps the most crucial aspect of the model to be initialised. This is because this distribution will initially impact the selection of the agenda and the policy instrument at the end of the process directly. All agents set the agenda while the policymakers can only select the policy instruments. The initial balance of power introduced through the agent distribution can influence the selection of both agenda and policy early on the model, something that can have repercussions throughout the rest of the simulation. This has the potential to be an issue only early on in the simulation as, later on, influence between the agents could erase this power dynamic.

To simplify the initialisation of the simulations, the concept of political affiliation is introduced. This concept is used as a way to systematically initialise the beliefs of a large number of agents along with two different groups: affiliation 0 and affiliation 1. This allows not to have to specify the beliefs of each of the agents one by one. This grouping is also used to analyse the results, allowing for easier comparisons of trends and behaviours than if lone agents were considered. These affiliations have no impact on the mechanisms of the model. They are not considered in any part of the agents' decision making. This only happens for the imperfect information add-on, the fourth level of complexity.

The agents' preferred states are provided in Table 11.1, depending on their affiliations. Their causal beliefs are provided in Table 11.2. The affiliation 0 agents are pro-sheep. Affiliation 1 agents, on the other hand, would like to see a 1:1 sheep-wolf ratio. Agents in both affiliations have the same causal beliefs, which are set as realistic causal beliefs. For all scenarios, a set distribution of agent roles is set. This includes two policymakers and four policy entrepreneurs for affiliation 0, and one policymaker and four policy entrepreneurs for affiliation 1. This selection guarantees a certain initial power distribution where affiliation 0 agents will select both the agenda and the policy instrument. The rest of the details on each of the scenarios are provided as follows:

- Scenario 0 (SIM) - This is a scenario where no influence action is taken into account. It is a repeat of the simplest implementation model (SIM) simulation but with the belief and agent distribution used in this section. This scenario is run as a comparison point for the other scenarios to assess what impact the introduction of complexity has.
- Scenario 1 (Benchmark +PL) - The benchmark scenario is used to provide a baseline for the other scenarios to compare. For this scenario, affiliation 0 agents are provided with more resources (75%) than affiliation 1 agents (25%). Comparatively, this makes affiliation 0 agents three times as impactful as affiliation 1 agents when it comes to influencing agents.
- Scenario 2 - For this scenario, a change in the resources distribution is made

to test how the resources affect policy learning and policy change. Affiliation 0 agents are provided with only 10% of resources while affiliation 1 agents are provided with 100% resources, making them ten times as impactful when it comes to influencing agents.

- Scenario 3 - For this scenario, an external event is introduced. This event affects the preferred state of the policy core issues PC1 and PC3 for agents from affiliation 0 at time step 4 of the policy process. The change is detailed in Table 11.1.

- Scenario 4 - This scenario similarly introduces an external event. This time it is on the preferred state of secondary issue S1 for agents from affiliation 0. It is also triggered at time step 4 of the policy process. The details are provided in Table 11.1.

	Issues					
	PC1 Sheep	PC2 Wolves	PC3 Grass	S1 Sheep growth	S2 Wolves growth	S3 Grass growth
Agents						
All scenarios						
Affiliation 0	400	50	2000	75	-50	200
Affiliation 1	200	175	1700	50	25	120
Scenario 3 - policy process time step 4						
Affiliation 0	50	-	250	-	-	-
Scenario 4 - policy process time step 4						
Affiliation 0	-	-	-	-80	-	-

Table 11.1 – Preferred states for the policy agents on the interval [0,1] for all scenarios.

All scenarios			
	PC1	PC2	PC3
-S1	1.00	0.75	-0.75
-S2	-0.75	1.00	0.25
-S3	0.50	0.75	1.00

Table 11.2 – Causal beliefs for the policy agents on an interval [-1, 1] for all scenarios.

Note that a Monte Carlo simulation could be run for the +PL model, but it would be computationally burdensome to set up and would complicate the analysis of the results. This is a result of the many parameters that would have to be taken into consideration. Beyond these technical issues, and this was true for the simplest implementation model, small variations of the parameters often do not lead to observable behaviour changes in the present model. Its complexity, the inherent inertia present resulting from the interaction of two models, and the limited stepwise policy instruments would most likely make it challenging to find meaningful behaviours in the results if a Monte Carlo simulation was run. Instead, the approach followed is one where a set of key parameters are changed to the extreme to attempt to generate new and unexpected behaviours that can help better illustrate how the model works.

11.1.3 The results

The results are presented for all scenarios to answer the research questions aforementioned. Four types of figures are used to display the results. First, there are the preferred state's figures. These figures show how the agents' preferred states, maybe it be for policy core or secondary issues, change over time. They are split colour-wise depending on the agents' affiliation and are plotted using percentile time series. These figures are used to illustrate policy learning.

The second type of figures plots the issues preferred, and therefore selected for influence actions, by the agents. They are presented as bar plots summing up the issues selected by each agent throughout the simulations. The policy core issues are plotted independently of the secondary issues. Once again, affiliations are differentiated to understand better the difference in selection between agents in each affiliation.

The rest of the results are plotted using similar methods as in the previous chapters. This includes bar plots for the selection of the policy instruments and percentile time series for the predation population results.

Simplest implementation model vs. ACF implementation model

First, it is important to show if policy change is affected by the introduction of policy entrepreneurs and influence actions, and, in turn, if it affects the predation model. The policy instruments implemented for both models are shown in Figure 11.1, while the sheep and wolf population counts are given in Figure 11.2. With similar initial beliefs and agent distributions for both scenarios, one would expect that the impact of the introduction of additional complexity is limited if not inexistent. Even though agents can now influence each other, these influences should only reinforce the trends present in the model where these influences were not possible because of the belief similarities and the power dynamics resulting from the agent distributions.

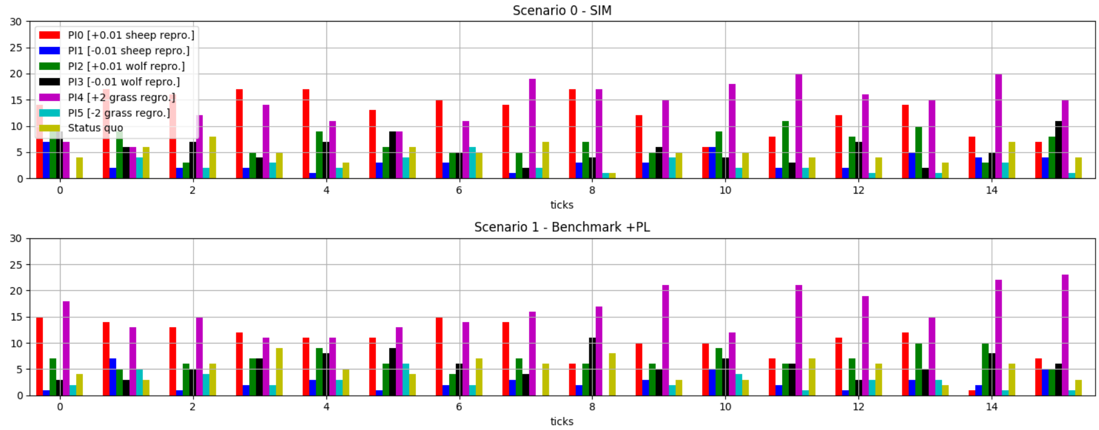

Figure 11.1 – Policy instruments selected for the SIM and benchmark scenarios.

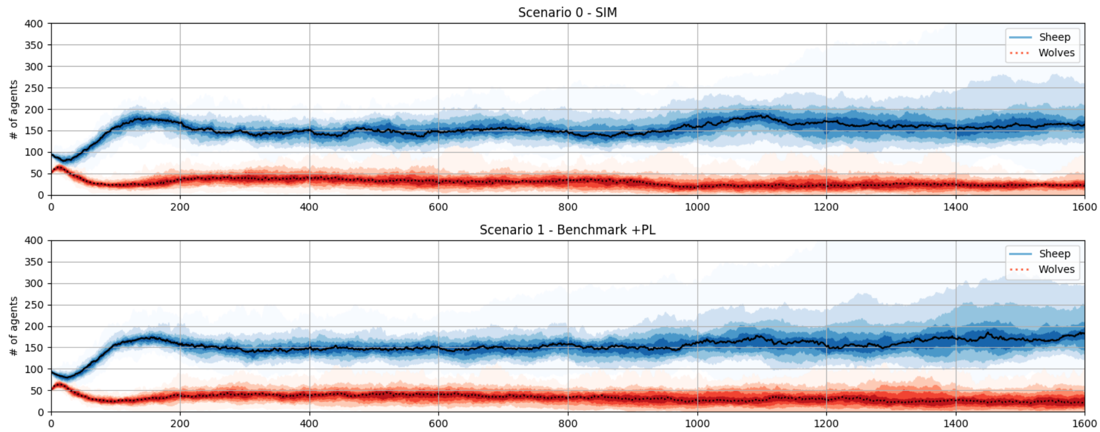

Figure 11.2 – Predation model results for wolf and sheep counts for the SIM and benchmark scenarios.

These expectations are met upon looking at the results. Figure 11.1 shows that the differences in policy selection between the simplest implementation scenario and the +PL scenario are limited. Having similar beliefs, the agents choose similar policy instruments. The introduction of a minority coalition (a minority in resources and agent numbers) has no tangible impact on the results. The agents tend to prefer policy instrument PI4, followed by policy instrument PI0 in both scenarios. These observations, in turn, help explain why the predation model results are also similar. In both scenarios, the wolf population is low, hovering below fifty, while the sheep population is mostly concentrated around 175.

Figure 11.2 shows that in some simulations, the wolf population goes to zero before the end of the simulation while the sheep population remains stable. This trend seems more prevalent for the +PL scenario compared to the simplest implementation

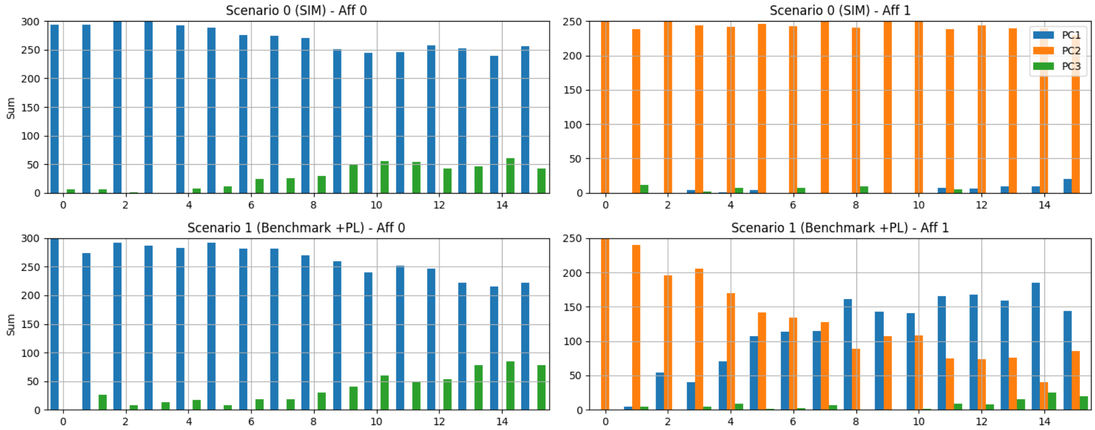

Figure 11.3 – Policy core issues selected per affiliation for the SIM and benchmark scenarios.

scenario. This suggests that, due to policy learning, inaccurate beliefs are being reinforced at a greater speed, leading the agents to select policy instruments that they perceive will fix their issues, but that ultimately do not. This policy learning is prevalent as is displayed in Figure 11.3. The agents from affiliation 1 are being influenced out of selecting policy core issue PC2 in favour of issue PC1 by affiliation 0 agents, something that is effective from time step 6 onwards. The reverse influence is not visible in this graph because of the difference in resources between the agents in each affiliation: the influence of affiliation 1 agents onto affiliation 0 agents is not strong enough and does not result in a change of their preferred policy core issues. Overall, the policy learning observed has a limited impact on policy change because the preferences of the dominant group of agents are only reinforced.

Though policy learning is present and significant for the policy core issues, it is practically non-existent for the secondary issues. Throughout the simulations, agents in both affiliations interchangeably prefer secondary issues S1 and S3. This means that they only influence one another on these two secondary issues where the gap in belief is around 0.1 initially. This preference is a result of the significant difference between their actual beliefs and their preferred states. However, as shown in Figure 11.4, agents have their most significant gap in belief, and therefore the largest potential for policy learning, in their secondary issue S2 preferred state. The gap in belief between agents of both affiliations is a little over 0.3, more than three times that of the other secondary issues. However, as a result of their focus on issues S1 and S3, every agent fails to seize the opportunity of influencing others on this issue, in turn limiting the policy learning at the secondary issue level.

This limited policy learning behaviour helps explain why the policy instruments barely

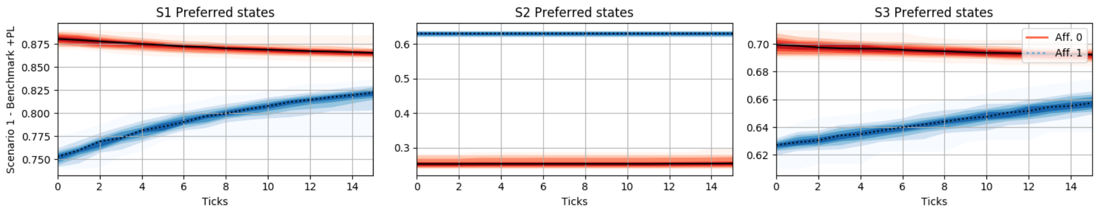

Figure 11.4 – Secondary issue preferred states for the benchmark scenario.

change throughout the simulation and between the two scenarios. This is because the secondary issue selection barely changes, which is, in turn, a result of the agents' preferred states barely changing. The secondary issues are most important when it comes to the selection of the policy instruments. If they remain stable, then the policy instrument selection will remain the same between the two different models regardless of the other parameters considered, and despite other changes present in the +PL scenario.

The impact of resources

The impact of the agents' resources is investigated by changing the distribution of resources between affiliation 0 and 1 agents. This is done in scenario 2. The results are then compared to the benchmark +PL scenario. One can expect that a change in the resources as drastic as the one performed in scenario 2 will see the policy learning reversed for the policy core issues at least. This time, it will be the affiliation 1 agents that will induce the majority of the policy learning and not affiliation 0 agents. As was outlined in the previous section, policy learning at the secondary level is minimal, and a change of the agents' resources is not expected to affect this trend.

From Figure 11.5, it is possible to see that this is indeed what is happening. The affiliation 1 agents are influencing the affiliation 0 agents out of their issue PC2 preferred states (the wolf count). This is an acceleration of the trend observed for issue PC2 in the benchmark +PL scenario. When it comes to the affiliation 0 agents' influence on affiliation 1 agents, it is still present. It is much slower for issue PC1, and the agents do not have enough resources to perform this influence for issue PC3, focusing their efforts on issue PC1 throughout the simulation instead.

However, the outcomes of this scenario are not as straight forward as one might expect. The influence of affiliation 1 agents on the policy core issue PC2 of affiliation 0 agents has no impact on the outcome of the policy core issue preference selection for the affiliation 0 agents. This is visible in Figure 11.6. Despite the evolution of their PC2 issue preferred states, they keep selecting issue PC1 as their issue of focus.

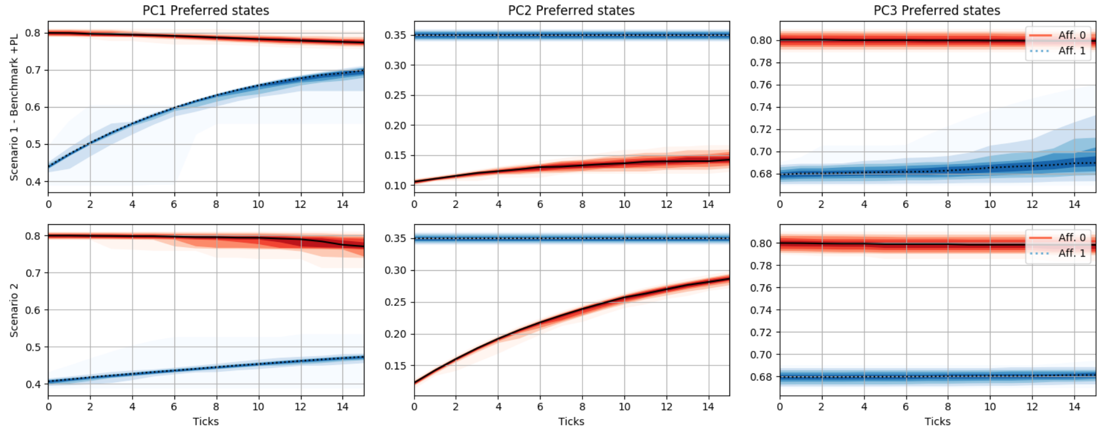

Figure 11.5 – Policy core issue preferred states for the benchmark +PL scenario and scenario 2.

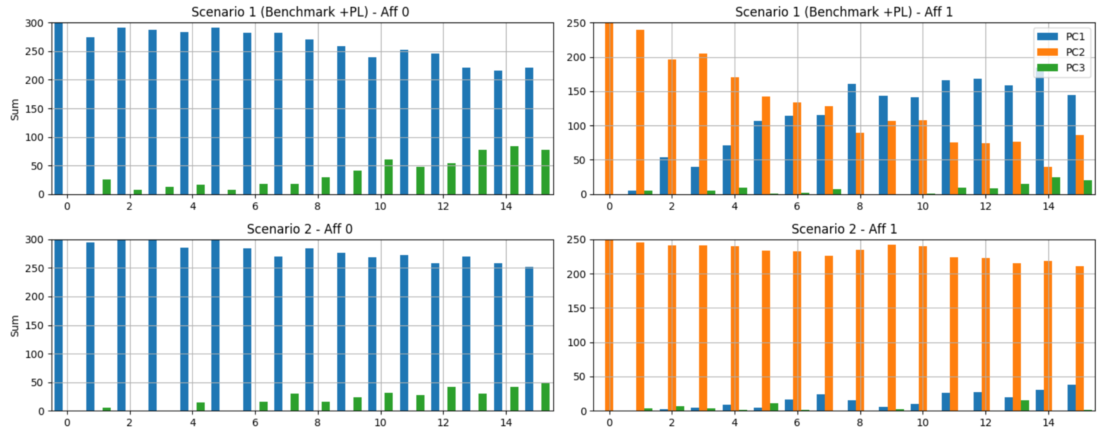

Figure 11.6 – Policy core issues selected per affiliation for the benchmark +PL scenario and scenario 2.

This outcome demonstrates that a change in the resources alone is not sufficient to affect the agents' preference selection effectively. This change in resources would have had a more substantial impact if the influence actions of affiliation 1 agents had been focused on issue PC1 instead of issue PC2. However, for this to happen, a change in the model that would affect the preference selection of affiliation 1 agents would have been needed. This could potentially be achieved by changing the initial preferred states of the agents.

Similarly to the previous results between the SIM and the benchmark +PL scenarios, the secondary issues are not affected by this change in the agents' resources. This is again due to the similar beliefs of the two affiliations' agents for the S1 and S3 issues preferred states. This results in preferences that remain unaffected. In turn, this leads to a selection of policy instruments that remains broadly the same between the two scenarios. As a result, the predation model results change only mildly. There are still a lot of the simulations that result in the wolf population going to zero while a large number of simulations remain stable until the last time step.

From these results, one can see that a change in resources can have, but does not necessarily lead to, an impact on the policies being selected. The right conditions also need to be present for this to happen. As is shown in this example, this change did not make a difference in the model's outcome because agents of both affiliations were focused on different issues where they had similar beliefs. They influenced one another on issues for which they already agreed on. To be more impactful, compounded to a change in the resources, a change in the initial preferred states would also have been needed. The next scenario looks into the impact of such a change in preferred states, although without considering a change in resources.

Changing the preferred states

The last two scenarios focus on a change in the preferred states linked directly to the sheep population (issues PC1 and S1) and the grass (issue PC3) for affiliation 0 agents, reducing them drastically. Scenario 3 focuses on the policy core issue changes while scenario 4 on the secondary issue changes. One would expect that a change in a preferred state will affect the model significantly at least when it comes to the policy learning behaviours. Based on the previous observations, it is likely that changes at the secondary issue level will have a much more significant impact on policy change than if changes are contained to the policy core issue level.

These expectations are confirmed by the results obtained. Looking first at the change in the policy instruments selected by the policymakers (see Figure 11.7), one can already note several differences between the benchmark +PL scenario and scenarios 3 and 4. For scenario 3, there is a significant increase in the selection of policy instrument PI4, once the external events introducing a change in the preferred states of issues PC1 and PC3 have occurred. This reinforces the selection of this same instrument seen in the benchmark +PL scenario. The agents also select policy instrument PI0 a lot less than in the benchmark +PL scenario.

As expected, the most striking changes occur with the results of scenario 4. Here a change of the S1 preferred state is introduced. The outcome is drastically different, with policy instrument PI4 being selected only a couple of times. Instead, initially, policy instrument PI0 dominates the selection of the policymakers. This changes around the mid-way point, at time step 7 of the policy process. From then on, a mix of

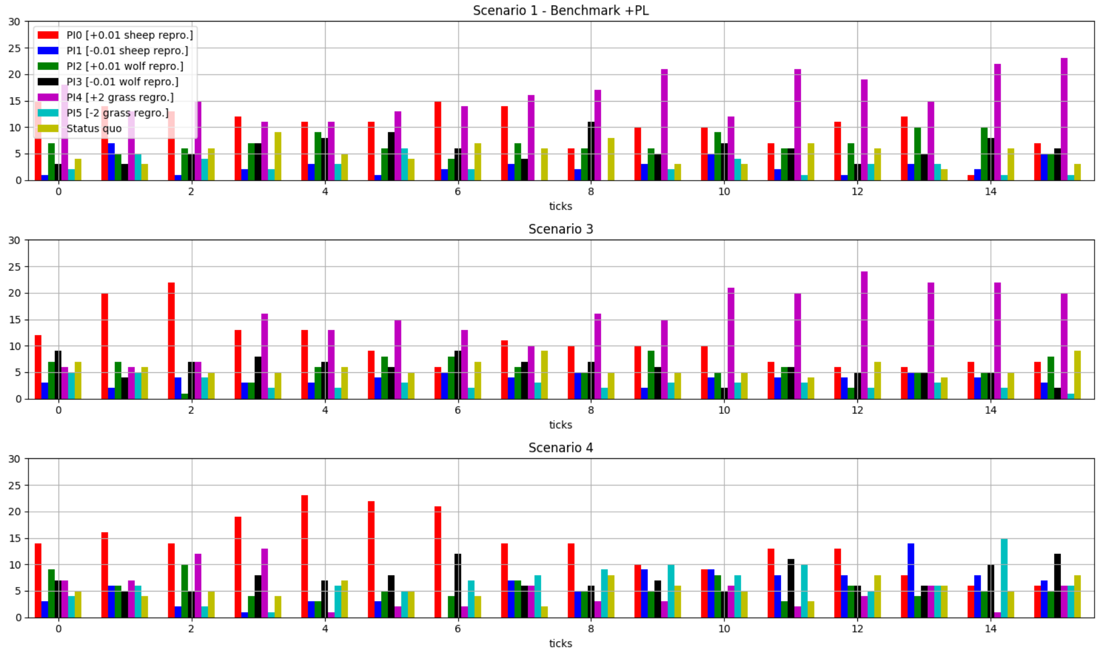

Figure 11.7 – Policy instruments selected for the benchmark scenario, and scenario 3 and 4.

all instruments is selected at a roughly equal level.

One of the outcomes brought on by this lack of one dominant policy instrument in scenario 4 is a collapse of the wolf population in the predation model (see Figure 11.8). It seems that this change of the S1 issue preferred states has led the agents into a trap. They are trying to strike an unrealistic balance because of their preferred states. Affiliation 0 agents would like to see eight times as many sheep as wolves in the predation model but at the same time, strive for a constant decrease of sheep population of 80 for every time step. A similar trend is present for the affiliation 1 agents. All the agents can do is to try to reconcile these preferred states with the instruments at their disposal. The results illustrate this reconciliation.

At first, they increase the sheep reproduction rate using policy instrument 0, but they soon find themselves too far from their sheep reproduction rate reduction objective. At the same time, the wolf population is plummeting. They attempt to correct this with any policy they can find. This leads to an almost random selection of policy instruments, most likely due to the assessments they have of the policies which guide their decision-making process. These assessments show that some policies lead to changes in the general direction of their preferred states, though no instrument is significantly better to get there. They all have roughly the same impact on the predation model.

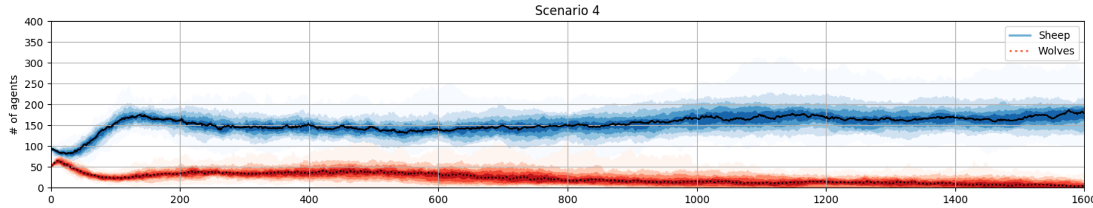

Figure 11.8 – Predation model results for wolf and sheep counts for scenario 4.

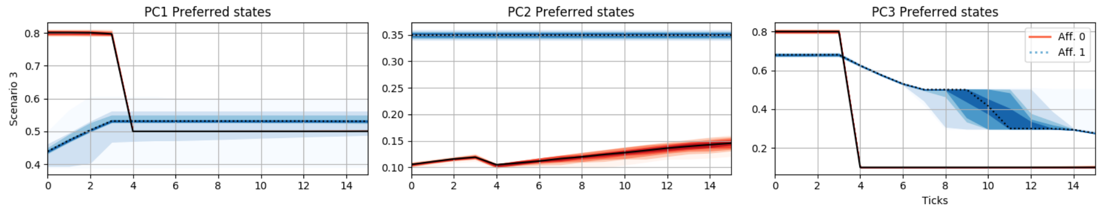

Figure 11.9 – Policy core issue preferred states for scenario 3.

Their efforts are not fruitless. Through the introduction of these policy instruments, the agents can find a stable situation with a model that essentially retains only two agent types: sheep and grass patch agents. This also means that the policy agents can mostly satisfy their preferred states for the issues related to the sheep and grass patches while having sacrificed their wolf objectives.

These policy changes can be traced back to the policy learning behaviour that is triggered by the introduction of the new preferred states. Figure 11.9 and Figure 11.10 illustrate the policy learning for both policy core and secondary issues. For scenario 3, this happens for issues PC1 and PC3 while for scenario 4, this happens for issue S1. In both cases, policy learning is a result of the influence of affiliation 0 agents on affiliation 1 agents. This influence ultimately results in a change of the secondary issue selection for affiliation 1 agents. They go from choosing issues S1 and S3 at the same rate to preferring issue S3 around the mid-point of the simulation. This secondary issue selection will ultimately influence the selection of the policy instruments for agents in both affiliations, leading to the outcomes described previously.

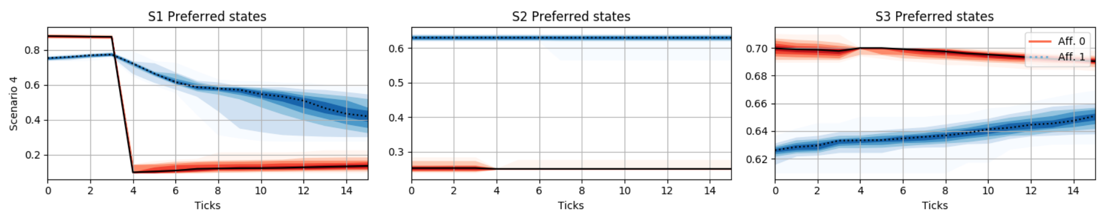

Figure 11.10 – Secondary issue preferred states for scenario 4.

The results from scenario 3 and 4 show the fragility and sensitivity of the overall hybrid model. It also shows that the agents' decision-making process is entirely rational. When they are provided with extreme beliefs, the agents will select policy instruments that lead to the collapse of the wolf population so that they can satisfy their other beliefs. Moreover, this will remain the best outcome they can achieve, considering their preferred states and the policy instruments at their disposal. The results also show that the agents will prioritise at any cost and without foreseeing potential adverse consequences, illustrating their rationality.

In this case, they prefer the sheep population, and they will not mind if the wolf population collapses. The agents have a target of fifty wolves, but they cannot take into account the fact that if they overshoot this target, and the population goes to zero, they will never again be able to bring back wolves into the model. This is what ultimately happens, with a majority of the simulations seeing the wolf population disappear entirely. Once the wolves have disappeared, the agents will move on, unable to meet that specific preferred state. Instead, they focus on meeting their other four preferred states relating to the sheep population and the number of grass patches, and the rates at which they grow.

11.2 Adding coalitions (+Co)

Adding coalitions is the second step in the addition of complexity to the ACF implementation model. It comes on top of the policy learning that was added in the previous section and for which the results have been presented and analysed. Coalitions are approached as super-agents who have an oversized role in the policy subsystem. Their influence is more effective due to their ability to pool resources from their members. Similarly to the previous section, this section outlines the experiments that are run to test this complexity level and presents the results obtained from them.

11.2.1 The research question

The goal with the addition of coalitions is similar to the one for the policy learning complexity extension: understanding how coalitions influence policy change. However, because of the high number of assumptions that were made to simulate the coalitions, the goal will be focused on understanding how these assumptions might affect policy change.

To this effect, the main research question to answer is the following: *What impact does allowing coalition formation in the policy process model have on policy change?*

Behind this research question lies another goal: finding out whether the addition of coalitions is useful. As was mentioned in a previous section, introducing coalitions into a model requires a large amount of data for its initialisation. Considering this, and depending on the results obtained in this section, the disadvantages brought by the data collection needs might outweigh the advantages brought by the addition of coalitions. The results should hopefully help better understand whether that is the case.

11.2.2 The scenarios

The coalitions introduce three new main parameters to the model. First, the coalition creation threshold parameter that defines the threshold of belief difference between agents for the creation of new coalitions. Second, the shared resources parameter that defines the number of resources that agents contribute to the coalition they join. Third, the salient (policy core) issue parameter that defines which policy core issue is the policy core issue around which coalitions assemble in a specific policy subsystem. These three parameters are at the centre of the scenarios used to test the present model complexity increase.

Note that all other parameters retain the same initial values as for the benchmark +PL scenario - scenario 1 - to allow for easier comparison. This includes the preferred states as provided in Table 11.1 and the causal beliefs in Table 11.2. The agents and resource distributions also remain the same for both affiliations.

- Scenario 5 (Benchmark +Co) - The first scenario considered for the +Co model is a benchmark. It will be used as a point of comparison against the benchmark +PL model and the other scenarios. The salient issue is the sheep population issue PC1, the coalition creation threshold is set at 0.15, and the shared resources are set to 50%.

- Scenario 6 - The second coalition scenario changes the salient issue to the wolf population issue PC2. The rest of the parameters remain the same. The goal with this scenario is to understand better what impact the salient issue parameter has on the rest of the policy subsystem and on policy change.

- Scenario 7 - The third coalition scenario changes the salient issue to the grass patch counter issue PC3. The rest of the parameters are also initialised similarly. Once again, the goal of this scenario is to understand the impact of the salient issue parameter.

- Scenario 8 - This eighth scenario is used to run a Monte Carlo simulation of the model to test the impact of the two remaining parameters: the coalition creation threshold and shared resources parameters. The first is varied between the values of 0.01 and 0.25 and the second varied between 5% and 75%. Latin Hypercube sampling is used to create a set of 50 different experiments. Each is then simulated ten times, leading to a total of 500 simulations. The goal here is to understand how each of these two parameters influences the behaviours in the simulation.

11.2.3 The results

Similarly to the previous section, four different types of figures are used to illustrate the results. They go from graphs outlining the agent populations in the predation model to graphs detailing the policy agents' preferred states and preferences throughout the policy process. The results are presented below in three parts. First, the impact that coalitions have on policy change is looked at. Then the impact that a changing salient issue has on coalition creation is analysed. Finally, I look at how the coalition creation threshold and shared resources parameters impact policy learning.

The impact of coalitions

When first looking at the policy instruments selected in Figure 11.11, it is hard to see a difference between the two benchmark scenarios - +PL versus +Co. The main difference is that in the benchmark +PL scenario, the policy agents seem to be choosing the fourth policy instrument a little more. Regardless, for both scenarios, it is the dominant policy instrument.

The difference between the two models becomes more apparent when looking at the policy core issue preferred states. One would expect policy learning to be reinforced by the coalitions as they facilitate a common strategy. This is indeed observed. In Figure 11.12, it is possible to discern a faster rate of influence for policy core issue PC1 where affiliation 0 agents influence affiliation 1 agents. It remains a subtle change that is hard to spot.

The accelerated policy learning has another impact: the grouping of the coalitions. As can be seen in the first plot of Figure 11.13, all of the fifty simulations start with two well-defined coalitions. Due to the influence of one coalition, the one made of the affiliation 0 agents - the one with most resources, the members of the other coalitions are drawn away and attracted to the dominant coalition over time. In all

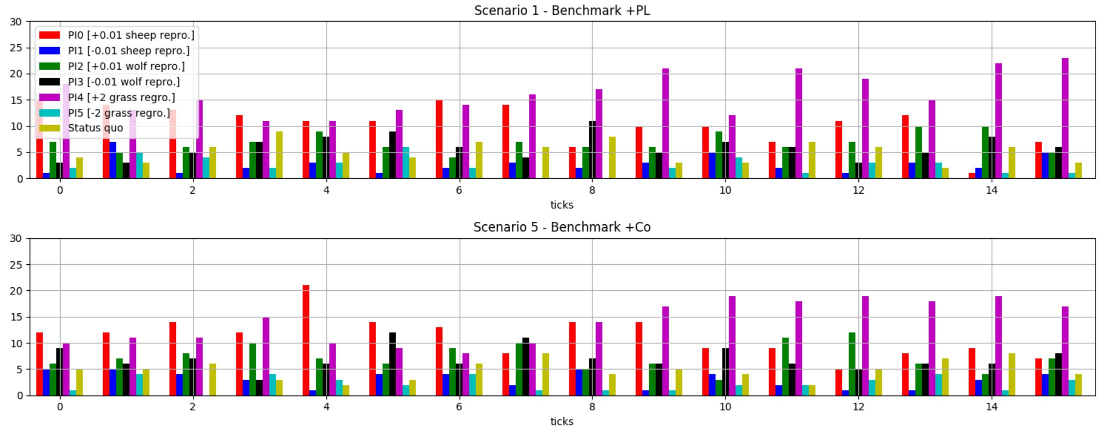

Figure 11.11 – Policy instruments selected for the benchmark scenarios.

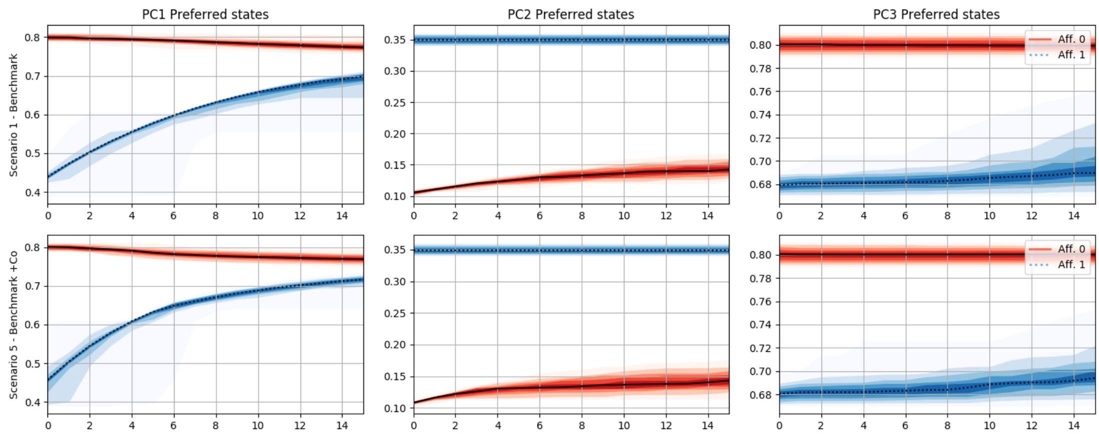

Figure 11.12 – Policy core issue preferred states for the benchmark scenarios.

simulations, the switch from two to one coalition happens between time step 6 and 10. The coalitions merge once the agents are within the value of the coalition creation threshold for the salient issue. So in Figure 11.12, when the preferred states of the agents are all within 0.15, only one coalition remains.

The salient issue

The matter of the salient issue around which the agents coalesce to create the coalitions is an important one. It is one of the parameters that can affect the number of coalitions that will be created. The number of coalitions per scenario is shown in Figure 11.13. As can be seen from the figure, in scenario 6 where the salient issue is issue PC2, two coalitions are always present. For scenario 7, where the salient issue is issue PC3, only one coalition is present throughout.

The preferred states of the agents can explain this difference between the scenarios. The preferred states of issue PC2 of agents in both affiliations are different and distant. Furthermore, they do not change over time as no agent considers that they are essential enough to select for influencing actions. In the case of issue PC3, the salient issue in scenario 7, the preferred states of the agents are close enough that they can all assemble into one big coalition at all time. In this case, policy learning would only reinforce the proximity in beliefs between the agents of both affiliations.

The coalition consolidation dynamic appearing in scenario 5 is not present for the other scenarios. Only in scenario 5 are the coalitions impacted by policy learning. This is because the agents' preferred sates are initially distant. As the agents prefer to influence one another on this policy core issue, policy learning ensues. Ultimately, this leads to a similarity in beliefs that do not justify the presence of two coalitions.

It remains to be defined how a modeller should decide which policy core issue should be chosen as salient. When making this choice, the modeller should consider that this selection is static and cannot change based on what is happening in the simulation. In most cases, s/he will have to decide based on interviews which policy core issue is most important to the creation of coalitions in that specific policy subsystem and base the selection of the salient issue on those answers.

These findings from the salient issue assumption bring up several questions on how coalitions are created and how they should be simulated. The current assumption is that their creation is based on a salient issue that is set by the modeller, and this salient issue does not change over time. This assumption could be changed in favour of a system where the salient issue can vary. For example, the salient issue could change depending on the coalitions' priorities or the policy subsystem as a whole. However,

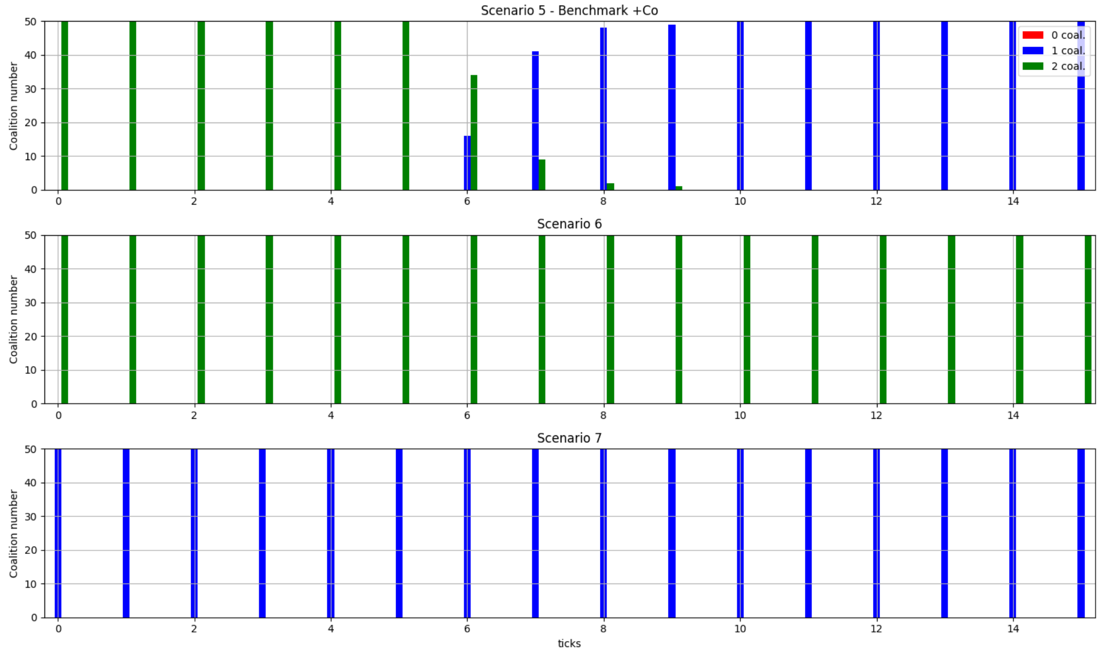

Figure 11.13 – Number of coalitions in the simulations for the +Co benchmark scenario and scenario 6 and 7.

this would then become a chicken and egg problem: for a coalition to know what issue is salient it first needs to be created and to be created it needs to know what issue is salient.

Finally, note that despite the change in the salient issue for scenario 6 and 7, the policy instruments selected by the agents barely changed. The coalitions kept advocating for what they would have regardless of the salient issue because the belief system drives their actions, and not the salient issue in both the agenda setting and the policy formulation steps. This means that affiliation 0 agents continued to focus on influencing affiliation 1 agents on issue PC1.

At the secondary issue level, this influence was concentrated on issue S1 and to a lesser extent on issue S3 for all three scenarios. The policy instrument selection was not impacted, and the agents continued to select the same instruments over time. In effect, this observation shows that the introduction of coalitions using the present assumptions has only a limited impact on the simulation outcomes. In the present example, it does not provide a lot of additional insights into the policy subsystem.

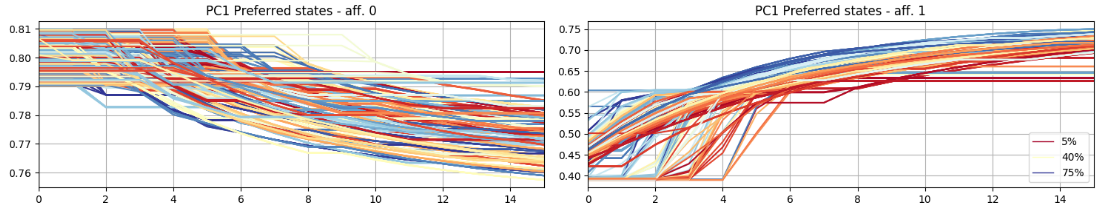

Figure 11.14 – Policy core issue 1 states for the scenario 8 with the colours representing the different levels of share resource parameter.

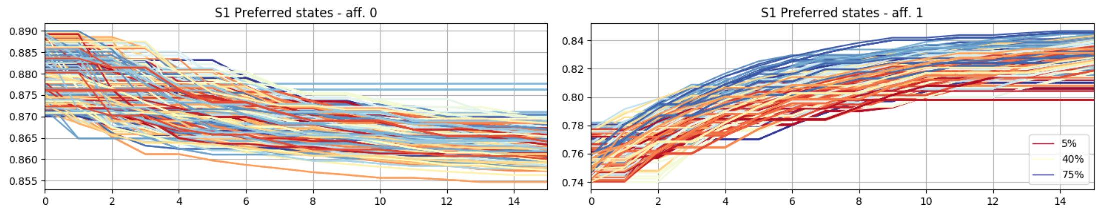

Figure 11.15 – Secondary issue 1 states for the scenario 8 with the colours representing the different levels of share resources parameter.

The Monte Carlo simulation

To study the results obtained from the Monte Carlo simulation in scenario 8, a different set of figures is chosen. The aim is to observe what impact the changes of the two parameters, namely the coalition creation threshold and shared resources parameters, have on the model. The expectation is that policy learning will be accelerated when agents provide more resources to their coalitions. The expectation for the coalition creation threshold parameter relates to the number of coalitions that will be created. It is unlikely to change for the same reasons discussed in the previous section. This is also discussed further on in the discussion section of this chapter.

The figures used to plot all 500 simulations at once are coloured depending on the values of the two parameters from blue (low value) to red (high value). Figure 11.14 and Figure 11.15 are two such examples specifically coloured for the shared resources parameter.

The plots confirm the expectations but only to a limited extent. The policy learning speed is influenced by the number of resources that are shared by the agents. The higher the amount of resource shared, the faster the policy learning in affiliation 1 agents. This is outlined for the PC1 and the S1 issues in Figure 11.14 and Figure 11.15. This is mostly visible for values above 65%. This acceleration trend of policy learning is, however, limited. It also does not happen for the influence of affiliation 1 agents onto affiliation 0 agents due to their more limited resources in the first place.

Finally, no link between the coalition creation threshold parameter and the policy learning can be observed in the results. More surprisingly, no link can be found between either of the parameters and the predation model population outcomes. Regardless of the parameter combinations, and even considering the accelerated policy learning for high values of the shared resources parameter, the results remain unaffected. The agents mostly continue selecting the same policy instruments, affecting the predation model similarly.

11.3 Including partial knowledge (+PK)

Considering partial knowledge within the policy process model is the third complexity addition. This addition can be seen as an add-on as it can be added to either the +PL or the +Co models. The partial knowledge's main direct impact is on the grading of the influence actions that the agents and coalitions can perform. In this section, I perform several experiments to see if the introduction of the partial knowledge impacts the selection of the issues, the selection of the policy instruments and ultimately the predation model.

11.3.1 The research question

The main research question, following in the footsteps of the previous complexity additions, is given as follows: *What is the impact of the introduction of partial knowledge on the decision making process of the agents and, subsequently, on policy change?*

This is investigated in several scenarios that are detailed below. The goal is to better understand the model that was created and how bounded rationality might affect the results obtained so far.

11.3.2 The scenarios

The partial knowledge only introduces one new parameter to the model: the catch-up parameter. It defines how much knowledge the agents will gain from their peers after an interaction. However, the most critical addition resulting from the partial knowledge introduction is the fact that it changes the procedures that are followed to grade the influencing actions. The grading of these actions is now dependent on the knowledge that the agents have on one another's beliefs. This knowledge also needs to be initialised. In the present simulations, this is done randomly to avoid including new assumptions that would need to be tested as well.

For the rest, the parameters are kept the same as they were for the +PL benchmark scenario - scenario 1. The agents' preferred states and causal beliefs are given in Table 11.1 and Table 11.2 respectively. The resource and agent distributions are also the same with affiliation 0 agents being given more resources and being more numerous than affiliation 1 agents.

The scenarios are built such that it is possible to compare the impact of the partial knowledge introduction with the results obtained for the +PL and +Co benchmark scenarios. For this, two scenarios are built: scenario 9, which is run as an add-on for the +PL model, and scenario 10 for the +Co model. The catch-up parameter value is the same for both scenarios. It is set at 20%.

11.3.3 The results

One can expect that the results will only be mildly affected by the introduction of partial knowledge. In the most simple terms, the incorrect knowledge that agents might have initially will lead them to perform the wrong actions because of incorrect high grades. Nevertheless, as they gain knowledge through interactions, they this will stop assigning incorrect grades to their influencing actions and will ultimately perform the right actions returning to the policy learning trends observed in previous scenarios. The main difference expected is a delay in the results.

This is indeed what is observed when the partial knowledge add-on is included. The policy learning trends are slowed down significantly. This is most visible for the PC1 issue preferred states, as shown in Figure 11.16. In some cases, and for some agents, policy learning seems not to be present because of the agents' inability to gain any knowledge before the end of the simulation due to a lack of interactions. However, over time, one can expect that policy learning would happen regardless, once interactions are present. The impact of partial knowledge is minimal on the rest of the model. The policy instrument selection is barely affected, and the predation model results also remain unaffected.

Testing different catch-up speeds for the partial knowledge add-on would not have significantly affected the results and the trends presented here. Increasing the knowledge gain speed would have lead to results even closer to the benchmark results for both scenarios. Slowing it down would have resulted in the opposite trend: slowing down policy learning.

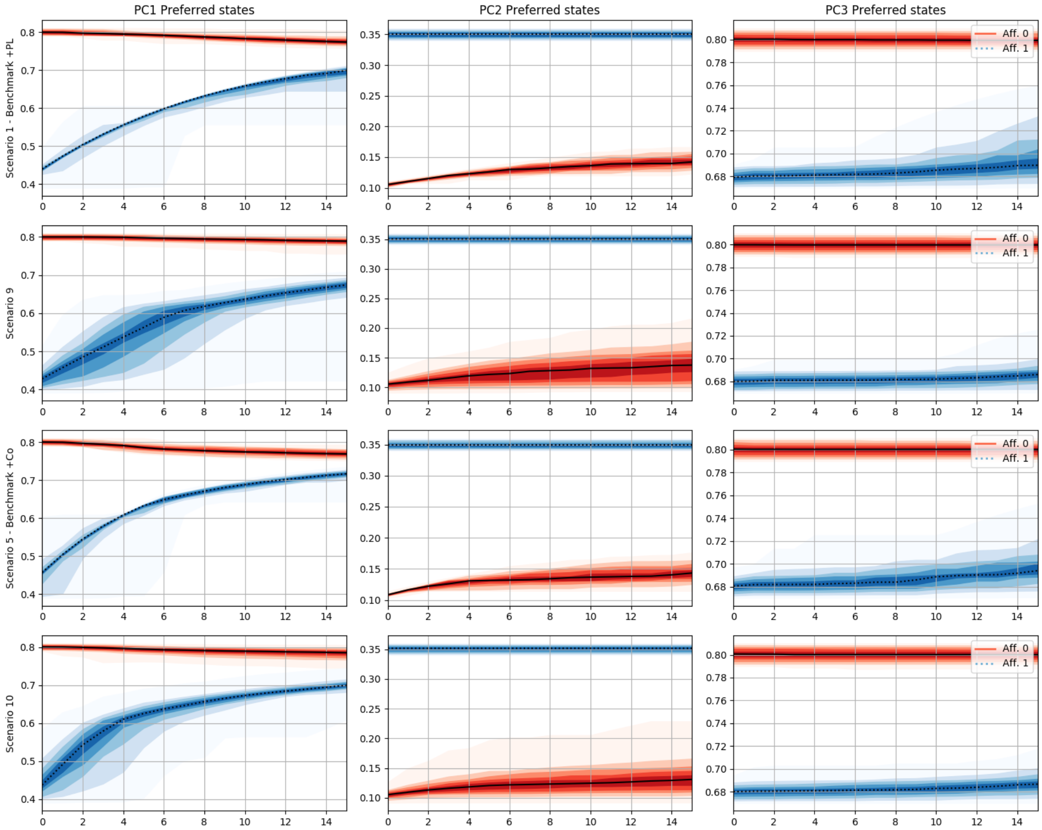

Figure 11.16 – Policy core issue preferred states for scenarios 1, 9, 5 and 10.

11.4 Considering imperfect information (+PI)

The introduction of the imperfect information is perhaps the most complex and uncertain increase of complexity so far. This is because the approach presented in the formalisation, and introduced in the simulation, is not based on any theoretical framework. This complexity level increase is more a thought experiment and demonstration of the common language's robustness and potential. It shows that the common language can be used to test a wide range of additional complexity, some of which is not present in the literature yet.

Similarly to the partial knowledge complexity addition, it is considered as an add-on to the model, as it can be introduced to either the +PL or the +Co models. It can also be considered in combination with the partial knowledge add-on. To test the imperfect information complexity addition, a large set of scenarios and simulations are used. They are presented in this section.

11.4.1 The research question

The goal of the experiments is to explore the impact of the different new parameters and mechanisms introduced. The research question is formulated as follows: *What is the impact of the introduction of the bias, the initial actual beliefs and inter-affiliation trust on policy change when considering imperfect information throughout the policy process?*

Similarly to previous complexity increases, another goal of these experiments is to understand whether this addition of complexity is warranted considering the additional data required to initialise and simulate it. The overall goal being exploration to demonstrate how theory exploration could be performed in the future using the approach presented in this thesis.

11.4.2 The scenarios

The scenarios are designed to test the different new parameters that were introduced with the addition of the imperfect information add-on. Overall, the parameters that are carried over from the benchmark +PL scenario are kept the same. This includes the preferred states, causal beliefs, agent distributions and resources distributions.

The three new parameters are varied between the different scenarios. The bias values are provided in Table 11.3 while the initial actual beliefs are given in Table 11.4. The political affiliation trust weights are given as follows: when communications happen,

	Issues					
	PC1 Sheep	PC2 Wolves	PC3 Grass	S1 Sheep growth	S2 Wolves growth	S3 Grass growth
Agents						
Scenario 11-14						
Affiliation 0	5	0	5	5	0	5
Affiliation 1	5	5	5	5	5	5
Scenario 15-16						
Affiliation 0	0	5	5	0	5	5
Affiliation 1	5	5	5	5	5	5

Table 11.3 – Biased values for the external parties for the scenarios 11 to 16. For a value of 5, the external party informs itself directly from the truth agent. For a value of 0, the external party does not change its beliefs.

the agents in the same political affiliation provide one another with exact information. Agents in different political affiliations only share 75% of the information. When there is no communication, then agents in different political affiliations share 0% of the information.

Note that for the first time, the political affiliations of the agents have an impact on the model simulation with the introduction of imperfect information. It impacts the transfers of information. It remains the primary way the results will be split and analysed later on in this section as well.

The imperfect information add-on (+PI) can be tested with both the +PL and +Co models, along with the +PK add-on. However, because previous results have shown that the +PK add-on has only a little impact on the results, it was decided not to include it in the +PI scenarios.

The focus is to test the imperfect information add-on with the +PL model, following the values of the benchmark +PL scenario - scenario 1. This is done for the sake of comparison between the results obtained in this section and previous sections. The +Co model is also used but only to a lesser extent as results analysed in a previous section have shown its limited impact on the model's outcomes.

The different scenarios considered are provided below with some explanations on why they are considered:

- Scenario 11 - In this scenario, the +PL model is used. It assumes that there is

	Issues					
	PC1 Sheep	PC2 Wolves	PC3 Grass	S1 Sheep growth	S2 Wolves growth	S3 Grass growth
Agents						
Scenario 11-12/14-16						
Affiliation 0	150	50	1225	70	45	200
Affiliation 1	150	50	1225	70	45	200
Scenario 13						
Affiliation 0	150	350	1225	70	95	200
Affiliation 1	150	50	1225	70	45	200

Table 11.4 – Initial actual beliefs on an interval [0,1] for the scenarios 11 to 16.

information communication, though limited, between the two political affiliations. Affiliation 0 agents are not receiving any information on their wolf-related actual beliefs (issues PC2 and S2). The initial actual beliefs of all agents are normal.

- Scenario 12 - This scenario has the same premise as scenario 11 but it uses the +Co model.

- Scenario 13 - This scenario uses the +PL model. It assumes that there is information communication, though limited, between the two political affiliations. Affiliation 0 agents are not receiving any information on their wolf-related actual beliefs (issues PC2 and S2). The initial actual beliefs of all agents are unreasonable.

- Scenario 14 - In this scenario, the +PL model is used. It assumes that there is no communication between agents belonging to different political affiliations. Affiliation 0 agents are not receiving any information on their wolf-related actual beliefs (PC2 and S2). The initial actual beliefs of all affiliations are normal.

- Scenario 15 - In this scenario, the +PL model is used, and it is assumed that there is information communication, though limited, between the two political affiliations. Affiliation 0 agents are not receiving any information on their sheep related actual beliefs (issues PC1 and S1). The initial actual beliefs of all agents are normal.

- Scenario 16 - In this scenario, the +PL model is used. It assumes that there is no information communication between agents of different affiliations. Affiliation

0 agents are not receiving any information on their sheep related actual beliefs (issue PC1 and S1). The initial actual beliefs of all agents are normal.

11.4.3 The results

The results are split into four parts to study the different aspects introduced with the imperfect information add-on. First, I look at the impact of including imperfect information compared to the benchmark results previously analysed. This consists of comparing the results of the benchmark +PL scenario (scenario 1) with the results from scenarios 11, 12 and 15. Then I look at what the impact of considering different initial actual beliefs is through the analysis of scenarios 11 and 13. This is followed by an analysis of scenarios 11, 14 and 16 to understand the impact of imperfect communication between the political affiliations. Finally, the impact of the biases is studied with an analysis of scenarios 1, 11 and 15.

The impact of imperfect information

To look at the impact of imperfect information on the results, scenarios 1, 11, 12 and 15 are compared. The main difference between the scenarios relates to the way the agents obtain information on wolf-related issues (issues PC2 and S2). In the benchmark +PL scenario, the agents have perfect information as they obtain this information directly from the predation model. For scenarios 11 and 12, the affiliation 0 agents are not informing themselves at all on these two issues. In scenario 15, the affiliation 0 agents do not inform themselves on the sheep-related issues (issues PC1 and S1).

The results of this differentiation in information communication are visible in Figure 11.17. Whenever there is perfect information transfer, the actual beliefs of the agents in both affiliation match exactly. The figures also show that whenever information is not transferred, in the case of affiliation 0 agents, the actual beliefs lag, although it appears only to have a minor impact. The reason why the difference between the curves is not larger is that the agents can still obtain some information from the external parties belonging to the opposite affiliation. Furthermore, the initial actual beliefs are believable, reducing the initial gap between the agents in both affiliations.

Overall, this lack of information transfer does not impact the agents' preference selection in any of the scenarios considered. This is logical for scenario 11 and 12 as the actual beliefs impacted are wolf-related issues, issues that are rarely selected by the agents. This lack of information transfer could have had a more substantial impact in scenario 15. Surprisingly, no impact is found in predation model either. The affiliation

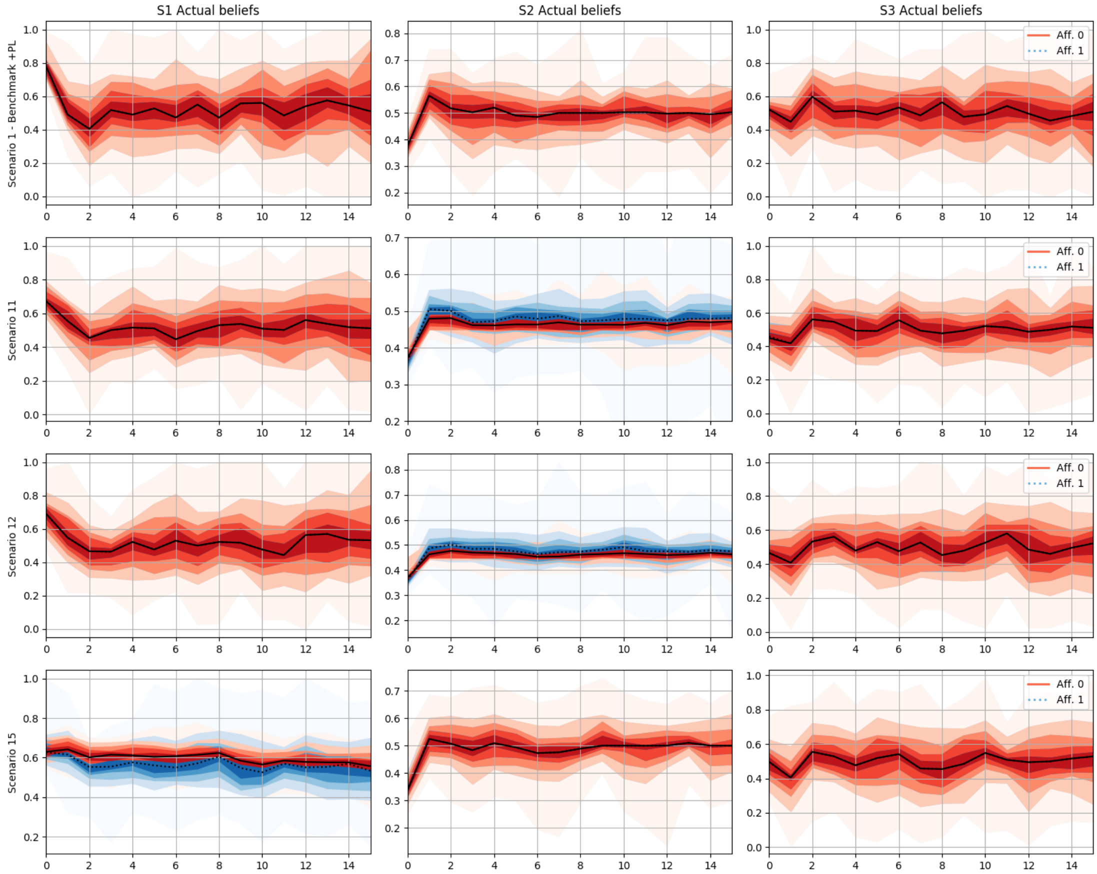

Figure 11.17 – Actual beliefs of secondary issues of scenarios 1, 11, 12 and 15.

0 agents keep selecting the same issues may it be at the policy core or secondary issue levels, even considering the inaccuracies within their decision-making process. This can be attributed to the small difference in actual beliefs between the informed and uninformed agents. This difference is smaller than 0.1, barely affecting the calculation of the preferences.

Finally, one can see from Figure 11.17 that there is no difference between the +PL and +Co results (scenarios 11 and 12). This comforts earlier results showing that the introduction of coalitions only quicken policy learning but does not affect the model significantly. Introducing imperfect information with the coalition does not change this observation.

The influence of the initial actual beliefs

The initialisation of the actual beliefs is expected to play an important role in the outcomes of the model. This is justified by the fact that the initial actual beliefs used are extreme ones and are therefore expected to disrupt the agents' decision-making process significantly. To analyse their impact, scenario 13 is analysed.

The results show that there is indeed an impact, through the entire duration of the simulation. This is shown for both the policy core and secondary issues in Figure 11.18. It is not visible because the figures are plotted using percentiles, but the actual beliefs of the affiliation 0 agents are brought to their extreme levels for each time step as the incorrect information is being conveyed to the agents. Through influencing actions between the agents, the actual beliefs of agents in both affiliations are brought back to a more reasonable level. This is visible in the results. In effect, the agents have no real idea of where these two issues, relating to the wolves, stand. They then base their decisions on the inaccurate values that they have.

This fuzzy knowledge has large repercussions on the rest of the model. First, it impacts the agents' preference selection to the extent where they are unable to even agree on an agenda in a large number of simulations. This means that the agents can never reach the policy formulation step. This is shown in Figure 11.19. In the figure, in the first time steps, the agents are unable to agree, and no instrument is selected. This is why the bar plots no longer add up to fifty at each time step. This, in turn, has disastrous consequences on the predation model results which see the wolf populations collapse in more instances than for any other scenario so far.

The agents do end up agreeing on an agenda after several steps. The agents' decision-making process pushes them to influence one another on their actual beliefs first because they are so different and therefore graded higher. This leaves only a limited

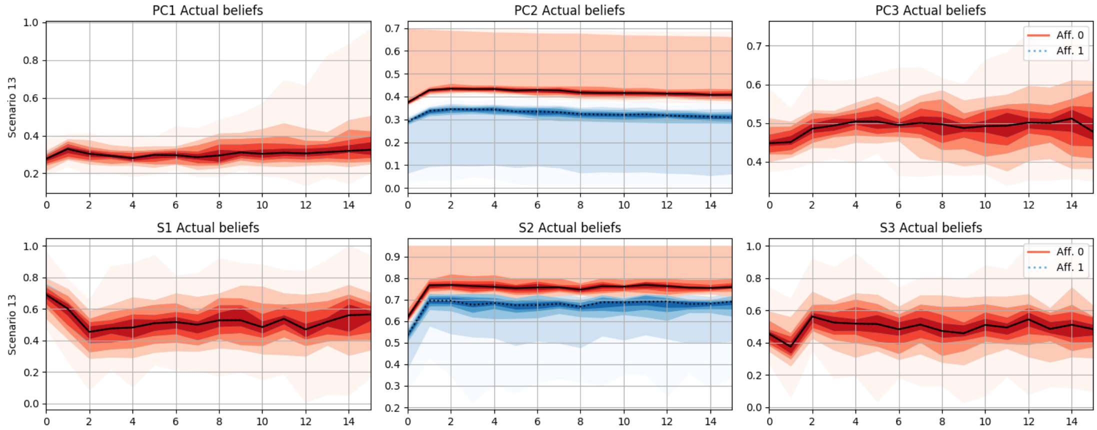

Figure 11.18 – Actual beliefs of the policy core issues of scenarios 13.

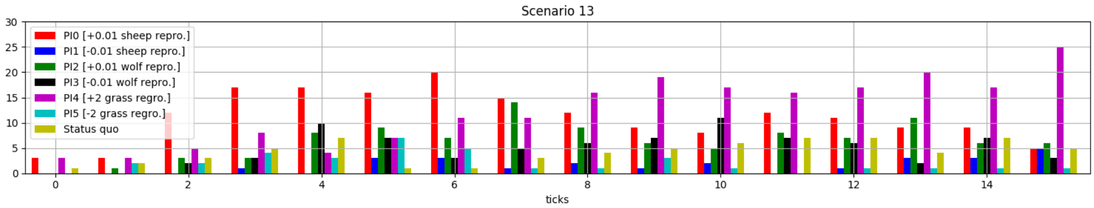

Figure 11.19 – Policy instruments selected for scenario 13.

amount of resources every time step for other influencing actions. This slows down the policy learning happening for the other preferred states but does not halt it entirely. After some time, this policy learning starts to have an impact and the agents return to the consensus observed in prior scenarios. It is only then that the agents manage to form an agenda and get to select policy instruments again. This takes about three-time steps.

Ultimately, after a period of uncertainty, the policy process model returns to some normalcy where the agents prefer policy instrument PI4 like for most of the other scenarios as shown in Figure 11.19. The individual switch from policy instrument PI0 to policy instrument PI4 about midway through the simulation observed in prior scenarios is once again present. This suggests that despite the extreme initial actual beliefs, the model returns to normal because of the robust policy learning driven by the agents and because the agents' actual beliefs are only a small part of their decision-making process.

At the policy process model level, this shows that extreme initial actual beliefs lead to a delay in the trends observed in scenarios where the initial actual beliefs are not extreme. For the predation model, the impact is much more consequential. In most

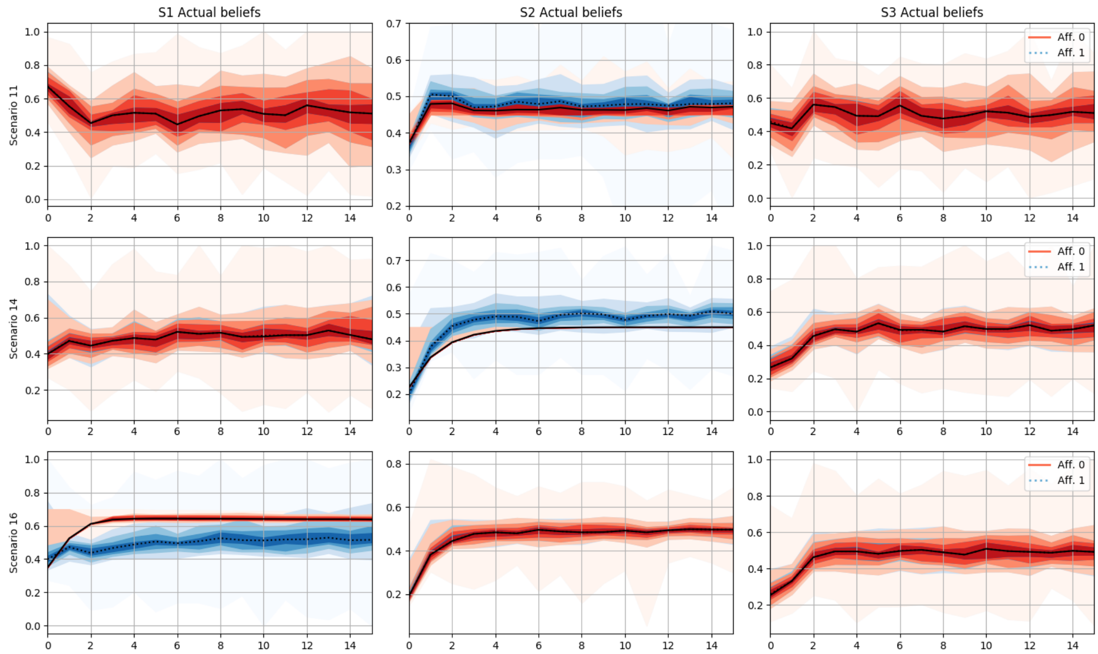

Figure 11.20 – Actual beliefs of secondary issues of scenarios 11, 14 and 16.

simulations, this lack of decision making from the part of the policy agents leads to situations where it is often too late for the predation model to recover, and the wolf populations remain at zero.

Communication between affiliations

To analyse the impact of the communication between the two political affiliations, results from scenarios 11, 14 and 16 are analysed. Scenarios 14 and 16 do not allow inter-political affiliation communication on the actual beliefs for the wolves and sheep issues respectively. The expectation here is an impact though limited as the agents are still able to obtain information from their political coalitions.

Once again, one can see that this lack of communication has an impact on actual beliefs. The actual beliefs concerned by the lack of communication for the affiliation 0 agents stay mostly flat at the level at which they are initialised. This is illustrated in Figure 11.20. In the graph, this is particularly visible for scenario 14, as no policy learning influences the behaviour observed. However, it is possible to see influencing in scenario 16, there is a limited influence from affiliation 1 agents which leads to some changes in the actual beliefs values.

In this case, it is difficult to track how these changes go on to affect the bigger picture.

For the issue preference selection, this lack of information sharing between the affiliations seems to affect greatly affiliation 0 agents in scenario 16, where the lack of information is present for the sheep issues. Affiliation 0 agents go on to continually select PC1 as their preferred issue. This is a result of a large gap between preferred states and actual beliefs for issue PC1. This trend is also present but to a lesser extent for scenario 14, where the lack of information is related to the wolf issues.

These effects can also be partially observed at the secondary issue selection level. For scenario 14, the affiliation 0 agents tend to consistently prefer issue S1, relating to the sheep, though their lack of information is on the wolves. While for scenario 16, the opposite is valid with the agents consistently prefer S3, which deals with the grass patches growth though their lack of information is related to the sheep issues. It is difficult to see any impact of these slight changes in secondary issue selection beyond the policy process model. The policy instrument selection does not change throughout the different scenarios. It is, therefore, impossible to observe changes in the predation model results.

One point that seems to be of importance is that policy learning is affected by these changes in communication and the introduction of imperfect information at large. From the results, one can see that agents from both affiliations will tend to influence the preferred states in the agenda-setting step while focusing on the actual beliefs in the policy formulation step. This limits the changes in the policy formulation step compared to previous scenarios. This is a direct impact of the introduction of imperfect information but can also be attributed to the consistently smaller gaps between agents of different affiliations in the preferred states of the secondary issues compared to the policy core issues.

Biases impact

Finally, the last part of the imperfect information results to analyse relates to the bias parameter. For this, scenario 11 and 15 are considered. One would expect the agents' preferences to be affected by a change of bias from the wolf issues to the sheep issues. This is not observed in the results. There is little difference between the results of the two scenarios.

This can mostly be explained by the fact that the external party instead informs affiliation 0 agents that should not be getting proper information of the other political affiliation. This is slower but only marginally and does not affect the results from the model.

11.5 Discussion

The results stemming for the four levels of complexity and their accompanying sixteen scenarios have been presented and analysed throughout this chapter. The results have allowed for the confirmation of expected behaviours, for the discovery of unexpected behaviours, and they have shown that the behaviours obtained are not always as expected. They have also provided a template for future theory exploration for researcher wanting to use the common language and a modelling approach. In this section, I discuss what these results mean for each of the complexity increases and then look at the models in general, considering the bigger picture aspects of policy process modelling.

11.5.1 The +PL model

The introduction of the policy learning saw the introduction of a new role for the agents and, most importantly, influence actions. The results have shown that this increase in complexity helps simulate crucial new behaviours and gain a better understanding of the dynamics in the policy subsystem. The results have also shown that these new behaviours can affect the predation model, although these effects are often limited.

The introduction of policy learning opens the way to a range of different tests, simulations and experiments, may it be with a standard environment model to explore the policy process itself, or with a more complex environment model to better understand the environment model. In this chapter, the scenarios were build specifically for the testing and exploration of the policy process and the algorithms devised to emulate the ACF and policy learning.

For the case of the +PL model, these tests are possible because of the number of parameters that are introduced. Beyond the tests considered in this chapter, additional experiments can be run as well. For example, power dynamics can be tested in more depth by changing the agents' roles and the distributions of these roles. Resources imbalances can be tested by assigning different levels of resources to the policy agents. Role and resource distributions can also be combined in different manners to create peculiar situations where the agents might behave differently. Finally, there is also the possibility to have more variety in the initial beliefs of the agents. In this chapter, only two groups of agents were considered, but this could be more parcelled out.

The introduction of policy learning through agent influences to study a policy subsystem can be a great asset to any case study. It does require an additional number of inputs, including the preferred states of the agents, the agent distribution and the agents' resources. However, once the data has been obtained, it can be used to understand better how the subsystem might evolve and how the agents might behave. It can be used to test scores of assumptions, and understand the reasons behind agents' behaviours. The policy learning complexity addition is a crucial addition to the policy process model and should be widely considered for the study of the policy process.

11.5.2 The +Co model

The results obtained from the introduction of coalitions to the ACF model were more nuanced and provided only limited insights into coalitions or the policy process as a whole. The coalitions' impacts were in effect limited to policy learning: coalitions helped accelerate policy learning in specific circumstances. For the rest, the scenarios used showed under what circumstances coalitions would be created and their numbers.

Some issues arose with the introduction of coalitions. They included the speed at which coalitions changed. The results showed that the coalitions would combine or change significantly within a limited amount of time - ten-time steps. Though a measure of time is not specified within the simulation, one can assume that an entire simulation would go on for no longer than three to four decades. In such a period, and according to the literature, it is unlikely that coalitions would combine the way they have in the present model. Instead, coalitions have been shown to remain stable for several decades, though this can change depending on the conditions surrounding the policy subsystem (Jenkins-Smith et al., 1991; Pierce, 2011; Weible and Sabatier, 2005; Zafonte and Sabatier, 2004).

Changes to the way coalitions are implemented, including better tuning of the resources, could help deal with this issue. Such an adjustment could necessitate a differentiation between policy learning in the agenda-setting step, and in the policy formulation step, as for the latter, policy learning is assumed to be faster. Introduced into the model, reducing the impact of resources, for the agenda-setting step, for example, might result in the agents' influencing actions having less impact and therefore lead to a slow down of the policy learning as a whole.

I have also shown that these results are directly impacted by the assumptions made when formalising the +Co model. This follows the discussion on the assumptions that came in the previous chapter. These assumptions were kept as simple as possible for the ease of analysis. From the results, it seems this lack of complexity might have gone too far, preventing the emergence of more interesting behaviours that could have been expected from the literature (Weible et al., 2018). This suggests that the approach used to include the coalitions is not mature enough. In the future, the assumptions surrounding the theories might need to codify ways for coalitions to strategise, organise their actions or withhold resources at critical points in the process. This would allow for richer behaviours resulting from the inclusion of the coalitions, making such simulation all the more useful to researchers.

For now, with the current coalition assumptions, I would not advocate for the use of the +Co model even when studying a case where coalitions are present. It provides too little insights compared to the amount of data that is required. Furthermore, the way coalitions are modelled in this chapter can easily be reproduced without the use of explicit coalitions. By providing more resources to key agents that are known to be part of coalitions in the real world, one could simulate similar behaviours as ones that would be obtained through the introduction of coalitions. The use of the +Co model in its current form is only recommended for theory exploration, to test further and advance the model based on what is present in the literature. This includes the development and simulation of more complex assumptions.

11.5.3 The +PK model

The partial knowledge scenarios results showed little change compared to the benchmarks, both for the policy learning and coalition models. This was mostly attributed to the learning rate that was sufficiently high to bridge the initial lack of knowledge on the agents' part. In general, from the results observed, it is questionable whether the introduction of such complexity is useful to gain additional insights.

As was mentioned, the learning rate could be modified to test its impact on the outcomes of the model. However, it can be expected that it would only have limited effects except in cases where the learning rate is close to zero and agents have no way to learn about other agents' beliefs. Then the impact would be mostly dependent on how the partial knowledge values are initialised. Ultimately it would render the policy learning random and inconsistent, effectively hindering or neutralising it.

Considering this analysis, I would not recommend the use of the partial knowledge add-on in general cases. Specific cases and contexts, where partial knowledge is known to have a critical impact on the outcome of the policy process, might benefit from this add-on. In such a case, the initial partial knowledge data would have to be collected before the model can be simulated. This could be a case where the political arena is highly polarised and where agents refuse to learn from their peers. Assumptions such as differential knowledge gain speeds could also be considered in such circumstances, though at the costs of additional data requirements. Further work could be performed to enhance this add-on and add aspects that might have been omitted.

11.5.4 The +PI model

As mentioned previously, the introduction of the imperfect information add-on into the policy process was the most complex complexity increase presented within this chapter. An extensive set of scenarios was used to test it with mixed results. First, the use of this add-on requires a lot of data to be used properly. This includes knowledge on the agents' bias for the external parties but also all of the agents' initial actual beliefs, data that can be hard to obtain. Without proper data, and if the add-on is not used for theory exploration purposes, the use of this add-on is not expected to provide much value added to the researchers.

Similarly to the two previous levels of complexity, the results of the +PI model had mixed outcomes. Overall, it was difficult to notice significant changes to the wolf and sheep populations, or with the policy instruments selections. Some scenarios did manage to reveal specific cases where the policy process and the predation models were both significantly affected by imperfect information, but these were niche cases where several specific conditions needed to be present.

Throughout the scenarios simulated, it was clear that imperfect information can have an impact on the agents' actual beliefs. However, these are smoothed over quickly by the agents influencing one another. More often than not, these impacts failed to make any difference when it came to the agents' issue preferences selection. Because of this, they had no impact on the agenda, or on the policies that were selected.

There are several reasons why these behaviours are observed. One is linked to the predation model. The model is a stable one with a lot of inertia as it has been outlined several times before. Even when the wolf population reaches zero, the model can remain stable with only sheep and grass agents left. This is further reinforced by the incrementality of the policy instruments used. The indicators of the predation model, which help inform the agents' actual beliefs, are therefore limited in their change. In turn, this limits the change that can happen within the policy process itself, regardless of the events fed into the policy process model.

This issue of inertia can also be linked to the layering of complexity within the policy process model. If only one parameter in the agents' belief system changes, it will often be subdued by the amount of complexity in the model before it can affect policy change. For example, a change in the actual beliefs needs to be large enough to change the preference of the agents, and this preference change must be strong enough to help change the overall agents' preference selection. This then needs to happen in enough agents' belief system for the agenda or the policy selected to be affected. Change in policy is even harder to enact when the original change happened at the policy core level. If anything, a change in a parameter at the policy core level might affect the agenda without more impact onto the rest of the model.

These remarks come with a caveat intricately linked to the predation model. The simplicity of the predation model means that only three secondary issues could be identified, each being related to the policy core issues. This makes the agenda mostly irrelevant as, regardless of the policy core issue placed on the agenda, the agenda does not serve its primary purpose of narrowing down the policy instruments that can be considered. In effect, this means that a change in a policy core issue parameter that leads to a change in the agenda will not affect the rest of the model and is unlikely to lead to a change in the policy implemented. Only a direct change in a parameter at the secondary issue level could do that.

In an environment model with more than three secondary issues and where these are not all related to all policy core issues, this might not be necessarily the case. In such a case, a change of agenda resulting from a change in a policy core issue parameter would restrict the choice of policies that can be chosen from and might, in turn, affect the policies implemented. This would, in turn, be more likely to have an impact on the environment model.

The results of the +PI model have also shown that it would be difficult to simulate a fake news scenario with the mechanisms put in place. The model is too permissive. Even if false information makes it to the actual beliefs of the agents, they will be influenced away from their false beliefs by other agents, partially or completely depending on the resources of these agents. One way this could be fixed would be to restrict the influence actions or to remove entirely actual beliefs influencing actions. However, this would be hard to systematise into a framework and would require considerably more work and data to simulate. Another way would be to tweak the resources such that the agents have fewer resources to influence the actual beliefs. This relates to a topic that has been recurring throughout the scenario, and that is discussed later in this discussion: resources tuning.

Another aspect that this add-on has shown is the importance of the initial actual beliefs, placing more pressure on the modeller with the data collection effort. As was shown, faulty or improbable initial actual beliefs can lead to simulations where the agents do not reach an agreement on the agenda. It seems critical that, especially for use in a case study, the initial actual beliefs be collected from the agents directly to obtain trustworthy results.

Similarly to the two previous complexity increases, the imperfect information add-on suffers from several flaws and limitations. However, this add-on does show potential. It would be interesting to test it with a more complex environment model where imperfect information is more likely to influence the outcomes of the model. This would allow modellers to understand further how useful this add-on is, something that was not possible in this chapter with the predation model.

As a final remark, one of the initial reasons for the inclusion of partial knowledge was to demonstrate how one can study the impact of the media on the policy process - and maybe look into the impact of fake news and other new phenomenon. One assumption was that external parties often are, amongst other possible actors, media actors and they frame certain issues in ways that might distort reality. One aspect that is not considered with the present approach is the impact of the media on the electorate, a topic studied at length in the literature (Arceneaux et al., 2016; Birkland and Lawrence, 2009; Muis, 2010; Shanahan et al., 2011). Such studies could benefit from the use of modelling and the overall approach proposed in this thesis with the common language. It could be used to study in more depth how this impact might affect policymakers and policy change. Such study could use some of the parameters and behaviours that were also mentioned with the integration of the electorate in the simplest implementation model based on the work presented in Laver and Sergenti (2011).

11.5.5 General remarks

There are certain aspects of the insights gained throughout this chapter that should be discussed in general as they pertain to the ACF implementation model as a whole. They include some of the more detailed inner-workings of the algorithms that are used to simulate this model.

The policy instrument selection

One limitation relates to the agent-level policy instrument selection. This selection is flawed for two reasons. First, the policy core issues have minimal impact on the instrument's selection. Although this is not paramount to the model, it is questionable. This was already discussed in the +PI model discussion above.

Second, the policy instrument selection does not consider the agents' preferred secondary issue. Instead, it is only based on the policy's overall impact on all of the secondary issues that are linked to the agenda - in the present model this means all three secondary issues. This completely sidelines the agents' secondary issue selection. This is a problem for two reasons.

First, the agents can only influence the actual beliefs, preferred states and causal beliefs related to their preferred secondary issue. This means that if they decide they want to advocate for a particular policy instrument, they will only be able to affect a small part of the parameters that help shape the grade for that instrument in other agents' belief system - that is the actual belief, the preferred state and the causal relations linked to their preferred secondary issue.

Second, the agents cannot influence the impact of the policies on secondary issues. In that part of the model, full information remains regardless of the complexity increases included. Every agent has the same information on the impact of all policy instruments. This was not included in any of the models presented above, although it could have been included in the +PI model. This further limits the agents' decision making power and their ability to influence other agents to select their preferred policy instrument.

There are two ways this could be modified to address better the issues outlined. First, it could be possible only to grade policy instruments based on their impact on each agent's preferred secondary issue and not all issues related to the agenda. This might lead to more variation in the policy instruments selected, especially if the agents' secondary issue preferences vary over time. This is something that is not currently the case with most scenarios being dominated by a single policy instrument. However, it could also lead to the selection of less useful policy instruments. The instruments selected will not make the entire system better anymore, but they will be aimed at improving a single issue. The underlying assumption would then be that agents only care about their preferred issue when selecting a policy instrument, something that is questionable.

An alternative change would be to introduce a new influence action, one that could see agents influence one another on their policy instrument impact perception. Similarly to the other influence actions, the agents would be able to influence one another on their perception of the policy instruments' impacts. It would, in effect, remove the perfect information assumption that is currently present for the policy instruments. By adding influence actions on these impacts, the results would better represent the often different agent expectation of the policies' impacts. To better introduce such action, the policy instruments impacts would have to be operationalised in the same way as actual beliefs. This would be done similar to the actual beliefs. The policy instruments impacts would have to be reassessed and communicated to the agents every time step.

One could go further and place such a change into the +PI add-on. Agents could see their information on the policy instruments impacts be communicated to them by the external parties and be subject to biases the same way the actual beliefs are. This would be a better representation of reality, although it would introduce much complexity at once, making it difficult to test. Overall, this shows that there is no one way to operationalise the policy process but a range of approaches, each propped up by different assumptions.

The agent resources

Concerning the agent resources, several issues have been encountered regarding their tuning - the resources that are given to each agent at every time step. This is a topic that was already discussed in the ACF's formalisation chapter (see Section 10.5) before any simulations were performed. There it was decided to set a maximum amount of resources given to each agent such that the policy learning rate is limited. To a certain extent, this has worked, and the results obtained from three of the four complexity levels were acceptable. However, the +PI model showed that maybe this level was too high. The agents have enough resources to influence all of the actual beliefs of their peers and still be able to influence the preferred states beyond that. This does not have to be considered a problem, but it should be questioned.

I would suggest that, in future research, a systematic approach be developed to tune the resources provided to the agents. Such an approach can be tested using the model presented here but would need to be refined through testing on subsequent real cases. The resource parameter is one of the most important of it sets the rate of policy learning, and policy learning is the primary endogenous driver of policy change in the policy process model developed in this dissertation. The resources distribution can be used to calibrate models being simulated and, in some instances, could be used to validate that the models built are proper representations of reality.

All of this must also take into account the issue of definition and data collection. The resource parameter is at the centre of the model, but it has been difficult to define it up to now. As mentioned in the common language (see Chapter 3), resources in the literature can vary in definition from financial resources to political resources. Some of these can be quantified while it is much more challenging to quantify others as they mostly involve human resources. To simulate the policy process, all resources need to be melded into one parameter.

This lack of a strict definition also affects data collection. It is complicated to collect data on something that is wide-ranging and imperfectly defined. This, in turn, has repercussions on the data the modellers have available to use for the simulations. Ultimately, the goal might not be to have exact numbers for the resources of each agent in the policy subsystem. Instead, the goal is to have the correct relative resources between agents such that power relations, as far as resources are concerned, be well represented within the model.

The experimental limitations

The scenarios that have been explored in this chapter are limited in their scope due to the nature of the work presented here. The goal was to explore the models presented, but because of the size and possibilities present with the ACF implementation model, the goal was not to perform an exhaustive test of all possible scenarios. Based on the results obtained, it is possible to guide the future work, highlighting some of the parameter combinations that are expected to result in behaviours of interest. This includes, but is not limited to, further tests on the balance of power - this time by varying the number of agents in the different roles, the combination of biases for the +PI add-on or the selection of specific initial actual beliefs for this same add-on.

Another avenue for model testing is to use more Monte Carlo simulations. This should be done with care and the necessary boundaries. If not, considering the number of parameters present in the model, this could lead to a large number of simulations, making it more challenging to analyse the results. Furthermore, with two many parameters varied at the same time, it might be difficult to understand the links between parameter changes and observed behaviours. The model would be too complex to understand, effectively nullifying the use of Monte Carlo simulations.

Another approach would be to incrementally advance the model, in a way similar to what was presented in this chapter, but at a more detailed level. In this chapter, and through the previous two chapters, the ACF implementation model was built up through the addition of complexity. This was done to be better able to understand the impact of additions brought to the model. This has provided an overall architecture for what a model of the ACF could look like, but it has also led to several assumptions that were not entirely tested. Through a more incremental approach, different parts of the model could be reevaluated, and different choices could be made to see which assumptions are best suited and lead to the most literature-representative results. These assumptions would have to be informed by the literature, and when the literature is lacking, they could go to inform the literature. A number of these choices, and their alternatives, were already mentioned within this discussion.

Advanced agent strategies

The results have shown that when agents have different preferences, they will focus on different issues. For example, affiliation 0 agents all focusing on the sheep related issues while the affiliation 1 agents focused on the wolf-related issues. This can result in the agents only influencing one another on issues where they have the same beliefs. This means impact-less influence actions and a lot of wasted resources. This was observed many times with the predation model.

Such issues might justify the need for the consideration of more advanced decision making process algorithms to determine how the agents should spend their resources. This might include more advanced agent strategies such as strategies where agents would influence their peers based on the expected advancement of their interests as a whole.

At the moment agents perform actions that they presume will advance their interests. However, because these actions are limited to their preferences, this is not always the case. Within their decision-making process, the agents do not have enough visibility to assess the entire situation. In many cases, acting on a different issue might have had a more significant overall impact on the decision making of another agent, and therefore being more effective overall. Allowing such strategies would mean assuming that the agents have an advanced decision-making process and they can think ahead.

It is unclear whether this is something that could fit with the current advancement of the theories. It might fit with parallel literature focused on agency and decision making. This could be combined with strategies mentioned before, where agents do not necessarily spend all their resources for every step, but they might conserve resources when they are satisfied with the current state of the model.

The model limitations

Though the predation model is used to test the policy process, using several scenarios for different levels of complexity, it has become clear that it is not a perfect model for testing the policy process. There are probably no perfect environment models to test such a model of the policy process. Each environment model will have specificities that will affect and be affected by the policy process model presented in this dissertation and more specifically, in this ACF implementation model. As was explained before, the predation model has a certain inertia which means that in many cases, the scenarios will not lead the model to exhibit significant changes following substantial changes in the policy process model. It is therefore encouraged to test this policy process model with a range of environment models for theory exploration purposes.

The different contexts might exhibit different behaviours which would be beneficial to the advancement of the theories of the policy process.

11.6 Conclusion

In conclusion, this discussion has shown that this policy process demonstration with the predation model is only the start of the study of models of the ACF. It has also shown how theory exploration can be performed using a modelling and simulation approach. The ACF implementation model presented in this dissertation has the potential to exhibit more complex behaviours than the simplest implementation model. It can also be used with a wide range of complex environment models.

The advice I would give considering the results obtained in this chapter is further to test this model with a variety of environment models to gain a better grasp of the model. Furthermore, adding complexity is not necessarily needed or even welcomed, and it might detrimentally affect the results. This was clearly shown for several complexity increases in this chapter. Beyond this, when using this model, it is essential to understand its limitations and gauge whether these limitations affect the results that are being sought by the modeller. The next chapter is such an example of another implementation with a different environment model, the Swiss electricity market.

12 The electricity market model

To further illustrate the ACF implementation model, the model is now combined with the Swiss electricity market model. The goal of this simulation is to show how the ACF implementation model can be used with a complex environment model. The electricity market model was used before and therefore is not detailed again in this chapter. There are no changes to the hybridisation process since it was shown in the first sections of Chapter 7.

The results obtained in Chapter 7 have shown that, regardless of the policies being implemented, a large amount of CCGT power plants would be introduced with the demand growth assumptions considered and the emissions would likely ultimately rise. In the discussion it was highlighted that this would not be the case in real life as Switzerland would be unlikely to build new CCGT power plants. The model did show that the current system would have to change to avert such outcomes. It was also highlighted that the simplest implementation model suffered from a significant limitation: the agent's beliefs were constant throughout the entire simulation. This limitation is an artefact of the simplest implementation model that is not found in the ACF implementation model. The present chapter tests whether using the ACF implementation model might lead to different results and insights.

12.1 The research question

The goal of the ACF implementation model is to see whether the outcomes of the model are affected by the introduction of agent interactions and the possibility to create coalitions. Considering that this might not happen simply through agent interactions, a hypothesis is formulated as follows: without additional resources to the right coalitions or a shift in the views of the pro-economy coalition, changes from the simplest implementation model will be impossible to observe. This hypothesis takes into account the two main new parameters that are introduced in the model with the ACF implementation model.

The research questions follows from this hypothesis as: *Can the use of the ACF policy process model with coalitions help get more insights into the Swiss electricity market model compared to the simplest implementation model?*

12.2 The scenarios

A number of scenarios are built to answer the research questions formulated above. Similarly to what was done for the simplest implementation model, the scenarios need to take into account scenarios for both the electricity model and the policy process model.

For the electricity market model, the scenarios considered are the same as the ones that were considered previously. No change is made to the electricity market model, so there is no reason to change the scenarios used. This should also allow for a better comparison between the results obtained in this chapter and the ones obtained with the simplest implementation model. The scenarios reflect the drive for electrification. This translates into electricity demand growth scenarios of 0%, 1.5% and 3%. For the initialisation inputs, the same values are used as before.

For the policy process, the scenarios are overhauled considering the new parameters introduced by the ACF implementation model. Before the scenarios are constructed, it is essential to decide which ACF complexity level needs to be used. Four versions were presented, including two add-ons. Unlike for the predation model, the aim is not to test the potential of the policy process but instead to answer a question specifically on the electricity market model. Therefore, only one complexity level should be selected.

Considering the Swiss context, it is possible to assume that agents have perfect information and full knowledge. This is because the electricity or energy subsystem is not a polarising issue in Switzerland, at least not to the point that there might be knowledge obfuscation. Furthermore, a lot of the data on the electricity sector is obtained from the government and available online freely. Thus, it is harder to distort such realities. Of course, this assumption is a little optimistic, but it should not materially affect the outcomes of the model. This means that add-ons for partial knowledge and partial information do not need to be considered.

On the other hand, coalitions are present as is shown in Markard et al. (2016). Two prominent coalitions exist within the subsystem: a pro-ecology coalition and a pro-economy coalition. Because of this, it was decided to use the +Co model, which includes the policy learning aspects of the model[1]. Furthermore, the model initialisation is made with two affiliations that do not necessarily represent the Swiss political parties, but that represent the coalition beliefs following the insights gained in Markard et al. (2016).

Now that the +Co model has been selected, it is possible to detail the individual parameters. Contrary to the simplest implementation model, no electorate is considered here. The focus to answer the research question is to change the resource distributions and the beliefs of the agents. Additionally, several parameters such as agent role distribution, causal beliefs, policy core issue of interest for the coalition also need to be selected. These parameters are kept constant for all scenarios. According to Markard et al. (2016), the resources of the pro-economy agents are more extensive (this difference is not quantified in the paper), and they also benefit from a more significant number of actors. The scenarios devised here should test a change in this state of the subsystem to show whether this might impact the electricity market.

For this, three scenarios are considered:

- Scenario 0 (Benchmark) - This is a scenario that can be used as a benchmark. The initial parameters are not changed in any significant way, and no external events are considered.
- Scenario 1 - In this scenario, the preferred state of the agents are changed during the simulation. This is done in an attempt to see if a drastic change in the direction of the goals would make a difference in the model's outcomes. In real life, this could be associated to an unknown external event such as the country being affected by a crisis resulting from climate change-related events, for example.
- Scenario 2 - For the third scenario, the agent resources are changed. This is again done to see what impact this might have on the overall outcomes of the model. The pro-ecology agents are provided with more resources than the pro-economy agents. In real life, this could happen due to a "green wave" during an election cycle, for example.

For all scenarios, the preferred states are given in Table 12.1 including for the scenarios

[1] Note that this decision was taken before the conclusions from the predation model showing that the introduction of coalitions only marginally accelerates the policy learning were drawn. Due to the time it would have taken to simulate the electricity model again with only the +PL model (several weeks), it was decided not to rerun the model. This is taken into consideration within the results analysis.

with external event. The causal beliefs are given in Table 12.2. The salient policy core issue is issue PC1 - the economy. This is chosen because of the current state of affairs in Switzerland with agents coming together depending on their beliefs, mostly on the economy. To inform the agent distribution selection, the paper from Markard et al. (2016) is used. This allows for an estimated distribution of the agents that is not meant to reflect exactly the real life data. In the paper, a distribution of actors is provided based on surveys and interviews. The distribution of actors along the coalitions is given as follows:

- Pro-economy coalition:
  - Policymakers (4): BDP, CVP, FDP, SVP
  - Policy entrepreneurs (15): Economiesuisse, EV, SGV, Swissmem, Alpiq, Axpo, BKW, EWZ, Swissgrid, Swisspower, Energieforum, IGEB, Swisselectric, VSE, ETH-Rat RKGK
- Pro-ecology coalition:
  - Policymakers (3): GLP, GPS, SP
  - Policy entrepreneurs (8): SGB, Swisscleantech, Pro Natura, VCS, WWF, AEE Suisse, SES, AkadWiss

Here the political parties are considered as policymakers. This does not precisely translate the representativity that the parties have in parliament. However, it does represent the fact that, overall, the pro-economy coalition actors are more numerous than the pro-environment coalition actors policymaker wise. This places them in a position of power for the agenda and policy instrument selection in the model. For the policy entrepreneurs, I consider that the ratio of entrepreneurs per coalition is a good representation of the power dynamics in the Swiss electricity policy subsystem for the model. It can be used as is, as it would accurately represent the differences in the actors.

To conclude, the agent role distribution is given as follows: for the pro-economy agents, four policymakers and fifteen policy entrepreneurs, and the pro-ecology agents, three policymakers and eight policy entrepreneurs.

Markard et al. (2016) does not provide any information on the resources distribution between the coalitions. Instead, it is left as a further work aspect. Assumptions, therefore, need to be made for the present model and simulation. For scenario 0 and 1, the resources distribution are the same. It is assumed that the pro-economy agents

	Issues				
	S1 RES [%]	S2 Price [CHF/MWh]	S3 REI [%]	S4 Dom. em. [tons]	S5 Imp. em. [tons]
Affiliations					
Scenarios 0-2: In year 2016					
Pro-economy	60	50	70	4 000 000	6 000 000
Pro-ecology	100	75	100	0	5 000
Scenario 1: In year 2021					
Pro-economy	95	55	95	500 000	750 000
Pro-ecology	100	75	100	0	5 000

Table 12.1 – Preferred states for the policy agents on a the interval [0,1] for all scenarios.

Pro-economy agents			**Pro-ecology agents**		
	PC1	PC2		PC1	PC2
-S1	0.00	0.25	-S1	0.35	0.45
-S2	0.75	0.00	-S2	0.75	0.15
-S3	0.25	0.25	-S3	0.35	0.55
-S4	0.00	0.25	-S4	0.25	0.45
-S5	0.00	0.25	-S5	0.25	0.45

Table 12.2 – Causal beliefs for the policy agents on an interval [-1, 1] for all scenarios.

have more resources than the pro-ecology agents to the tune of 100:75. For scenario 2, it is considered a somewhat extreme reversal, with the pro-economy - pro-ecology resources distribution at 50:100.

12.3 The results

The results are presented in this section. Nine scenarios are being analysed with 50 simulations for each. This includes three demand scenarios for the electricity market model and the three scenarios that were devised for the policy process model. Similarly to Chapter 7, the results are displayed using some figures used before. For the electricity market model, this includes graphs displaying the amount of electricity production per technology along with the demand as a function of time and figures tracking the electricity prices. For the policy process, the graphs either show the policy instruments selection, the agents' issue preference selection or the agents' preferred beliefs evolution.

First, I look at the policy agents' preferences and how policy learning might influence them in different scenarios. This is followed by a description of the policy instruments' selections to gauge the impact of policy learning if any. Finally, this is linked to the electricity market model's effects and how those results evolve with the policy instruments selected by the policy agents.

Note that all figures are not shown in this chapter for conciseness. They can all be found in Appendix F.

Policy learning

Policy learning is present throughout all scenarios. At the policy core level, it is mostly limited. This is a result of the proximity between the preferred states of all of the agents. Most of the policy learning happens on the second policy core issue, for the ecology parameter (PC2). In the first two policy process scenarios, the pro-ecology agents are influenced towards the beliefs of the pro-economy agents. This is a result of the more significant amount of resources that these agents have.

For the third scenario (policy process scenario 2), where the resources are inverted, another trend can be observed (see Figure 12.1). Agents from both affiliations seem to converge to a mid-point though the pro-economy affiliation retains a slight edge. This is because, despite the additional resources provided to the pro-ecology agents, they have a deficit of agents. As a whole, they still have fewer resources by a narrow margin, the resources being distributed on a per-agent basis and not on a per-coalition basis.

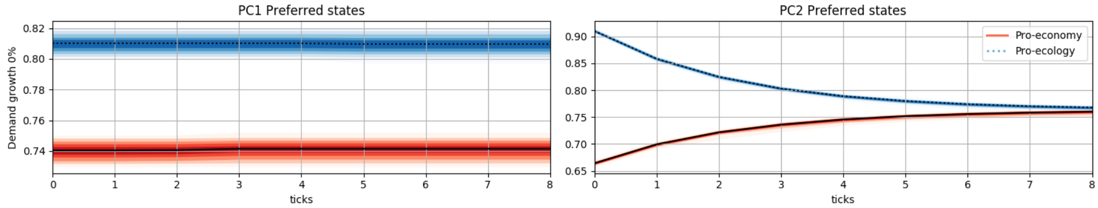

Figure 12.1 – Policy core issue preferred states for scenario 2 and a demand growth scenario of 0%.

Policy learning at the policy core level is not affected by the different electricity demand growth scenarios. The results remain the same regardless of what is happening in the model, possibly because the actual beliefs barely change throughout the simulation. Furthermore, all of the actual belief values are very close meaning any small change might not be made visible.

This is different for the secondary issues where policy learning is much more present, though less consistently, making the figures more challenging to decipher. For the secondary issues, policy learning is influenced both by the policy scenarios and, to a much lesser extent, by the growth demand scenarios. Scenario 1, with its change in the preferred states, has drastically different results compared to the other two scenarios.

In Scenario 1, the driving force of policy learning is linked to the change in the preferred states induced by the scenario. This is most visible with secondary issues S1 and S5. Initially, the pro-ecology agents are being influenced down on the imported emissions preferred states (issue S5), lowering their goals of lower emissions. At the same time, these agents are influencing the pro-economy agents on their renewable energy goals (issue S1), increasing their renewable production aspirations. Minimal influences accompany this on the other secondary issues. Once the pro-economy agents change their preferred states, at time step 2, the preferred states of all agents become much closer to each other. In effect, policy learning is then blunted as the agents agree with one another. From then on, there is no more policy learning.

For the other two scenarios, the initial trends are continued to the end. Policy learning is present throughout the entire simulation. However, first, it is important to look at the issues that are selected by the different agents depending on their affiliations. These choices are reflected in the actions that are performed. When an agent prefers a certain secondary issue, one can directly see that same issue being influenced figures. For example, for scenario 0 - the benchmark scenario - the pro-economy selects secondary issue S3 as a preferred issue (see Figure 12.2). This is reflected in the

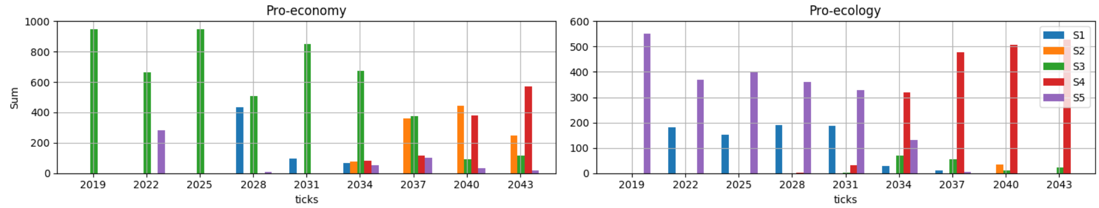

Figure 12.2 – Secondary issue selected per affiliation for scenario 0 and a demand growth scenario of 3%.

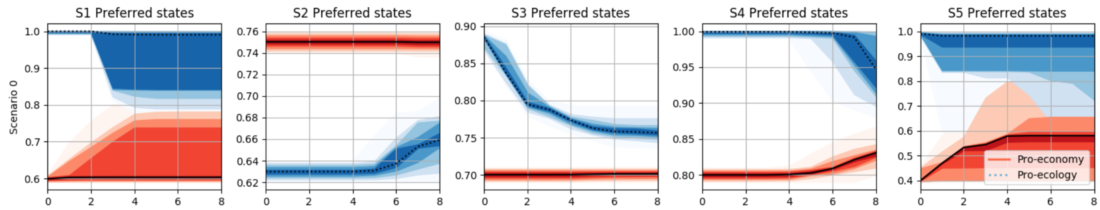

Figure 12.3 – Secondary issue preferred states for scenario 0 and a demand growth scenario of 3%.

progressive evolution of issue S3 of the pro-ecology agents (see Figure 12.3). At the moment the pro-economy agents select S5, in the year 2022, changes appear in the preferred states of the pro-ecology agents for secondary issue S5. Finally, looking at the pro-ecology agents, they select secondary issue S1 only in the second time step, the year 2022. This is also reflected in the evolution of the issue.

Overall, policy learning is most present at the secondary issue level, where the agents start with beliefs that are further apart. This does not mean that the agents are more convincing in these cases, in fact, the policy learning speed is an artefact of the implementation as the influence is based on the gap between the preferred states of both agents and not on a constant value. This explains why policy learning is particularly high for the S1 and S5 secondary issues in Figure 12.3.

Issue selection

Policy learning has only a minimal impact on the agents' issue selections. For example, one would expect for the pro-economy agents to influence the pro-ecology agents such that they end up choosing similar issues both at the policy core and the secondary levels. For all scenarios, it is not apparent that this is the case on either level. The pro-ecology agents either already have the same preferences as their peers, or they do not choose their preferences depending on the pro-economy agents' preferences, which would be a result of policy learning. There is no convergence in the preferences of the agents of both coalitions as is evident throughout the results, including in Figure 12.3. This suggests that the preferred states do not so much drive the preferences, but they are driven by the changes in the actual beliefs, a parameter that the agents only have an indirect influence on.

Policy instrument selection

The influence of the actual beliefs on the agents' preferences is partially confirmed when looking at the policy instruments that are being selected for each scenario. I analyse the selection of these policy instruments here, on a per-policy scenario basis.

In the case of the policy process benchmark scenario - scenario 0, and for the 0% demand growth scenario, the agents select a variety of policy instruments with no instruments being overwhelmingly chosen. This seems to be motivated by a dispersion of the agents on their beliefs and the fact that the goals of the majority of the agents seem to be broadly met. In this case, the agents see their instrument selection vary as they consider which instrument might lead to the best results concerning their goals.

This is very different for the 1.5% and 3% demand growth scenarios. The severity of the scenarios can be seen directly through the change in the policy instrument selection. In those scenarios after three to four time steps, still for policy process scenario 0, one policy instrument is being overwhelmingly selected. This is the policy instrument that increases the domestic carbon tax. This reflects an assessment by the agents that this is the only policy that can help them reach their goal, may they be pro-economy or pro-ecology. This assessment is the result of their actual beliefs more than their preferred states. The agents' actual beliefs come to dominate their belief system regardless of their preferred states because the amount of emissions in these high demand growth scenarios is so high. They drive the widening of the gap between preferred states and actual beliefs. In the 1.5% demand growth scenario, it seems that the agents first believe that the increase in the imported carbon tax can help. This is most likely due to the delay in the CCGT power plants construction that comes in later in that scenario. For the 3% demand growth, the construction of the CCGT power plans happens much earlier.

This pattern of choosing predominantly an increase in the domestic carbon tax for the 1.5% and the 3% growth demand scenarios is something that also happens for policy process scenarios 1 and 2. The reasoning behind the agents' choices is the same each time. This shows that this is a reaction driven by the electricity market model and that despite a change in the beliefs or resources of the agents, the electricity market system is in such a state that the agents have no other options but to choose a policy

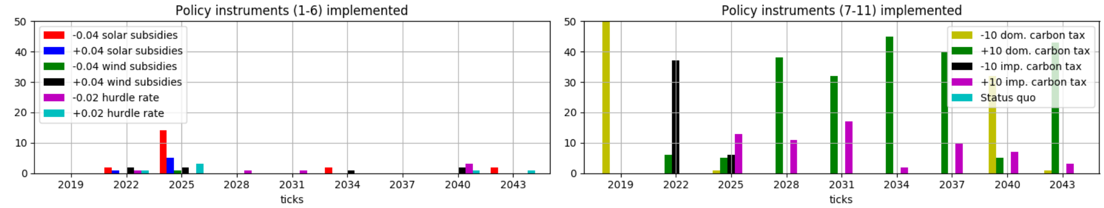

Figure 12.4 – Cumulative policy instrument selection for each demand growth for all fifty repetitions for scenario 1 and a demand growth scenario of 0%.

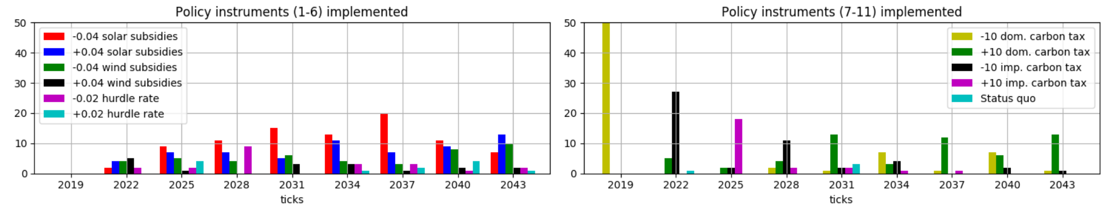

Figure 12.5 – Cumulative policy instrument selection for each demand growth for all fifty repetitions for scenario 2 and a demand growth scenario of 0%.

instrument related to the emissions' tax. This behaviour also illustrates the powerlessness of the agents regarding the policy instruments at their disposal. Whatever they implement, it has little impact on the outcome of the electricity market model. Regardless, the agents have no other choice but to keep selecting and implementing the same policy over and over.

Looking at the results from the two scenarios where the policy agents are modified, little significant differences can be found. The impact of the belief changes - scenario 1 - is mostly seen in the 0% demand growth scenario where the agents seem to come to the same conclusion but this time much earlier. This is illustrated in the much larger policy instrument selection majorities in Figure 12.4. The predominant policies are once again the ones related to the introduction of a carbon tax, either domestic or foreign. This is repeated and amplified for the other growth scenarios.

For scenario 2, where only the agents' resources are changed, the change in the agents' beliefs is much slower. This leads to a minimal change in the policy instruments selected in scenario 0 for the 0% demand growth scenario, as shown in Figure 12.5. This means that the change of resources has no impact on the policy instruments selected. This change is most likely not drastic enough, not affecting the power balance in the model. Only a drastic change, such as the one shown in scenario 1 makes a difference.

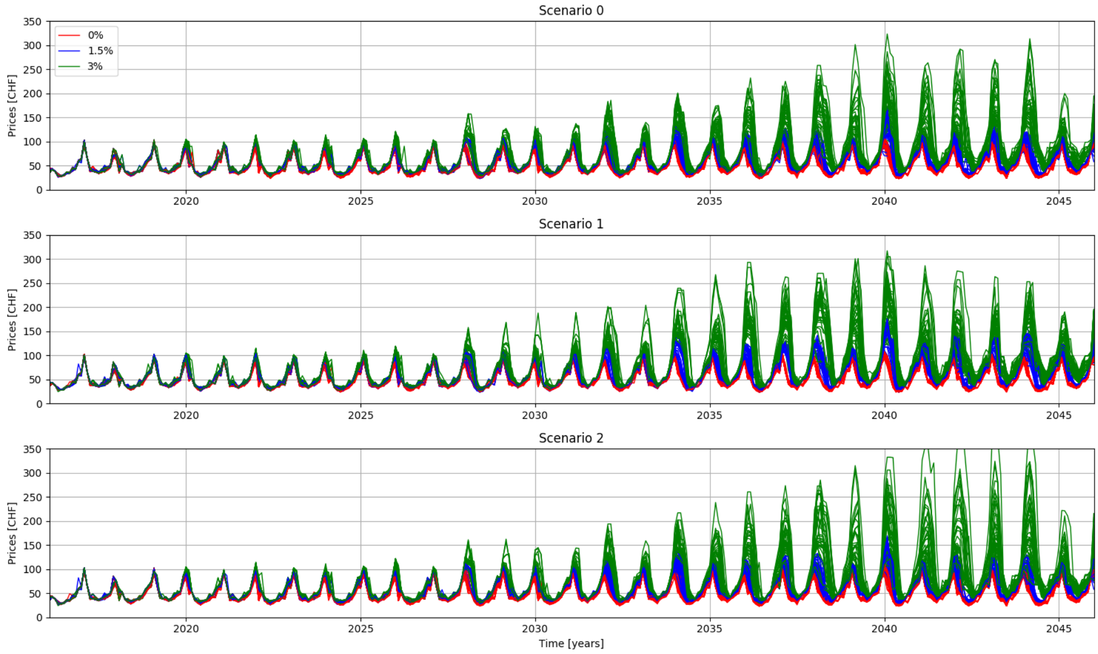

Figure 12.6 – Monthly averaged electricity prices, split according to the three policy process scenarios and for the three demand growth scenarios.

The electricity market model

This difference in policy instrument selection observed throughout the different scenarios and outlined above has not necessarily lead to a change in the trajectory of the electricity market model. In fact, the agents' preferred states are more often than not unsatisfied, highlighting how little impact these instruments had on the model. This is illustrated in Figure 12.6 with plots of the electricity prices for all scenarios as a function of time. In these plots, there is no clear discernible patterns between the different scenarios.

This trend is found across all possible performance indicators that are recorded for the electricity market model. This includes the investments into new assets or the imports and exports of electricity. The only minor change that can be found is a five years delay in the growth of the power supplied by CCGT plants for scenario 1 as can be seen in Figure 12.7. However, even this is only true for the most extreme of cases: the 3% demand growth scenario. For the other scenarios, this difference is not perceptible.

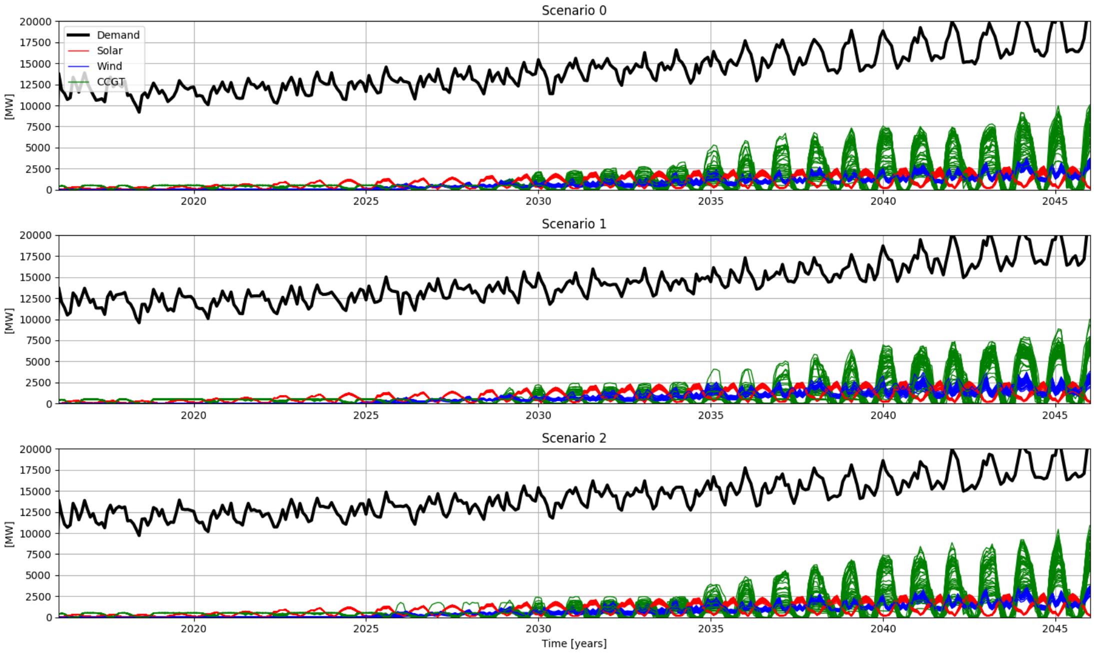

Figure 12.7 – Overall demand with the electricity supplied by solar, wind and CCGT sources for all policy scenarios and for the 3% demand growth scenario.

The coalitions

In this chapter, the +Co model version of the ACF implementation model was selected. Coalitions were present in this model. Unfortunately, and because the coalition creation criteria were too permissive, only one coalition ended up being created. This was not an outcome from policy learning, as was the case for the predation model. This coalition was formed along from the beginning. The main reason for this is that the agents, despite their differences, have very similar beliefs regarding the salient policy core issue. The threshold having been placed at 0.15, all agents fit in that one coalition. This difference in beliefs can be seen in Figure 12.1. It is often as close as 0.1. This outcome is unfortunate considering the initialisation of the model was made according to coalitions present in the real world. This issue is discussed later on.

12.4 Discussion

Overall, the results have shown how little the impact of the policy instruments is on the electricity market model. This discussion is meant to address this and also outlines some of the other limitations experienced, such as the belief mappings, the influencing actions, and the policy core issues.

The models

The results section has shown that using the ACF implementation model is not sufficient to resolve the limitations that were found for the simplest implementation model. They provide only limited additional insights. The model suffers from the same issues, especially on the electricity market model side. This includes limitations related to the policy instruments and their incrementality. It also relates to the fact that the model's scope is limited and restricted. This limits the avenues that can be explored to achieve the goals of the policy agents. As mentioned in the simplest implementation model chapter, this could include innovation, a relaxing of the solar and wind limits, new technologies such as the inclusion of batteries or the addition of demand-side management capabilities.

The knowledge gained from the ACF implementation model is mostly limited to the policy process. From the results of this model, the modeller is able to understand the dynamics between the policy agents better. The results have shown what would be needed to trigger change and that an initial shock would work much better than waiting for policy learning to affect the agents and the policy instrument selection. Even so, using this same model, policy learning and its impact could be further explored. It could be done through the introduction of an agent distribution change as time progresses. This would likely have shown that policy learning is sped up when the agent power dynamic is modified as the pro-ecology coalition gains power both through additional resources and additional agents. This was in part already shown with the second policy process scenario. The key objective would remain to determine whether that would be needed to achieve the goal set by the government.

The mapping to the belief system

The ACF implementation model suffers from other limitations that were exposed in this chapter with the electricity market model, aspects that were not found when using the predation model or the simplest implementation model. They are linked to the formalisation of the policy process itself and are only exposed with specific environment models. One such limitation relates to the issue of mapping the key indicators to the belief system. To transfer the indicators from the Swiss electricity market model to the policy agents' actual beliefs, the indicators have to be mapped to an [0,1] interval. For this, there is a need to define a minimum and a maximum for each indicator. This is where issues arise.

Several key indicators in the model are large or very large in absolute value and can vary over an extensive range. When selecting a minimum and maximum, one needs to consider the values these parameters might take throughout the entire simulation. For the case of the carbon emissions, this needs to take into account a case where CCGT power plants would be installed throughout the entire simulation, potentially increasing the amount of carbon emissions significantly. At the same time, the goal of most policy agents is to see these emissions go to zero. Concretely, the minimum for that specific indicator was set to 0 while the maximum was set to 20 million tons of CO_2. This range obtained through testing and considering worst-case scenarios.

The issue with the mapping is that a change of 0.1 in the beliefs corresponds to a change of 2 million tons for this particular parameter. This is also the amount for which the agents can influence one another on their preferred states each time step. This might seem unrealistic. One way to deal with this, as the mapping will always be required, would be too calibrate the resources' impact concerning the size of the mapping range of each parameter. However, this would mean that influence would become issue-specific and, therefore, would require more data to run. It would also be difficult to obtain data on the policy learning rate of different agents on the issues in their belief system.

Beyond affecting influence and resources, this mapping also needs to be reflected in other choices throughout the model. This was shown with the creation of the coalitions. The mapping was not considered when setting the coalition creation threshold. This led to a situation where the agents always had 'similar' beliefs and formed a single coalition, contrary to the literature. In that case, the preferred beliefs' range throughout the simulation should have been checked before the simulation, and the threshold adjusted to a value of 0.05 or less such that two coalitions would be created early on. This would have prevented the merging of the two coalitions down the line. The danger with this approach is that in the case of an unforeseen event, this very small threshold might lead to the total disbanding of all coalitions if the beliefs were to change too much.

The main problem with mapping is not necessarily the mapping itself. Mapping is necessary for the present approach. The problem is that it is difficult to systematise the approach. Every model will incorporate different issues, and modellers should be aware of this. In some environment models, there will clearly be a minimum and a maximum in which case no issue will be present. In others, that will not be the case, and modellers will have to be advised of the adverse consequences their choices can have on the results.

The impact of agent influencing actions

A second limitation relates to the disproportionate influence agents have in situations of high conflict levels. Agent influence actions are calculated using the difference in beliefs between them. Therefore, when two agents have a high conflict level - a significant difference in their beliefs, the action that is performed, if selected by the agent, will be disproportionately high. The model effectively assumes that agents with high levels of conflict are highly unlikely to talk to each other, but if they do, then their interactions will be highly influential. This is incoherent with the literature that states that agents with vastly different beliefs are unlikely even to talk or listen to one another in the first place. Currently, the only way to prevent such high conflict agent interactions is by not selecting the action in the first place. However, when agents have no other possibilities, they will resort to selecting influence actions on agents with whom they have a high conflict level with no other recourse possible.

There are different ways this could be modified, though none perfect. The first idea would be to simply ban any communication between agents that have very different beliefs. This would be a very drastic move, limiting interactions between agents arbitrarily too low and average levels of conflict. It might be justifiable in some very polarised subsystems, but the literature does not justify it.

A second approach would be to allow agents not to spend all their resources. Once an agent considers that it has done the best it can do and it is satisfied with the state of the system, that agent can keep its resources and stop influencing other agents. This would allow agents not to be forced into influencing agents that have vastly different beliefs to their own, depending on where the threshold for satisfaction is placed. The issue here would be to define, for each agent or the subsystem as a whole, what being satisfied means. This is likely to be a data-intensive process, and it could further complexify the model. However, it would be easy to use for model exploration.

A third and final approach could be to have a constant influence value instead of the influence being based on the belief gap between the agents. This would be a constant, independent of the gap and dependent only on the resources that are being spent by the influencing agent. This could be applied in two ways: either this constant value influence is only used when two agents have a high conflict level - i.e., they have a wide gap in beliefs, or it is used for all influences. The latter would have the inconvenience of prolonged influence for influence with agents with mid to low levels of conflicts.

One question that should be asked, maybe in relation to the literature, is whether there are interactions between high conflict level agents in the first place. If not, then this could simply be removed from the model. If there is, then maybe these only happen under certain specific conditions, like an external event, for example. In this case, such interactions could also be restricted per the literature. The model presented here can allow any of these changes, depending on the case being studied. The problem remains to find what change is appropriate.

The policy core issues

The selection of the policy core issues can also be a limitation or at least a difficulty. This applies to both the simplest implementation and the ACF implementation models. This remark only applies to specific environment models, and it is the case for the Swiss electricity market model. The boundaries of the model are rather small, and the model often does not consider what could be seen as ’the bigger picture’.

Because of these narrow boundaries, the policy core issues are not explicitly present in the model. Instead, they are calculated using other indicators present in the model. In the case of the electricity market model, this was done for the economy and the ecology policy core issues using a linear function of the secondary issues. This is a limitation because these issues do not actually represent what the economy and the ecology are, contrary to the secondary issues that are directly linked to model parameters. The policy core issues are abstract constructs that have nothing to do with actual economy or ecology.

Beyond the abstraction element, the linear function used to calculate the policy core issues can become a problem in itself. As was shown in the electricity market model, it can lead to agents having similar actual beliefs, even if they have a rather strong disagreement in their secondary issues. As a side note, this also links to the mapping limitations and the issue found with the coalition creations. The only way for the policy core issues to diverge in this case would be to have agents with opposed secondary issues, something that is not likely to happen in any policy subsystem. Even in such instances, there would be a possibility for the agents to have similar policy core beliefs.

An alternative approach that could be used to mitigate this limitation is difficult to consider. In effect, it comes down to the creation and/or selection of the environment model. Researchers should prefer models where policy core issues can be easily identified within the environment model itself. This would avoid the present dilemma. If no model can be found, then the linear function approach can be used, but most of the focus of the results’ analysis should then be on the secondary issues and not on the agenda and policy core issue selections. If coalitions are considered, then appropriate dispositions should be taken as mentioned before.

The salient policy core issue

The electricity market model has also brought to the fore another issue relating to the salient policy core issue and the creation of coalitions. In the present application, the choice is made to have one salient policy core issue based on which the agents assemble into coalitions. This is similar to the approach used for the predation model in the previous chapter. However, the case could be made that this salient issue is allowed to change following the agents' belief changes. An issue then arises on how this might be dealt with in the model. A systematic dynamic selection of the coalitions should be possible using criteria that would have to be determined. Clustering algorithms could be used for this as has been demonstrated in the literature, though with the necessary changes and adaptations (Pierce, 2011).

Other limitations

Finally, a number of the limitations that were outlined with the predation model are still present in this implementation. This includes the uselessness of the policy core issue and the agenda. This is because, similarly to the predation model, the policy instruments have no impact on most of the secondary issues or their selection, making all policy instruments eligible to selection regardless of the agenda that has been selected.

12.5 Conclusions

The goal of this chapter was to demonstrate the ACF implementation model with the Swiss electricity market model and to determine whether the resulting hybrid model can provide more insights than the simplest implementation model. The chapter has first shown how using the ACF implementation model does not lead to added complexity when it comes to the integration of the hybrid model. It has also shown that only little additional insights are gained from the use of the ACF. The main reason behind these limited results is the fact that policy learning, and the presence of coalitions, only reinforces the trends observed in the simplest implementation version of the model. Even through the introduction of external events, the results have shown that, barring expanding the electricity market model to have more possibilities for the agents, the results obtained are unlikely to change in any significant way.

The additional insights that were gained were mostly related to the ACF model. These were insights on what can be gained from the use of the model along with the limitations that come with the model. The model can provide a better understanding on the dynamics within the policy process, may that be on how policy learning is evolving or on why and when coalitions form. The model also has a number of limitations that need to be taken into account when analysing the results. These were detailed in the discussion and were mostly operationalisation-related limitations. Dealing with these limitation, and validating potential solutions, could also help advance the literature.

This chapter marks the last demonstration of the ACF policy process model for this dissertation. As outlined throughout the different chapters of this third part, this only shows some of the potential of this approach. It calls for further studies to outline how the policy process can be used in modelling and simulation, to validate the policy process models used here and to use these models for theory exploration.

Conclusions Part IV

13 Discussions

Throughout the dissertation, chapter-specific discussions have focused on discussing aspects directly related to the respective chapters and the results obtained. This chapter looks at the bigger picture, discussing the general approach used to construct the policy process models, the methods that have been followed, including the common language, and then outlining the results obtained. This discussion includes limitations that were observed and that affect both the simplest and ACF implementation models.

In this chapter, I also propose alternative approaches, when possible, and I outline what the trade-offs are in considered these approaches. The goal, by presenting these alternatives, is to highlight the lack of a silver bullet solution for the modelling of the policy process. Instead, combinations of assumptions should be considered, assumptions that might often be only well suited for one or a set of cases.

A n-step policy process

One of the initial assumptions for the construction of any model using the common language relates to the number of steps that the policy process should follow. It was decided early on that for the simplest implementation model a one-step approach would be sufficient. For the ACF implementation model, a two-step approach was preferred. I will show here that though crucial, this assumption was undermined by subsequent assumptions and behaviours, and can also be undermined by the environment model. Considering this, I will also show that a one-step approach, containing only a policy formulation step, would have been appropriate for both models.

Before I expand on this, I would like to highlight the adequate number of steps considered in the models built for this thesis. Though the policy processes considered only had one or two steps, the real number of steps is much closer to two- or three-step policy processes when considering the entire hybrid model. This follows the theory of the policy cycle (Jann and Wegrich, 2007). By considering the environment model, the policy implementation and evaluation steps are included in the overall hybrid simulation, though the agents are unable to affect these steps. Policy implementation happens when the policy is implemented in the environment model. The policy evaluation occurs when the indicators are passed along to the policy process model to inform the actual beliefs of the policy agents.

Though the agents cannot influence the implementation of policies, the model does not prevent a future modeller from adding a module that could bring policy implementation into the decision-making process of policy agents. This can be done by explicitly adding an implementation step after a policy has been selected. Then, depending on the goal of the research and the data obtained for the case, the modeller could make sure that implemented policies are translated partially or inadequately into the environment model, in the same way as it might happen in real life where the actions, or inactions, of certain actors, will impact how thoroughly and effectively a new policy is implemented. Depending on the complexity considered, there is also the possibility to add an entire model for this step, making it the focus of the policy process and simulating the decision made by the actors in charge of policy implementation. The same goes for any other step that a researcher might want to study in the policy process, and that is not currently considered.

The one-step approach used for the simplest implementation model can be justified along different lines depending on the theory that is being simulated but also based on the results that have been analysed so far. On the one hand, one could assume that the agenda is a higher-level issue and that the selection of the agenda is made with the selection of the policy subsystem. This would effectively exclude agenda setting from the model entirely. For example, actors in parliament might select to either discuss healthcare or security. This would push the decision of the agenda out of the scope of the simulation boundaries as those two topics will often belong to two different policy subsystems.

Additionally, most environment model would likely be focused on only one of these widely different topics. Having a one-step model would make perfect sense in such an instance. It would leave the selection of the policy instruments as the lone decision that the agents have to perform throughout the simulated policy process. Coincidentally, the ACF literature does not mention any agenda or the agenda setting within its framework, suggesting that this one-step approach might be more representative of the framework than its two-step counterpart.

Another reason why a one-step approach can be sufficient is related to the structure of the modelling implementation. The results have shown that the agenda has a minimal, if any, impact on the policy instruments implemented. In the present approach, the agenda is used to limit the set of secondary issues that the agents consider in the policy formulation step. This limitation would be expected to trickle down to a restriction in the selection of the policy instruments. However, in practice, this does not happen. This is because the agents can select policy instruments if they have an impact on the secondary issues chosen by the agents, and all policy instruments always have an impact on the secondary issues. Even if that impact is negligible, it is never exactly zero.

To understand why that is, it is necessary to go over how the policy instrument impacts are calculated. They are obtained through a separate simulation where each policy instrument is implemented into a copy of the environment model, and the model is simulated for a given period. The impact on each of the secondary issue is calculated based on the changes in the indicators from the time before the instrument is implemented to the end of this separate simulation. The impacts are then used in the broader model to inform the decision-making process of the policy agents. Because all indicators always change, even if for a small amount, this results in all policy instruments always having an impact on all secondary issues related to the policy core issue on the agenda. This happens regardless of the agenda and the restrictions on the secondary issues chosen. In effect, this renders the agenda setting step useless beyond the policy learning behaviours that might partially influence the preferences of the agents' secondary issues.

Different approaches can be used to remedy this problem and give more weight to the agenda-setting step. One would be, as mentioned above, to only consider a one-step approach without changing the other parts of the model. In this way, the agenda would be considered outside the boundaries of the policy process. A one-step approach would also remove the limitation related to the mapping of policy core issues in environment models with no explicit policy core issues - a limitation mentioned in Section 12.4). However, this approach would also introduce new limitations. This includes difficulties, or the impossibility, of choosing coalitions, the removal of policy learning at the policy core level influencing the agent preferences calculations for secondary issues, and it could lead to the - almost - trivialisation of the policy process model. Some insights would be lost along with these limitations.

The second option would keep the two-step approach and better represent reality through the model. This can be done by rethinking the way the policy instruments' impacts are communicated to the agents. At the moment, this is fully transparent, meaning that all agents have full, unbiased information about what the impacts are. In real life, these impacts would not necessarily be known by all agents, and different agents would draw different conclusions from different studies assessing these impacts, focusing on how a few specific issues are affected by the different instruments.

With this in mind, the policy instrument impacts communicated to the agents could be modified. One approach could see the impacts communicated to the agents only if they are above a certain threshold. This would help weed out instruments that have a limited impact on important issues. In the case of the Swiss electricity market, for example, a policy instrument like the hurdle rate might have only a limited impact on imported electricity prices. In that case, its impact would not be considered, such that if an agent selects the electricity price as its preferred issue, the hurdle rate policy instrument would not be proposed in the agent's policy options. Other approaches, maybe even bridging to the ones mentioned in the +PI model discussion (see Section 12.4), could also be used here.

In summary, there are different ways to justify the number of steps that a model of the policy process should contain. In general, I would advise keeping a two-step approach for more complex studies while considering steps to help limit the set of policy instruments considered in the policy instrument selection, and in this way making agenda setting more critical. As mentioned, the simplest way to do this would be by including a threshold. Maintaining this two-step approach guarantees a certain amount of insights that would be otherwise impossible to get. It also allows for more emergent behaviours.

The policy instruments incrementalism

The policy instruments' incrementalism is an aspect that has been shown to impact the results throughout the different implementations presented in this dissertation. They limit the amount of change that can happen in the model at any one time. For the Swiss electricity market model, this had a real impact as with only incremental policy instruments it was shown that it is impossible to reach the goals set by the government - though it was not the sole issue responsible for this outcome. Despite the limitations brought by such incrementalism, the incremental approach tracks with the literature. As it is mentioned in Lindblom (1959), policies are often selected when they "differ in relatively small degree from policies presently in effect."

However, there might be some merit in investigating what would happen if policy instruments were not as incremental as they currently are. It is technically possible to test more variations of the same instruments, but a trade-off quickly appears. First, there is a computational element. Having to assess the impact of a hundred policy instruments for each time step would require much computational power and would significantly slow down the model. Note that currently, the policy instruments' impact assessment is the main bottleneck of the model despite using parallel processing methods. Second, having many policy instruments will inevitably splinter the vote of the policymakers. There will be a much higher probability for the status quo to be maintained because of a lack of a majority. To deal with this, a new selection algorithm would be needed to make sure that agents can agree on policy instruments when needed, even it ten agents have to choose from two hundred policies. Other approaches have also been proposed in previous chapters' discussions, such as including a module that would have agents build up policy packages so that they can reach a majority.

The merits of considering less incremental policy instruments might be made most apparent when investigating transitions, and more specifically, the energy transition. Policy instruments that are incremental are often insufficient, as was shown with the results from the Swiss electricity market model. By considering policy instruments with more significant changes, it might be possible to investigate what would be needed to achieve a transition in a reasonable time, whether that is even possible or whether the policies proposed are never sufficient alone.

The model validation

The validation of the models presented within this dissertation has yet to be discussed at length. The validation of a model consists of making sure that what has been modelled matches with what was supposed to be modelled. Validation can be done using many different approaches. Validation approaches may differ from model to model, but here I only consider the validation of agent-based models as this is the modelling paradigm used for the policy process model. The environment models used - the predation, electricity market and flood safety models - were validated independently, and their validation is considered out of the scope. The different validation techniques considered are the following: face validity, Turing tests, historical validity, even validity, validity of simulation output, solution space exploration and participatory approaches (Barreteau et al., 2013; Nuno et al., 2017; Pahl-Wostl, 2002).

In the present case, the validation can have two goals. The first is to validate the fact that the theories are used appropriately, according to the literature. The second is to see whether the overall hybrid model, resulting from the integration of the policy process with environmental models, produces realistic results. These two goals can be fulfilled by using two different approaches. The technique used to validate the policy process model is a participatory approach, and to be more precise: expert validation. To validate the overall hybrid model and the results obtained, another participatory approach is considered: serious gaming (Susi et al., 2007; Michael and Chen, 2005). In circumstances where the environment model has been the focus of the work, and therefore data has been gathered for both the policy process and the environment model, a historical validation would most likely be more appropriate, though I will argue here that a serious game can be just as useful.

To validate the policy process, expert validation is used. This consists of asking key academics who are prominent researchers in the field, the original makers of the theories used if possible, whether they agree with the way the model has been conceptualised. This was done early on in the construction of the policy process model. Four key researchers who have written a large number of papers to advance the theories reviewed the original model. Furthermore, the feedback was obtained at several conferences where a more recent version of the models where presented.

For the most part, the original model was well received. The researchers agreed on the overall approach while noting some more delicate points. Most of these were related to the requirements for simulation. This includes, for example, the need for having a linear policy process, an assumption frowned upon when considering qualitative studies of the theories used. In some instances, there was a rejection of the approach taken. This was the case for an earlier version of the model, though not the original version, that followed a combined holistic approach. This approach assumed that the ACF and the MSF could be combined, following research presented with the five streams framework (Howlett et al., 2016a,b). Most of the researchers asked about this approach mentioned that this five streams approach was not well received within the field and that I should stay clear of it. This ultimately led to the approach presented in the dissertation where different theories are simulated separately, though all built using the common language which was developed after this feedback. Finally, expert validation can be considered a continuous process of validation for as long as the model is changed and refined. Note for full transparency that the final models were not revalidated because of time constraints.

The overall hybrid model can be validated using another participatory approach: serious gaming (Pahl-Wostl, 2002). In this context, serious gaming consists of taking the simulation into a workshop setting and replacing the policy process model by the actors that are present at the workshop. The participants can decide on policies, in a controlled setting, that will influence the environment model, which remains a computational simulation. The environment model is simulated for a certain period, and the necessary information is provided to the participants. They can then decide on further policies depending on their respective goals. In effect, the policy process model is brought to life in a small room where the serious game facilitator can record and study the behaviour of the participants.

Though not perfect, such an approach can be used to validate many aspects of the model. First, some of the assumptions of the model itself can be validated. If the participants are real-life actors of the policy subsystem, then their decision-making process can be studied live during the workshop. Depending on the model to be validated, the constraints placed on the participants would be different. If the simplest implementation model is to be validated, for example, participants could be prevented from discussing with one another. After recording their initial beliefs, there would just be a need to observe the decisions that they take.

In the case of the ACF implementation model, depending on the complexity considered, the participants could be allowed to perform actions. They could be allowed to talk to one another in an attempt to influence one another. This could also be influenced by either providing all actors with full knowledge or with partial knowledge. Many possibilities are present. It would then be up to the serious game facilitator to study all these interactions and see whether they fit with what has been included in the model. Of course, this would also have to be reconciled with the formal theories considering that a workshop setting might not be the best place to observe genuine negotiations, for example.

Serious gaming can also be used to verify results when it comes to simple abstract models such as the predation model. A set of preferred states can be given to the participants. They then have to select policies such that these goals are met. They decide based on where the different animal populations are at the time. It was planned to perform such a workshop with participants from the university to validate the simplest implementation model at least. However, due to unforeseen last-minute circumstances, it became impossible to set up the workshop. This is work that could be performed at a later stage and could be used later on.

Beyond validation, serious games can also be used for two other purposes: data collection and dissemination. The former is discussed later on. Serious games can also be used to inform participants. This can be done through specific scenarios that the participants are run through. By following scenarios, the participants can be made to experience particular situations on which the facilitator would like to educate them. So, for example, a serious game workshop could be a way to convey the results of the Swiss electricity market model practically. By making the participants react to the Swiss electricity market model with a set of restricted policy instruments, participants could be made aware that incremental policies are insufficient and that there is a need for out-of-the-box thinking and innovation if Switzerland is to meet its energy transition goals. They would not only be provided with this information but also the reasoning why that these policy instruments are insufficient. This could be particularly valuable if the participants are policymakers.

The data collection

The data collection aspect of the model has also been mostly avoided throughout the entire dissertation. The main reason behind this omission is the fact that, because of time constraints, no data collection was performed and the simulations were mostly run with symbolic data or partially representative data. Although no data was collected, there were indeed mentions of the fact that the models used were data intensive. I detail this further here.

To simulate a specific case study, data is an integral part of the process. There is a need for information on who the actors are, what their preferred states are, and if possible, what their understanding of the world is - to set up the causal beliefs. This is the least amount of data required to run the simplest implementation model. Simulating the ACF implementation model, depending on the complexity level, would require even more data. This includes an understanding of the power dynamics to define the resources, an understanding of the coalitions and the issues around which the actors come together. For further complexity, there might also be a need for understanding the actors' actual beliefs or the way they inform themselves about the world. Overall, this requires a large amount of data to be collected, data that might not always be easy to obtain from actors that have little time to provide it.

The serious games approach was proposed previously as a method to collect this data. This can be done in the following way. Key actors can be enrolled in a series of two workshops, spaced by a few months. The first serious gaming workshop is one where actors play their regular roles. By doing that, the facilitator can study their behaviours, but also gather their beliefs and observe the dynamics in the room. This is data that can go on to inform the model later on. The second session is used to inform policymakers as a dissemination exercise, as was outlined in the previous section. In this way, the modeller gains the data s/he needs from the first session to simulate the model, draw conclusions, and then disseminate those conclusions through the second workshop. This two-pronged approach also makes it more likely for policymakers and actors, in general, to join as their presence will also allow them to learn.

Other avenues to collect data exist. The preferred states of the agents can be obtained through surveys using Likert scales. Previous qualitative researches have obtained data on the beliefs through parliamentary transcripts or on on-the-record documents as well (Pierce, 2011). Such data collection efforts can be time-consuming but can provide an insight into the beliefs of an actor or organisation and understand what strategy they are using. There is also the option of using data obtained from previous studies. At the moment, this would include mostly qualitative studies. There are some caveats in this case as the model, and the way the policy process theories have been approached here might differ from the operationalisation that qualitative studies might have used in the past. For example, in Pierce (2011), the beliefs of the actors are collected. However, these beliefs are not defined in the same way as they are in the present models. As defined, the beliefs can only take three discrete values, depending on whether an agent has the belief, does not have the belief, or it is unknown. Any data collected for this example study would not be usable in the present model where continuous values are needed.

Finally, when considering a full in-depth study of a specific case, the data collection required for modelling might not be as cumbersome as expected. In any qualitative case study, there is a need for a lot of data obtained through a variety of methods, including surveys and interviews. If there is prior knowledge that the model is to be used at a later stage of the research, then, while collecting data for the qualitative study, data can also be collected for the initialisation of the policy process model. This requires the addition of a few survey questions and potentially specific interview questions. However, such an effort would significantly enhance the study as a whole, allowing for more in-depth insights. As part of a whole study, the data collection process would be limited to a minimum.

Shifting scenario limitations

In the introduction, I have outlined one of the main limitations of the current approaches being used for modelling and simulation: the way current models deal with deep uncertainty by using exogenous policy implementation. By this, I meant that current approaches introduce policies as input to socio-technical models for policy analysis purposes. These are often introduced to the simulation through scenarios that are designed by the modeller and that do not take into account what is happening in the simulation at the time a policy is implemented. This can limit the results obtained from such simulations as the policies might be inadequate or ineffective. Furthermore, it does not take into account whether these policies would have been able to go through the policy process or not.

The approach taken here, through endogenisation, removes these limitations. Scenarios for the policies are not required anymore as it is the agents that select the policies based on their beliefs and the states of the environment model. However, this is not a silver bullet that solves all issues related to deep uncertainty and modelling. Instead, different limitations are brought in with the endogenisation of the policy process.

Scenarios are now required to simulate the policy process itself. These can be scenarios for the evolution of the agents' resources or their preferred states, as was shown in the different scenarios explored within this dissertation. It can also include a change in the agents' role distribution and any changes linked to the numerous parameters present in the model. In essence, the burden of scenario creation is moved from policy scenarios to policy agent or policy process scenarios. This shift is unavoidable for simulations. Models that do not require scenarios do not exist. There are always external events that can trigger a multitude of behaviours. However, this shift can be used to the modellers' advantage to gain more insight into the agents, or policy pathways or other elements that might not have been thought of in this dissertation. This approach opens the way to numerous new applications.

The art of modelling

Modelling is often referred to as being more an art than a skill because building models requires a lot of decisions and assumptions to be made from the part of the modeller. One of the most significant decisions that have to be made is related to the setting of the boundaries and the definition of the scope. A lot of that is dependent on the research question that is being answered but is often not limited to it. It is up to the modeller to decide how much should be included in the model.

This fact is just as crucial with the hybrid modelling approach presented here. Decisions have to be taken at three different points in time regarding the boundaries: for the creation of the environment model, for the creation of the policy process model and the hybrid model resulting of the two previous models. The creation of the environment model, if explicitly built to be then combined with the policy process model - not as it was done in the implementation models presented here, requires special attention on the potential feedback effects and behaviours that might be desired or expected. As was shown with the electricity market model, a model that falls short in scope might not provide as the insights expected compared to a model that would have had wider boundaries.

For the policy process model, this decision making from the part of the modeller is most important at the beginning of the model building process, after the modeller has selected the theories that s/he wants to emulate. The modeller needs to decide which concepts of which theories will need to be explicitly included, which concepts will be emergent behaviours and which should not be included or would not fit with the essence of the model. One such example might be the strategies that are followed by the coalitions. In the present models, these were very limited and highly dependent on what the belief system allowed organically without the need for extra complexity. Other approaches might be preferred in different contexts.

The qualitative literature goes further. Sometimes coalitions go to extreme length to advance their interests, concocting sophisticated strategies like occupying their opponents on the negotiation of a policy such that they have fewer resources on another (Winkel and Sotirov, 2011). One can decide to model such aspects, but that would often be falling in a trap where unnecessary complexity is included, complexity that would hardly be generalisable to other cases. In the present case, the goal was to build models that can approximate a majority of cases and with the lowest amount of complexity. However, there are cases where there might be a need to research specifically coalitions strategies, in which case the strategies used by the coalitions in Winkel and Sotirov (2011) might be necessary.

For the hybrid model, the creation of the belief system and the design of the policy instruments are crucial in establishing the scope of the simulation and the expected results. They should be based both on what is the focus of the study, but also consider the availability of parameters within the environmental model and should make sure that they do not consider elements that are not present in the environmental model, even implicitly.

To conclude this point, I would like to highlight that, though this dissertation has attempted time and again to provide a set of steps to simplify the construction of models of the policy process, no modeller will avoid having to take a certain amount of decisions, especially regarding the model boundaries. Every case is unique, and every case will require a different focus. Anyone that would like to use this work and build a model of the policy process should be aware of this fact.

14 Conclusions

In the introduction, I began by outlining the two main challenges that this thesis had to tackle. The first one related to the current inability to perform theory exploration with the policy process theories using current methods, along with the current lack of explanatory power of these same theories. The second one related to the methodological limitations encountered when using modelling and simulation to deal with deep uncertainty for policy analysis. A solution to deal with both of these very different challenges was proposed: the modelling and simulation of the policy process.

To demonstrate this solution, several steps were considered. This first consisted of developing a common language for models of the policy process. Using four building blocks, the common language established a set of requirements for the building of policy process models. These building blocks were based entirely on what can be found in the theories of the policy process. The agent-based modelling paradigm was chosen because of its ability to deal with agency and process, aspects that are pillars of the policy process. This selection was completed with the presentation of a hybrid modelling approach needed to simulate the policy process within a specific context. For this, it was proposed to couple a model of the policy process with environment models, often socio-technical models.

Moving beyond the theory, the second step demonstrated how the common language could be used to build a model of the policy process. This was done with the simplest implementation model. The goal of this implementation was to construct the simplest model of the policy process possible using the four building blocks of the common language. A one-step policy process was conceptualised, along with the role of the policymaker as sole agent in the model. This simplest implementation policy process model was then simulated with three different environment models, each satisfying a different purpose. First, a predation model was used to explore the possibilities of this novel model. Second, the model was simulated alongside a model of the Swiss electricity market model. The goal of that simulation was to show how the policy process model could be used to better understand and get new insights into a real case study. Finally, the policy process model was implemented with a system dynamics flood safety model, the goal being to demonstrate that the hybrid modelling approach works across different modelling paradigms, making the model highly versatile and reusable.

The third step was to push the boundaries further and to build a model that can emulate a specific theory. In this case, it was done for the advocacy coalition framework. As a result of the ACF's complexity, this was approached through four models with increasing levels of complexity. It included first a model that introduced policy learning through agent interactions (the +PL model). It was followed by a model introducing coalitions (the +Co model), then a model considering the impact of partial knowledge (the +PK model) and finally a model exploring the effects of imperfect information (the +PI model).

The ACF implementation model was also tested with the predation model with exploration in mind. The predation model was simulated with every variation of this model to demonstrate the model's potential but also to outline the limits in the theories' assumptions and the need for further research. Finally, the model was simulated with the electricity market model once again. The +Co model was used to understand whether the limitations found within the simplest implementation model could be overcome through an increase in the complexity on the part of the policy process model. This was not the case as the limitations of the electricity market model continued to be a significant hurdle to gain additional insights.

In the final part of the dissertation, I discussed the approach and the results that were obtained throughout this thesis. It included a number of the limitations that were observed through the use of the models created from the common language. This also included a discussion on the model's validation and the crucial data collection. In this conclusion, I will now explicitly answer the research questions that were formulated in the introduction. To complete the dissertation, the following chapter addresses future work. The present dissertation opens the door to a large number of potential studies focused on the policy process theories, modelling and simulation, and a range of other disciplines.

The principal research question of this dissertation was given in the introduction:

To what extent can simulating policymaking be used to explore the policy process theories and help study policy impacts on socio-technical systems?

Three sub-questions accompanied this research question. I intend to first answer these explicitly below before answering the main research question.

Can a systematic approach for the creation of policy process models be formulated?

Yes. This dissertation has shown that it is possible to create a systematic approach. This was made possible as a result of the development of the common language. Using this language, it is possible to create a broad set of models depending on the goals of the researchers. This common language takes into account what can be considered as the minimum building blocks needed to emulate the policy process according to the policy process theories. It leaves open the possibility of adding an almost infinite amount of complexity. Finally, the common language is an approach that was also developed for the ease of use of modellers that might not be familiar with the policy process theories. The hope is that it can be used by researchers further out from the field of policy sciences, therefore broadening the use of these theories.

In what way can a policy process model be simulated within a socio-technical system?

The work presented in this dissertation has followed the common language approach. Considering the time it can take to make either a socio-technical model or a policy process model, the goal was to provide a methodology that allows for the reuse of already existing socio-technical models. This way, the two models were simulated as a hybrid model where the models are simulated sequentially. The models communicate through a minimal set of parameters that are inputs and outputs to each other, making it easier to use a variety of socio-technical models. These parameters include indicators as outputs for the socio-technical model and inputs for the policy process model and policy instruments as outputs of the policy process and inputs for the socio-technical model.

What insight can be gained from the simulation of the policy process within a socio-technical system?

The insights can be categorised in three separate sets: the policy process, the socio-technical system and the feedback effects. Throughout the different implementations with the three different environment models considered, insights were obtained from the results.

Concerning the policy process, the insights come in two forms: on the policy process itself on a case per case basis, and the policy process theories. Each simulation provides information on the trajectory that the policy process and the policy agents might follow. The results of the simulations can also be used to get insights on the assumptions used to build the model and on their validity. On the other hand, through the construction of the models themselves, insights can be gained on the assumptions that are used to build the model, the vast majority of which are coming from the theories or are derivatives of concepts present in the theories. For the socio-technical systems, a lot of the insights are insights that a simple simulation - one without the policy process - would be able to provide. Such insights can also be obtained in a study before the use of the policy process.

The bulk of the insights come from the feedback effects between the two models. As was shown throughout this dissertation, most insights are obtained by looking at the reaction of the socio-technical system to the policy instruments implemented. The opposite is true, as well. Insights can be gathered by looking at the impact that changes in the socio-technical system have on the beliefs of the policy agents, and how this impacts their decision making process and ultimately the policy instruments that they implement.

Finally, one should also not underestimate the insights that can be obtained from the hybridisation step where the beliefs of the agents have to be selected based on the environment model scope and boundaries, along with the policy instruments or packages. This construction step of the hybrid model forces a reflection on essential aspects of the model that often come back in the analysis of the results and the subsequent discussions. This was the case for the Swiss electricity market model and the flood safety model for example.

To conclude this chapter, this dissertation has shown that the simulation of the policy-making process can help explore the policy process theories while providing in-depth insights into socio-technical systems with which they are simulated. The approach provided in this dissertation has shown that several different policy process theories can be simulated and has provided a stepping stone for other researchers to simulate further theories. It also allows for policy sciences' researchers to explore the theories by experimenting with a variety of assumptions for the different concepts that are present in the theories. On the other hand, the model can provide deep insights on the feedback effects between the policy process and the socio-technical system, these being limited only by the scope and boundaries of the socio-technical systems considered. Finally, the present approach has also provided a new method to deal with deep uncertainty in models through the endogenisation of the policy process.

15 Further research

Simulating the policy process is a novel approach. The work that has been presented in this dissertation is only just one step forward and only touches the surface of what can be done and achieved with this novel approach. As was discussed throughout this dissertation, starting with the motivations in the introduction, there is much potential with the use of modelling and simulation of the policy process theories. Based on the work presented here, I mention several avenues that could be explored to continue the present research, either by continuing along what was demonstrated in this report or beyond.

Confirming the current approach

This dissertation has presented many experiments to test the two models developed. However, these experiments are limited due to the scope of the thesis and because of time constraints. As was outlined in the discussions, there are still many experiments that can be performed on the models to understand better them and better test them. Such experiments can also be used to assess whether some of the details not looked at yet fit with the theories more widely.

Beyond this, several assumptions were challenged throughout this dissertation. For some of these, alternatives were proposed, either to expand on the behaviour observed or, because the theories are permissive enough, to use different approaches to solving the same problem. I would propose that these different assumptions be tested against the ones that were used in this dissertation and for which the results were obtained. Some valuable information can be gained by testing different assumptions, may it be for example to understand how it might affect the behaviours of the model or so that it can be used to advance the theories of the policy process further.

Testing more theories

This dissertation has focused on the simplest implementation model, which was not meant to emulate any theory in particular beyond considering concepts taken from several theories. It also included the ACF implementation model that was meant to emulate the ACF, using concepts mostly from the ACF. In effect, this means that only one theory was tested in this dissertation.

This can be pushed further. This was one of the goals for the development of the common language. Another logical theory that could be simulated would be the multiple streams framework (MSF). This framework has been mentioned throughout the dissertation, and some of its components appear in the two models developed. The main requirements to emulate the MSF would be to introduce the three streams. At the moment, in the ACF implementation model, the policy entrepreneurs are already present, and one can assume that the agenda can be conceptualised as the window of opportunity. When it comes to the three streams, the problem stream can be associated with the belief system and the politics stream to the policy agents at large. Only the policy stream is missing. One approach could see the policy streams introduced within a model of the MSF by considering a structure mirroring the belief hierarchy and regrouping policies. It could then be referred to as the policy hierarchy. This would be partly inspired by the work done on the five streams framework in Howlett et al. (2016a).

Going beyond the ACF and the MSF, other theories could be simulated using the common language approach as well. This might include the diffusion theory. Such a model could be used to explore multiple policy subsystems' interactions (Lubell et al., 2002). The feedback theory is another theory that could be considered. This would require the introduction of specific feedback effects with the introduction of every policy instruments. The common language allows for the implementation of each of these theories and more, well beyond the ones presented in this dissertation.

Delving into the details

Most of the implementation has so far stayed shy of going into the details of the theories. Future work could explore such aspects of the models by focusing on specific concepts or parts of the policy process or of certain theories. This would be dependent on the research questions proposed.

Such studies should gather all the information present in the literature on the part of the process that is of interest to the researcher. This would lead to modellers to having to make many assumptions. These could then be tested with the policy process model and in coordination with qualitative studies. All this would also have to keep in mind that the findings might be case-specific and not generalisable. An example of such a study would be to further detail the causal beliefs. One study has looked at them, Weible et al. (2004), but not in a way where these causal beliefs could be operationalised in a model. Having a study that considers modelling from the start could be beneficial to both the modelling and the policy sciences literature.

This could be extended to other concepts and parameters in the belief hierarchy. For the ACF implementation, there are numerous examples of details that need to be investigated, from the strategies that the coalitions use for their decision making process to the way coalitions come together and strategise. More details can also be considered for external parties better to study the propagation of false or incomplete information. That part of the model was under-developed and could be highly relevant in current times where the policy-making process can be influenced heavily by distorted information.

Expanding the model use

This dissertation has seen the policy process model being implemented with a Swiss electricity market model and a flood safety model as environment models. This is only a minimal set of examples. I would encourage the use of more varied environment models for further research. There are many benefits to doing so, beyond only studying these specific models. The application of the policy process model to more environment models would further demonstrate the versatility of this approach and its usefulness. It would allow researchers to systematise further aspects such as the hybridisation of the models, something that has been singled out as a limitation for the electricity market model. Beyond this, connecting the policy process model to different environment models would also allow to more thoroughly test the policy process model. As was shown with the predation and electricity market models, different issues will be brought up with different environment models. By having more examples, the model will be continuously improved and advanced, advancing the literature along with it.

Exploring alternatives

One other avenue to explore would be for modellers to go back to the common language and produce new models of the policy process, disregarding entirely the ones considered here. The assumptions in all of the implementations presented in this dissertation were that the models should use the belief system from the ACF and a one or two-step approaches of the policy cycle. Other researchers might find that these assumptions are incorrect and could propose alternatives, henceforth enriching the literature and testing new ideas. This could further the thinking around the policy process theories and hopefully contribute to them at the same time.

Simulating socio-technical transitions

Finally, the policy process could also be used in the field of transitions. This has been briefly mentioned in a previous chapter. Transitions are an important and current topic. Modelling and simulation is currently a nascent methodology in the field, following some qualitative theories and frameworks that have been developed (Hansen et al., 2019; Papachristos, 2014; Holtz et al., 2015). However, these simulations are currently limited, and often rely on scenarios for essential aspects such as the simulation of the socio-technical landscape and to deal with deep uncertainty (Papachristos and Adamides, 2016; Papachristos, 2017; Köhler et al., 2018). This relates directly to the justifications of this dissertation. There is an argument to be made for using the policy process model as a representation of a part of the multi-level perspective landscape. This way, it might be possible to simulate transitions using the same model as the one presented here. This could prove valuable to continue the advances in the field of modelling and simulation for the transition literature.

A The Swiss electricity market model

The model presented in this appendix is a transposition of a model created by Paul van Baal and Reinier Verhoog and first created by Paul van Baal for his master thesis (van Baal, 2016). The model has been further improved to look into the effect of a strategic reserve using a hybrid system dynamics - agent based model version (van Baal, 2019). In this report, the model, which was a system dynamics model and then a hybrid model, is turned into an agent based model. This appendix is the ODD presentation of that model (Grimm et al., 2010). All equations and additional details can be found on the GitHub where the model is saved[1].

A.1 Purpose of the model

The purpose of this model is to simulate the Swiss electricity market system. This includes the spot market, international trade with neighbouring countries, and investments.

A.2 Entities, state variables and scales

There are four types of agents within the model: the market operator, the firms (or investors), the supply agents and the demand agents. The market operator is the agent that is in charge of the spot market, making sure everything is going well. The firms are the agents that own the power plants and other assets present in the model. They are in control of the power plants. The demand agents are the agents that set the electricity demand in the model. This includes the inflexible demand (which is based on a historical scenario), and the demand created by hydro-pumping assets

[1] https://github.com/kleinrap/SwissElectricityMarket

and international trading.

The firms are characterised by the following attributes: assets owned, electricity supplied, planned assets, retired assets and constructed assets. The supply agents can either be the power plants (assets), the long term contracts (LTC) with France, or the net transfer capacities (NTC) from the different border countries. Each has a different set of attributes. The assets are characterised by the following attributes: owner, technology type, installed capacity, age, lifespan, capital costs, annual fixed costs, variable costs and utilisation factor. Additionally, depending on the technology type, some power plants have more parameters. For example, the thermal power plants, the nuclear power plants and the waste power plants all have a fuel cost. The thermal power plants also have an emission attribute. The nuclear power plants have attributes related to their maintenance requirements: maintenance month and maintenance time. Waste, hydro and hydro-pumping assets have attributes related to their reservoirs: reservoir level and maximum reservoir level. Hydro-pumping assets have an efficiency attribute related to their pumping efficiency.

The technology types are limited to: solar, wind, hydro power, hydro-pumping power, run of river, thermal, nuclear and waste. The firms can only invest in solar, wind and thermal technologies as it is assumed that other technologies are already maxed out in Switzerland or they cannot be used to produce significantly more electricity.

A.3 Process overview and scheduling

The model runs along two different scales, highlighting two parts of the model. The first part is the spot market, running on an hourly basis. It consists of all of the actions related to the spot market including all of the inputs, the calculation of the demand, the calculation of the spot price, the distribution of the revenues and electricity when the equilibrium is found, and the update of the NPV for all of the agents (that is later used for investments).

The second part is related to the investments that the firms can perform. This happens monthly. These are actions that are related to the firms. They decide whether to invest in new assets. They also decide whether they should reinvest in their current assets by extending their lifetimes or shuttering them temporarily or definitively. Then there are additional measures that include the end of life actions that occurs when an asset has reached its lifetime. It also includes scenario based events such as the closing of nuclear power plants according to a politically determined timeline.

A.4 Design concepts

Basic principles

The model is, in essence, a simple supply and demand model where electricity is demanded and supplied. The main added element is that instead of resolving this supply and demand every week or year as it is mostly done in current models, it is done on an hourly basis. The model also includes the possibility to invest into new capacity to satisfy a growing demand and deal with timed asset closures.

Emergence

The main outputs of the model relate to the energy mix that is needed to meet the Swiss electricity demand. The investments, their type and amount are also of interest for the purpose of the study and should emerge from the need to supply electricity.

Adaptation

There is no real adaptation programmed in this model beyond agents deciding on whether to discontinue their current assets and whether to invest in current or new ones.

Objectives

The objectives for the market operator is that there be a balanced spot market, demand should be met by supply. The objective for the firms is to earn as much money as possible. The objective of the supply agents is to supply as much energy as possible. The demand agents would like to see their demand met.

Prediction

The firm agents have to use prediction to decide whether to invest in new assets. They forecast the price of electricity for the next year, two years and five years based on historical data for each technology considered. This is then used in the profitability check and the NPV by the firms for their respective assets or future investments.

Sensing

The sensing of the agents is limited. Only the firm agents have sensing. They have a clear and full understanding of the performance of their assets. This includes the costs involved, the electricity generated and sold, and for some technologies, the reservoir related values. For investments, agents only inform their investment potential based on what assets are present in the system and the overall price of electricity. They do not have knowledge of other firm's assets in construction or planned. This can therefore lead to periodical supply surplus.

Stochasticity

Most of the model is deterministic. Some outages can occur randomly for each of the plants. Scenarios also provide some stochasticity to the simulation.

Observation

The model produces a large amount of data. Not all of it is necessary for testing, understanding and analysis. Some of the data needs to be collected to feed the policy process model. The agents in the policy process base their decision on what is going on with a set of key performance indicators in the electricity model. Beyond this, the interest for understanding and analysis is mostly focused on the electricity prices, the number of outages (if any), the supply mix, and the trade with foreign countries. Depending on the study being performed, the amount of investment is also of interest along with the type of investments and measures related to the goals of the Energy Strategy 2050.

A.5 Initialisation

The model is initialised with values from 2018 for all of the assets that are present in the model. This includes the 2018 Swiss electricity power plants distribution and costs. The initialisation state is always the same for all simulations. All the values considered are informed on the Swiss electricity sector directly. This is detailed further in van Baal (2016) and van Baal (2019).

A.6 Input data

There are a lot of input data required to simulate the electricity market system. The data used to run the model is given below:

- Asset investment (type, sizes and costs)
- The gas prices for thermal power plants (scenario based)
- The emission prices for thermal power plants (scenario based)
- The water inflow in Swiss reservoirs for hydro power plants yearly and hourly (scenario included)
- The waste inflow in Swiss waste management facilities yearly (scenario based)
- The price of nuclear fuel (scenario based)
- The amount of solar radiation hourly (based on the years 2015, 2016 and 2017)
- The amount of wind hourly (based on the years 2015, 2016 and 2017)
- The amount of run of river water (based on the years 2010, 2011, 2012, 2013 and 2014)
- The average hourly electricity price in France, Germany and Italy (based on the years 2015, 2016 and 2017)
- The average border capacity (import and export) with France, Germany and Italy (based on the years 2015, 2016 and 2017)

A.7 Submodels

There is a large number of submodels that are used to simulate the Swiss electricity market. They are all detailed qualitatively within this section. The equations used are present in the appendix for each submodel.

1. The spot market
2. The electricity price forecast
3. The profitability calculation

4. The NPV calculation
5. The end of life actions
6. The international trading
7. The demand aspect of storage in the model

The spot market

The spot market is at the centre of the model. Its role is to match supply with demand. Some of the demand is inelastic and has to always be met. Some of it is elastic and will be met depending on the supply price. The spot market includes all of the assets (supply and demand wise) and the international trading. It is cleared on an hourly basis using a merit order curve.

The spot price is calculated using a merit order curve. The cheapest technologies are first selected and then depending on demand, the price moves up to account for other technologies. In the cases where there is not enough supply, the Value of Lost Load (VOLL) is set at 3000 CHF per MWh.

The supply that is considered for the spot market is made of: hydropower (including run-of-river, reservoir and pumped storage), nuclear power, CCGT, solar and wind power, long term French nuclear import contracts, interruptible contracts (dischargeable generation option), and thermal power (including green CHP, waste burning power plants, other thermal).

The electricity price forecast

The electricity price forecast is used by the firms to gain an understanding of the market and helps them assess whether future investments are worth the expenses. This price forecasts consists of estimating a linear relation for the future in the form $y = mx + p$. Therefore finding a slope (m) and a constant (p) for future prices based on prices from the previous four years. This is done using a weighted average of the last three years of prices and is updated throughout the simulation based on the evolution of the price of electricity for each technology.

The profitability calculation

Towards the end of life of an asset, within ten years of the end of life, the one year and five profitability of the assets are assessed monthly by the owners. Then several options present themselves. If the one year profitability is negative and the asset has reached its lifetime, then it is decommissioned. If the five year profitability is higher than zero but the one year profitability is negative, then the asset is mothballed. If the one year profitability is positive and the asset has been renovated less than twice, it is renovated. If not, it is decommissioned when it reaches its final age.

The NPV calculation

The NPV calculation is used by the agents to assess potential new power plants for their portfolios. The NPV is used to assess the profitability of a future plant. If that profitability is higher than the hurdle rate of the agent, then it will consider investing in the plant.

The investment pipeline

The firms can invest in three main technologies: solar, wind and thermal power plants. These investments are discrete in capacity. Only one option per technology is provided to the investors. Every month, each firm is provided with the opportunity of investing in one of the three technologies. They test the NPV of each of the plants and the most positive, if there is one, is approved by the firm. Approval at this stage means that a permit is demanded. This is a process that takes a different amount of time depending on the technology. Its rate of success also depend on the technology with the rate of success of solar assets being affected by land scarcity while the rate of success of wind assets being affected by land scarcity and social acceptance.

Once the permit has been approved, the firms will once again assess the NPV of the investment on a monthly basis. If the NPV has changed and is now negative, the firm keeps the permit without building the plant. If it becomes positive, then construction is started. The plant then comes online only after the building period has been completed.

The international trading

International trading of electricity is introduced in the model. The import and export prices of the electricity are known from historical data for Germany, Italy and France. The supply of this electricity is then limited by the inter-connections to these different countries.

This international trading is supplemented by the long-term contracts that Switzerland has with France. Such contracts take a part of the capacity on the interconnections between France and Switzerland, limiting the potential for international trading.

The demand aspect of storage in the model

Demand is mostly present in the model through the inelastic demand of Swiss consumers. The demand of foreign countries and the demand of storage technologies such as hydro-pumping is also included in the model. All these aspects are taken into account in the spot market to make sure demand is met by supply.

B The Swiss electricity market model hybridisation equations

The present appendix outlines the equations that are used to connect the Swiss electricity market model to the policy process model. This consists of all the equations that are used to calculate the actual beliefs of the policy agents based on key indicators in the electricity market model.

To obtain the actual beliefs, the indicators from the Swiss electricity market model needs to be extracted from the model and mapped onto an interval $[0,1]$. The procedure to do this is shown here.

The secondary issues are the following: renewable energy production (S1), electricity prices (S2), renewable energy investment level (S3), domestic level emissions (S4) and imported emissions (S5).

The renewable energy production issue (S1) is calculated as follows. The total supply of electricity is given by (note that I only consider domestic production and no imports or exports):

$$S_{total} = S_{solar} + S_{CCGT} + S_{wind} + S_{nuclear} + S_{hydro} + S_{hydrop} + S_{waste} + S_{ROR} \quad \text{(B.1)}$$

The renewable supply is given by:

$$S_{RES} = S_{solar} + S_{wind} S_{hydro} + S_{hydrop} S_{ROR} \quad \text{(B.2)}$$

The indicator is then normalised using:

$$S1 = S_{RES}/S_{total} \tag{B.3}$$

The electricity prices issue (S2) is calculated based on the average electric price of the previous year. It is then normalised using an expected maximum electricity price ($P_{elec,max}$). The normalisation equation is given by:

$$S2 = \frac{P_{elec}}{P_{elec,max}} \tag{B.4}$$

$P_{elec,max}$ is defined in Table B.1.

The renewable energy investment level issue (S3) is calculated using the investment performed by all the investors in solar, wind and CCGT assets.

$$I_{total} = I_{wind} + I_{solar} + I_{CCGT} \tag{B.5}$$

$$I_{RES} = I_{wind} + I_{solar} \tag{B.6}$$

The indicator is normalised using:

$$S3 = I_{RES}/I_{total} \tag{B.7}$$

The domestic level emissions issue (S4) is calculated based on the CCGT emissions. The normalisation of this indicator is once again done using an assumed maximum for the emissions which is given as five times the emissions for year 1 of the simulation. This is an arbitrary value that can be tuned to make sure that the indicator is always between 0 and 1.

$$S4 = S_{CCGT}/S_{CCGT,max} \tag{B.8}$$

The imported emissions issue (S5) is calculated based on the imports and the policy mixes of the countries from which Switzerland imports. The policy mixes are scenarios that are obtained from technical report and goals for the different countries. For each country, the percentage of coal and gas production is considered to calculate the imported emissions.

$$E_{FR} = (S_{FR,NTC} + S_{LTC}) \cdot (M_{FR,CCGT} \cdot E_{CCGT} + M_{FR,coal} \cdot E_{coal}) \tag{B.9a}$$

$$E_{DE} = Ss_{DE,NTC} \cdot (M_{DE,CCGT} \cdot E_{CCGT} + M_{DE,coal} \cdot E_{coal}) \tag{B.9b}$$

$$E_{IT} = S_{IT,NTC} \cdot (M_{IT,CCGT} \cdot E_{CCGT} + M_{IT,coal} \cdot E_{coal}) \tag{B.9c}$$

where E are the emissions per type of technology and where M is the share of the mix for a specific technology in the country.

To normalise this indicator, we once again select a maximum amount of emissions. This is calculated as being 5% higher than the initial imported emissions in year 1 of all these countries. Considering the emissions should decrease in the scenarios, this means that the indicator should remain within the $[0, 1]$ interval. The exact value is provided in Table B.1.

$$S5 = \frac{E_{FR} + E_{DE} + E_{IT}}{E_{total}} \tag{B.10}$$

where $E_{total} = E_{FR,init} + E_{DE,init} + E_{IT,init}$

The policy core issues are given as the economy issue (PC1) and the environment issue (PC2). They are calculated using weighted averages of a number of secondary indicators. The equations used can be tuned but the ones implemented are given below:

$$PC1 = \frac{3}{4} \cdot S2 + \frac{1}{4} \cdot S3 \tag{B.11}$$

$$PC2 = \frac{1}{4} \cdot S1 + \frac{1}{4} \cdot S3 + \frac{1}{4} \cdot S4 + \frac{1}{4} \cdot S5 \tag{B.12}$$

	Issues				
	S1 RES [%]	S2 Price [CHF/MWh]	S3 REI [%]	S4 Dom. em. [tons]	S5 Imp. em. [tons]
Minimum	0	0	0	0	0
Maximum	100	200	100	20 000 000	90 00 000

Table B.1 – Minimum and maximum values used to map the indicators from the electricity market model to the secondary issues on an interval [0,1].

C The simplest implementation predation model results

This appendix collects all of the graphs that were plotted using the results from the simplest implementation model with the predation model. The graphs show the results for all of the scenarios in a .PNG compressed format. The results can also be found in higher quality on GitHub[1].

[1] https://github.com/kleinrap/policyemergence_SM/tree/master/3_PredationModel/0_ResultsGraphs

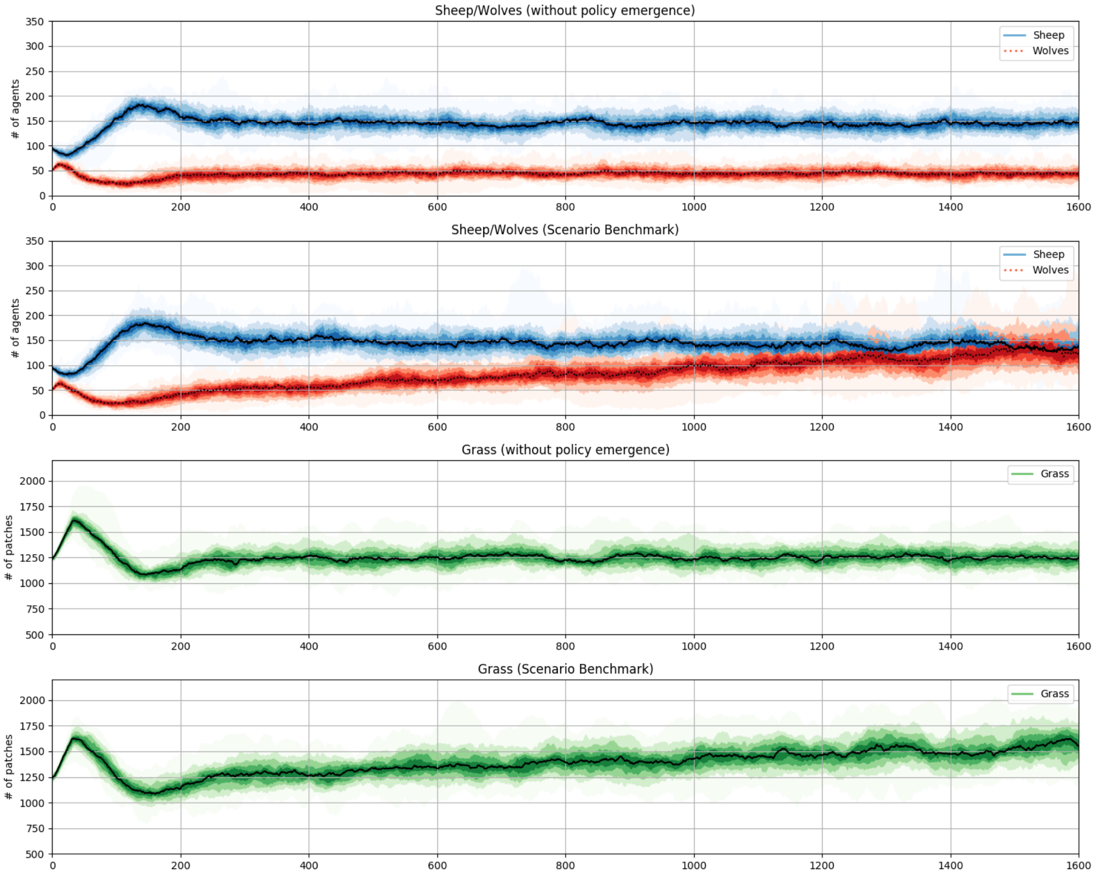

Figure C.1 – Predation model populations for the predation model and the benchmark scenario.

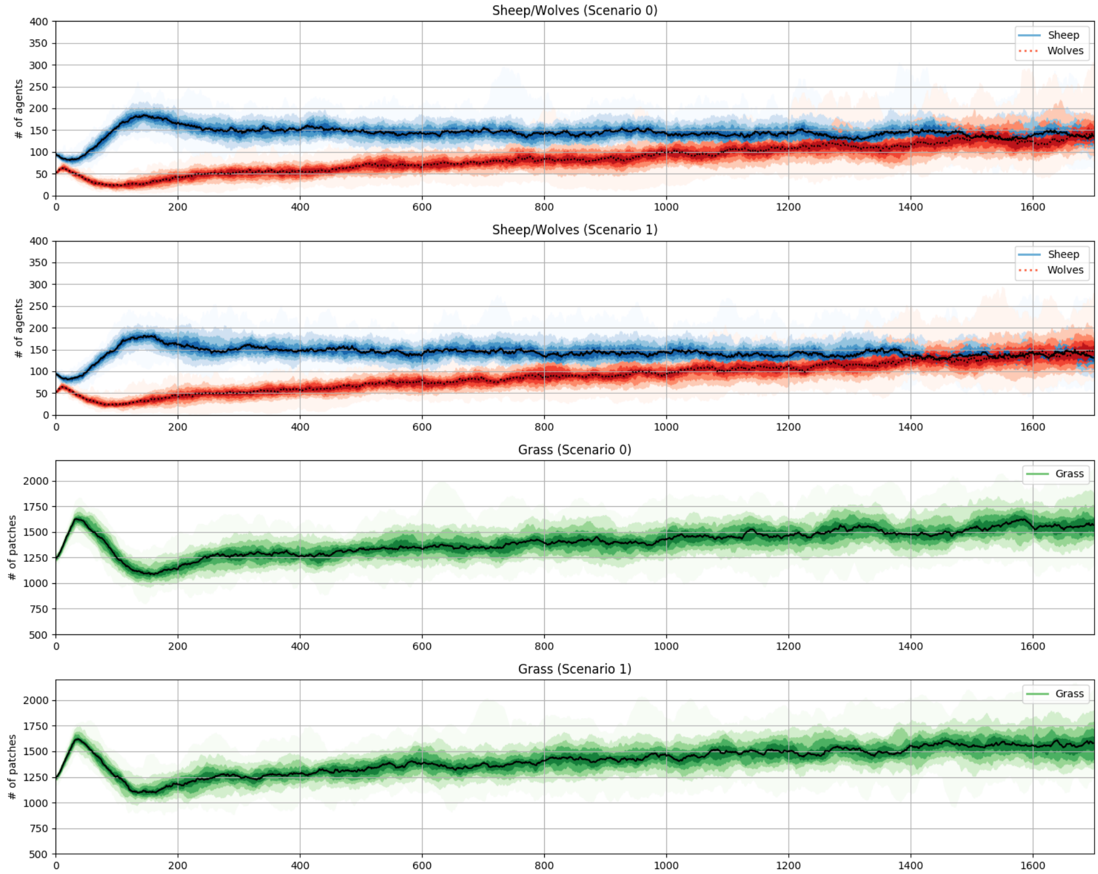

Figure C.2 – Predation model populations for the benchmark scenario and scenario 1.

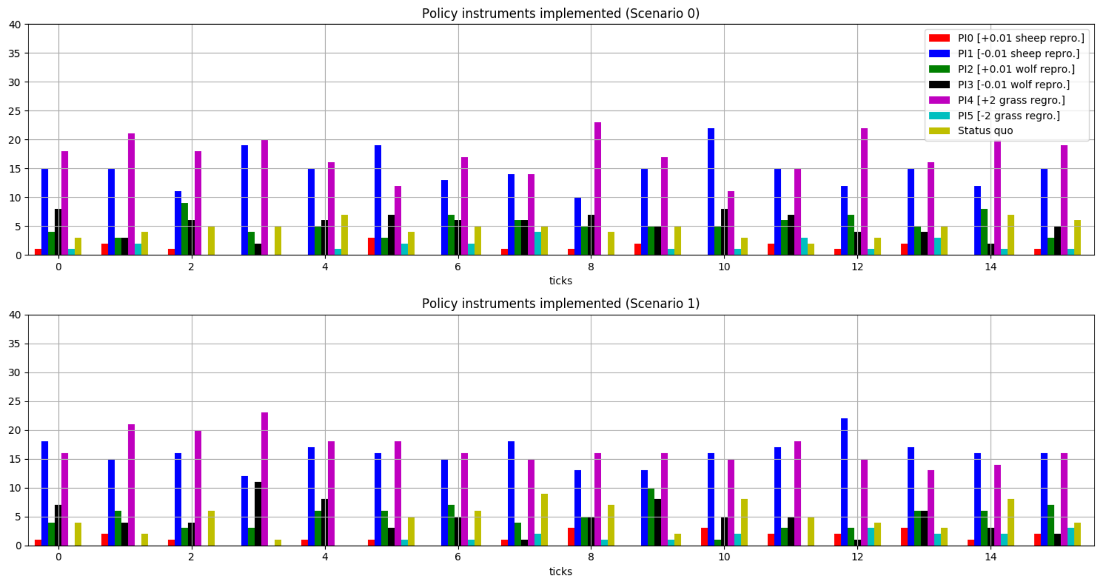

Figure C.3 – Policy instrument selection for the benchmark scenario and scenario 1.

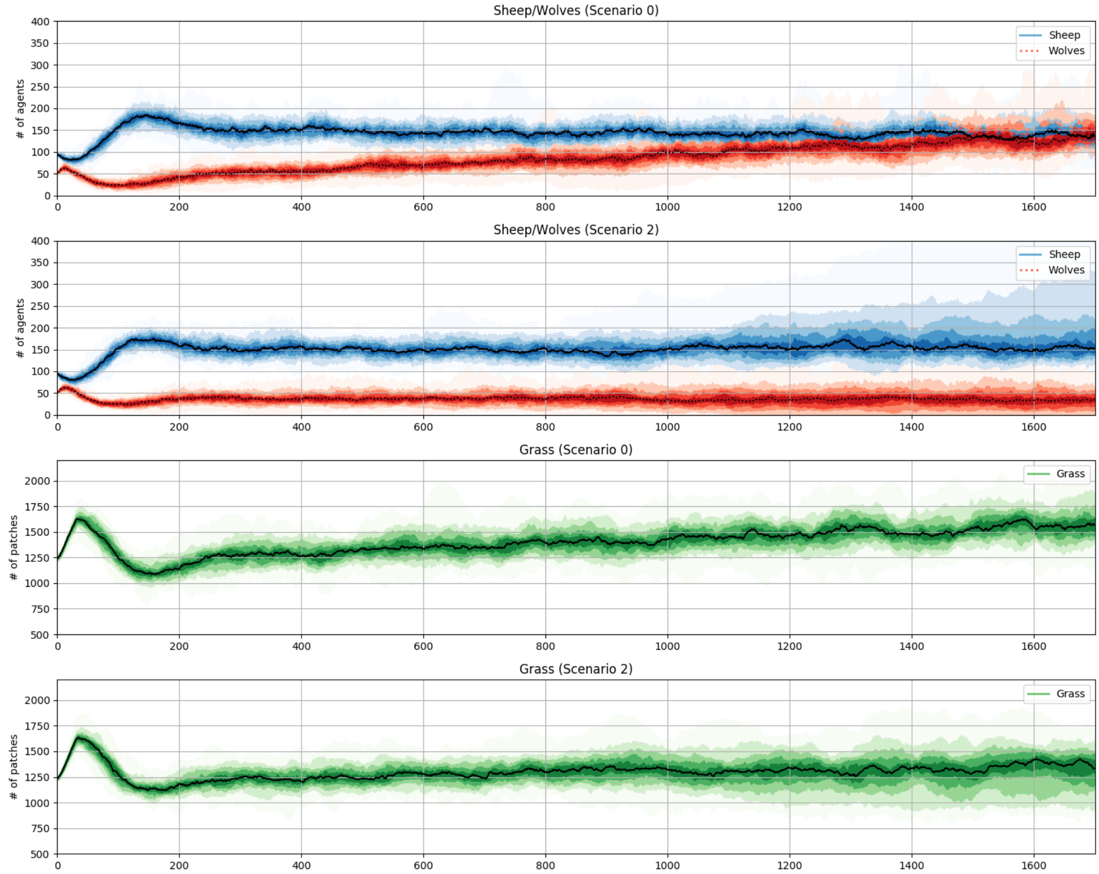

Figure C.4 – Predation model populations for the benchmark scenario and scenario 2.

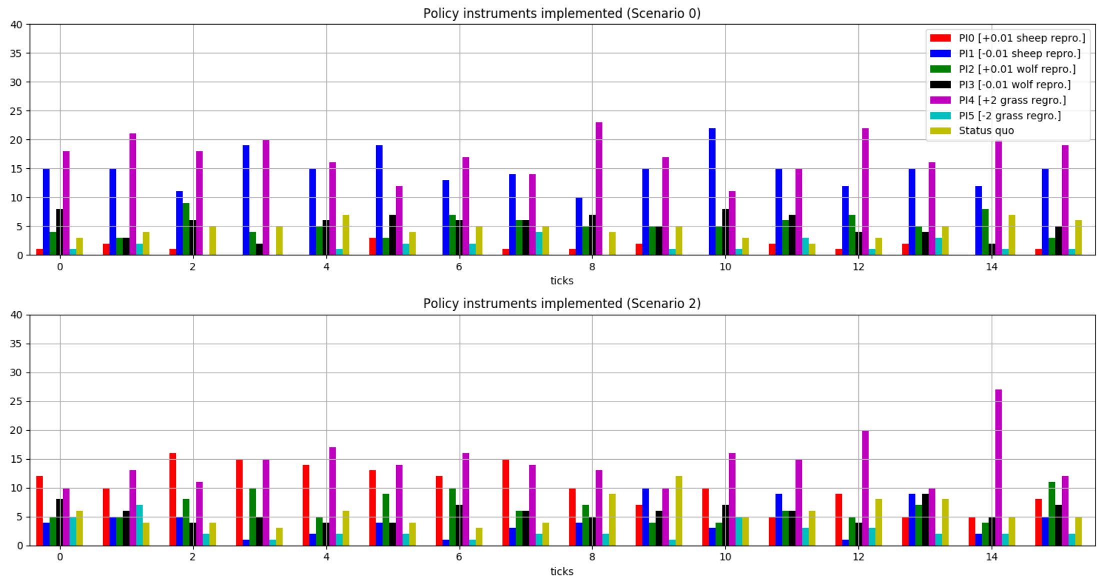

Figure C.5 – Policy instrument selection for the benchmark scenario and scenario 2.

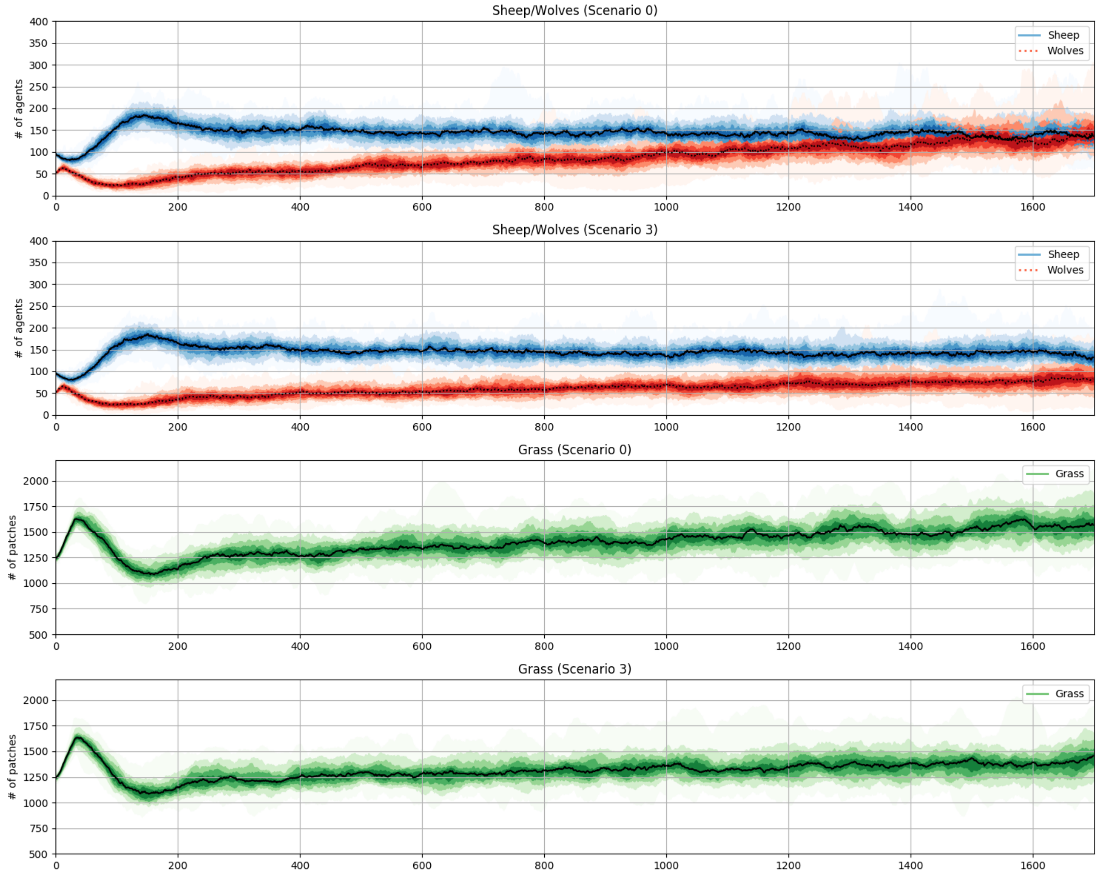

Figure C.6 – Predation model populations for the benchmark scenario and scenario 3.

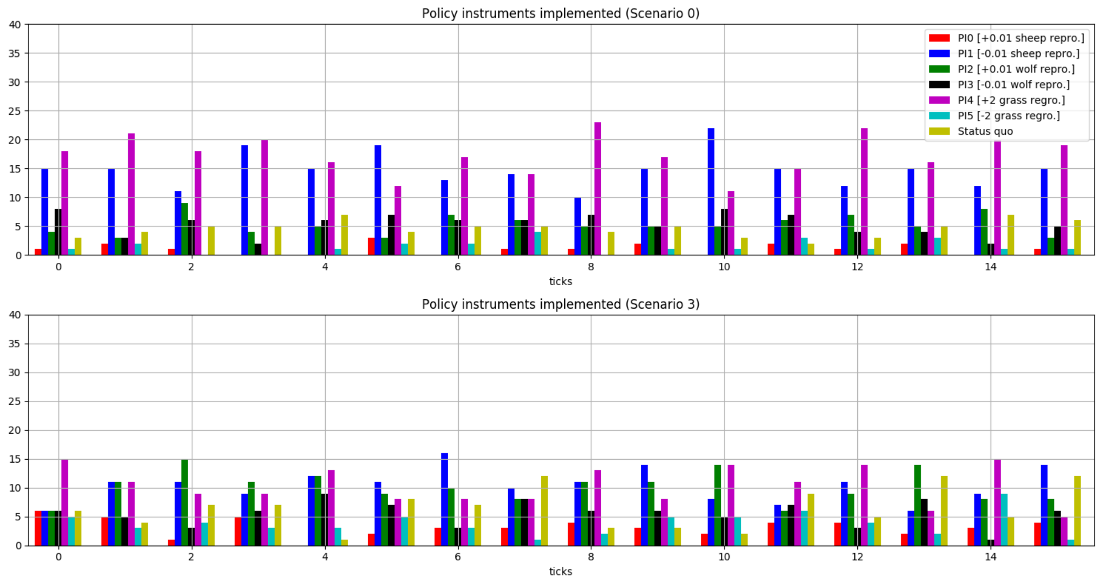

Figure C.7 – Policy instrument selection for the benchmark scenario and scenario 3.

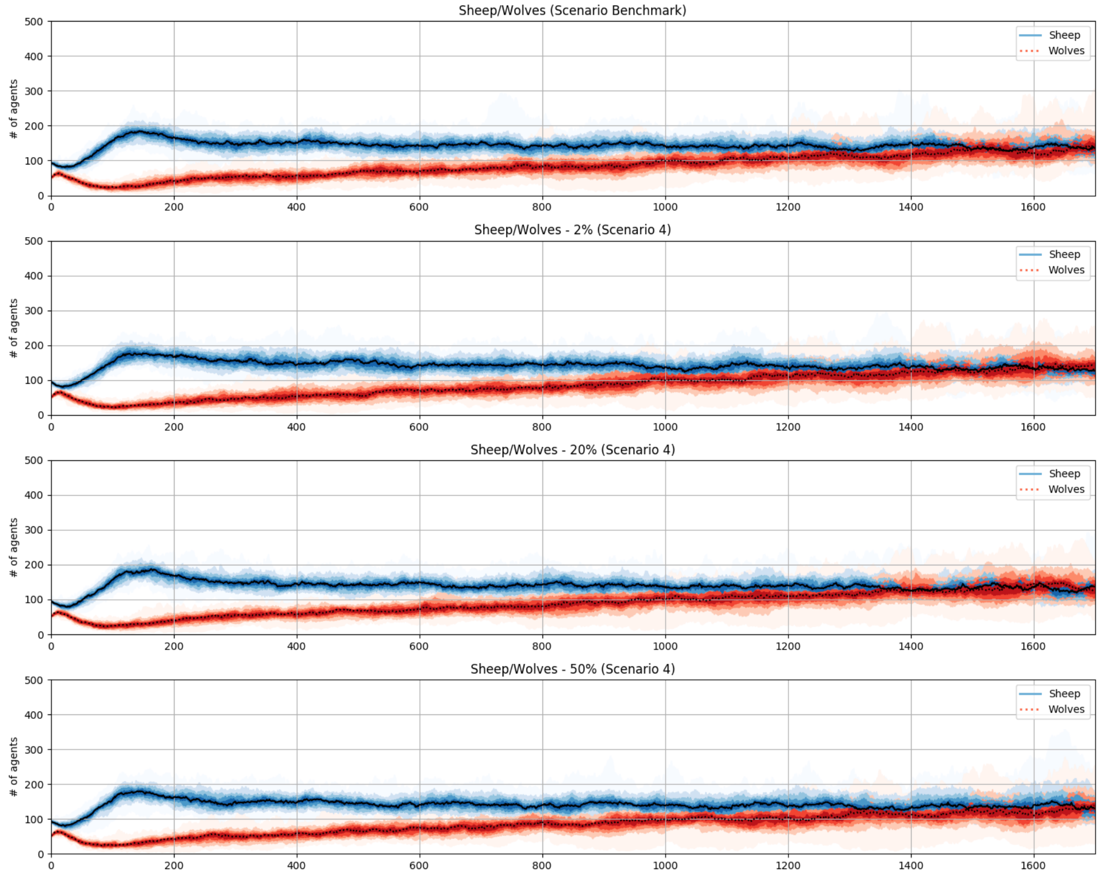

Figure C.8 – Predation model populations for the benchmark scenario and scenario 4.

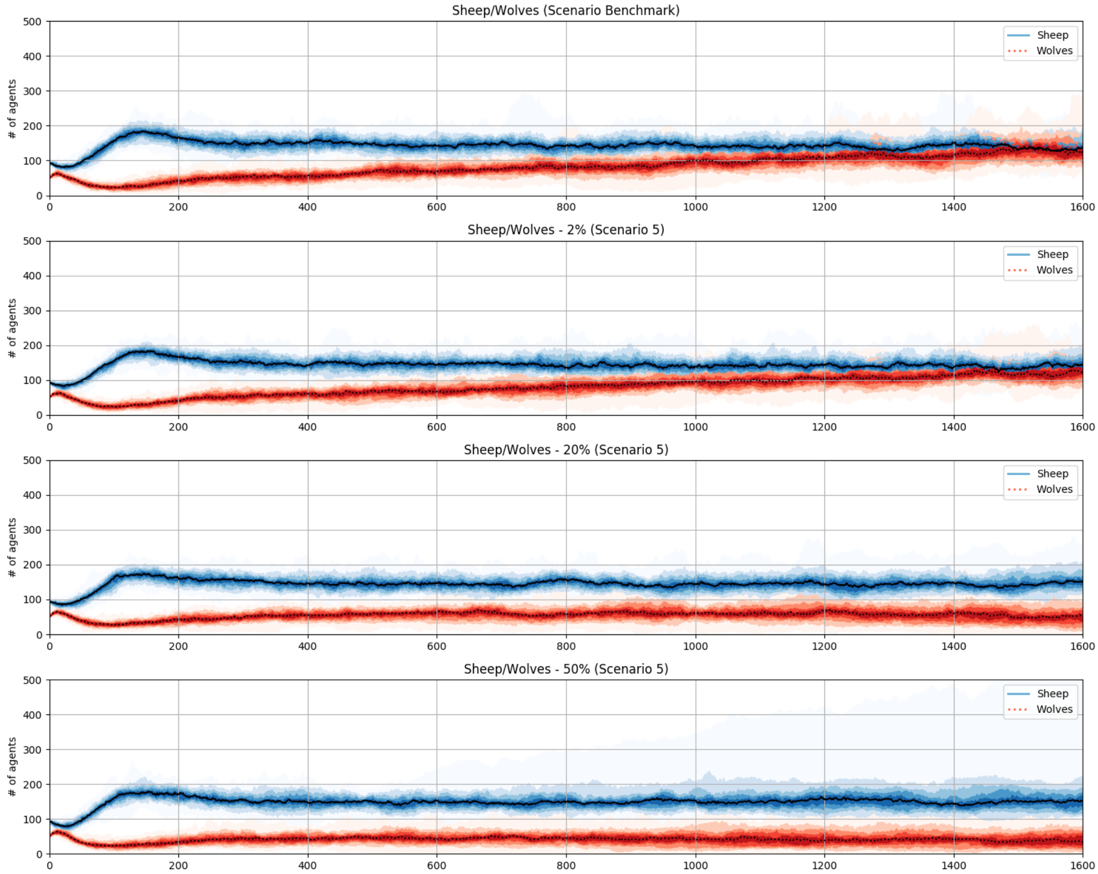

Figure C.9 – Predation model populations for the benchmark scenario and scenario 5.

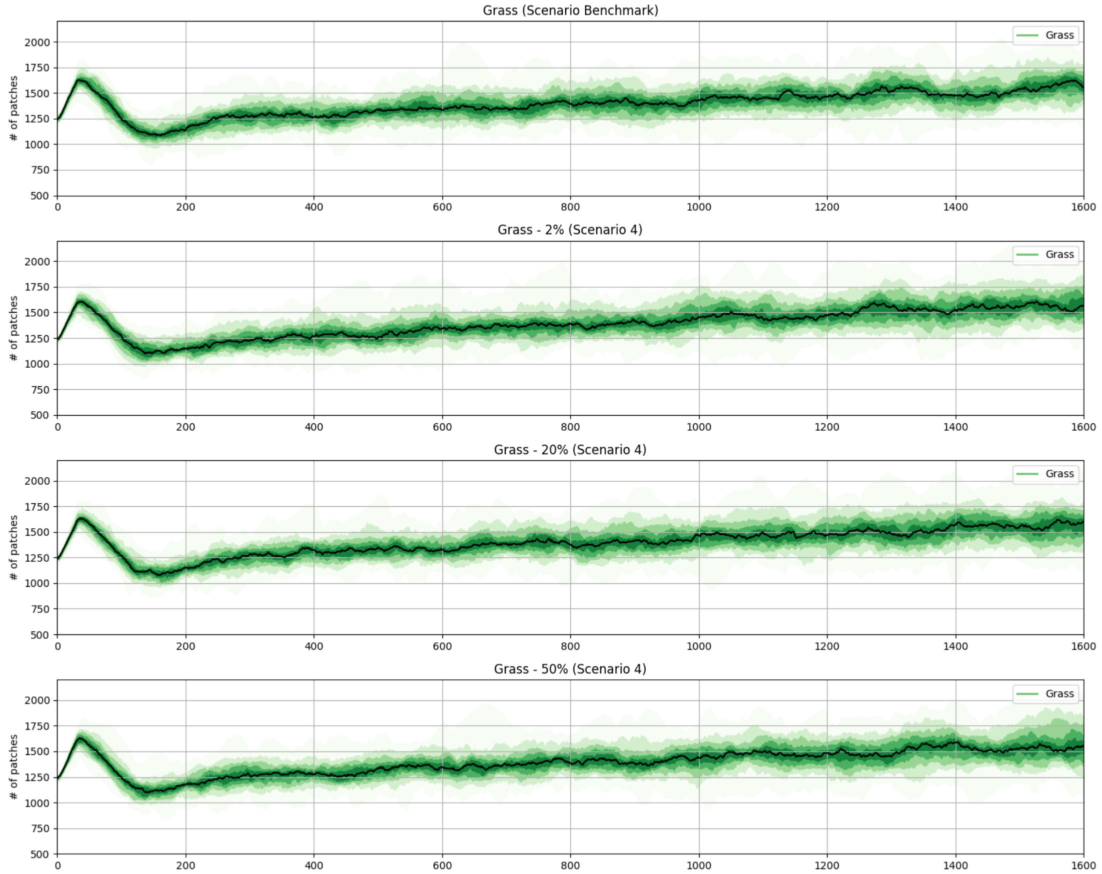

Figure C.10 – Predation model populations for the benchmark scenario and scenario 4.

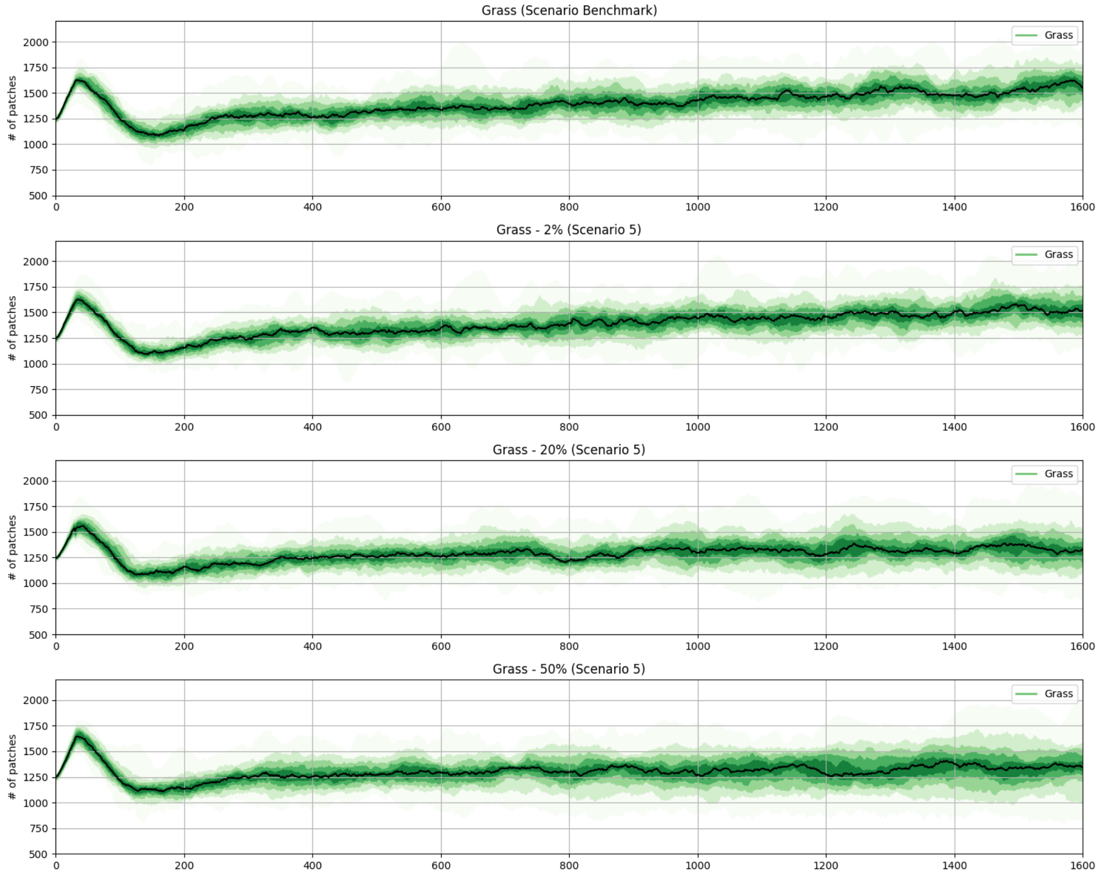

Figure C.11 – Predation model populations for the benchmark scenario and scenario 5.

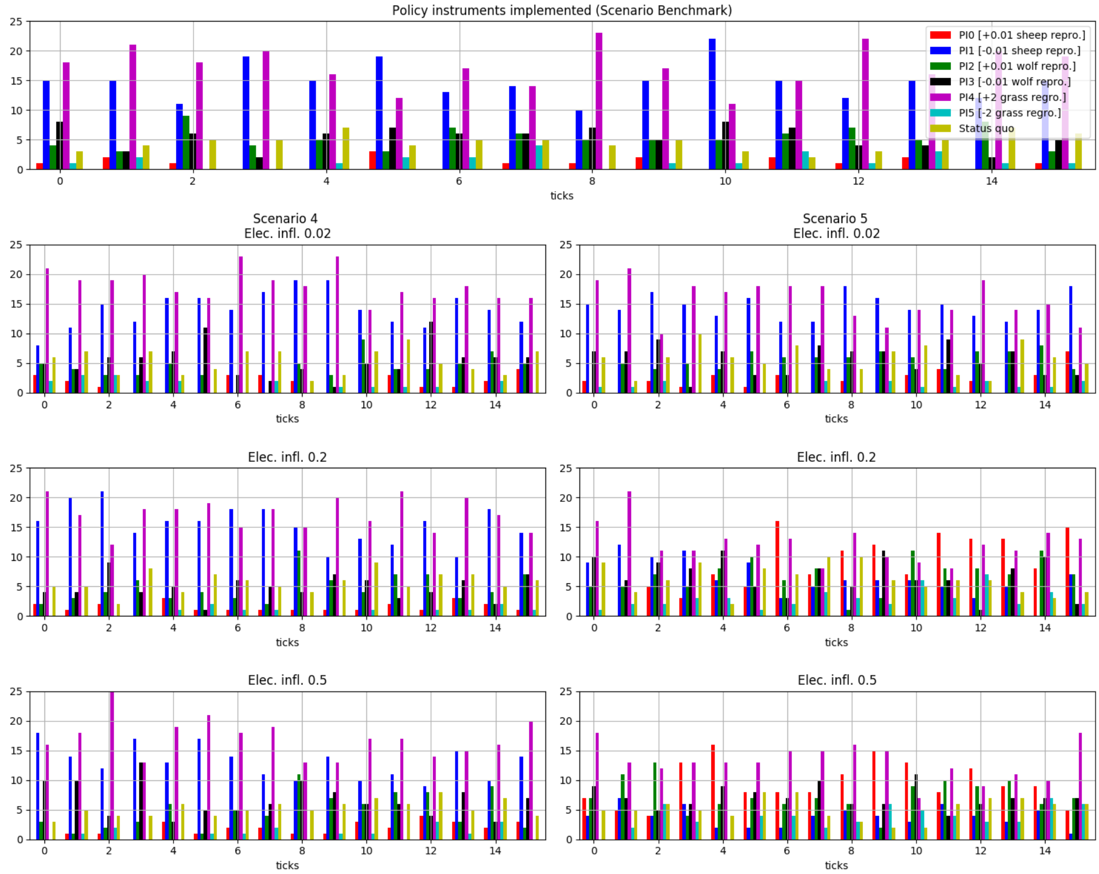

Figure C.12 – Policy instrument selection for the benchmark scenario, and scenarios 4 and 5.

D The simplest implementation electricity model results

This appendix collects all of the graphs that were plotted using the results from the simplest implementation model with the Swiss electricity market model. The graphs show the results for all of the scenarios in a .PNG compressed format. The results can also be found in higher quality on GitHub[1].

[1] https://github.com/kleinrap/policyemergence_SM/tree/master/2_ElectricityModel/0_ResultsGraphs

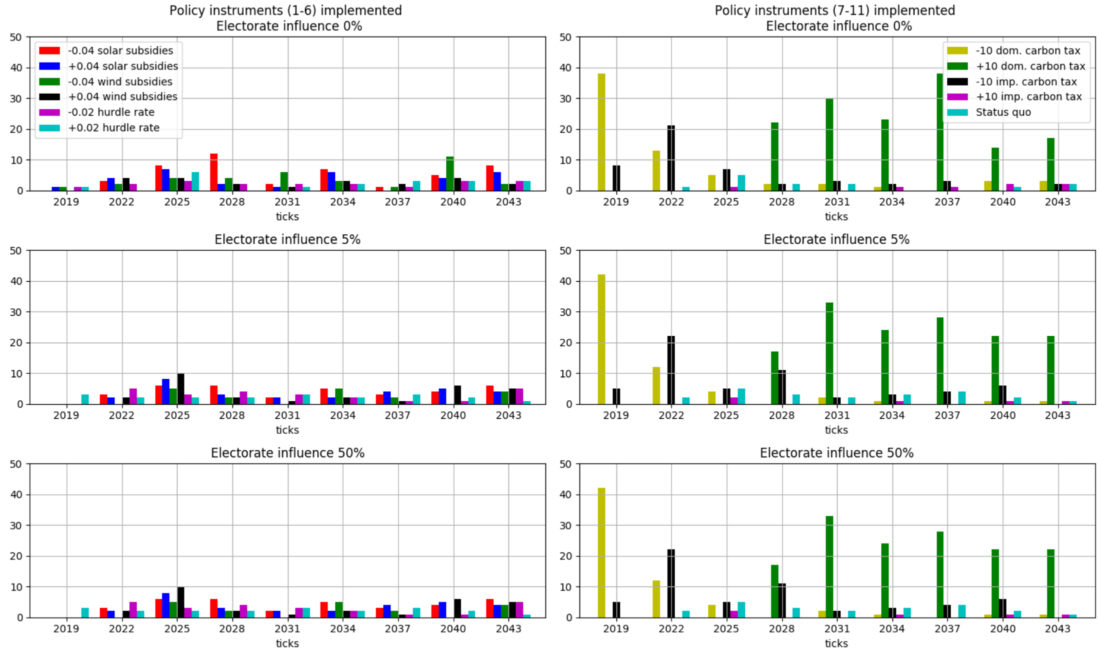

Figure D.1 – Cumulative policy instrument selection for a 3% electricity demand growth scenario for all fifty repetitions for each simulation.

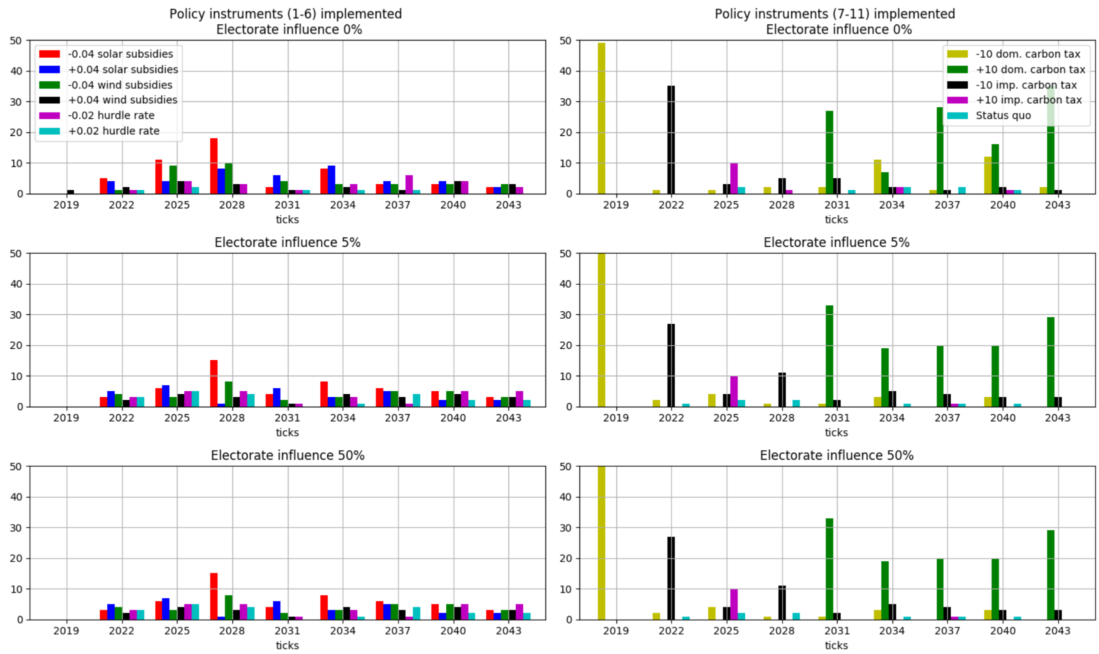

Figure D.2 – Cumulative policy instrument selection for a 1.5% electricity demand growth scenario for all fifty repetitions for each simulation.

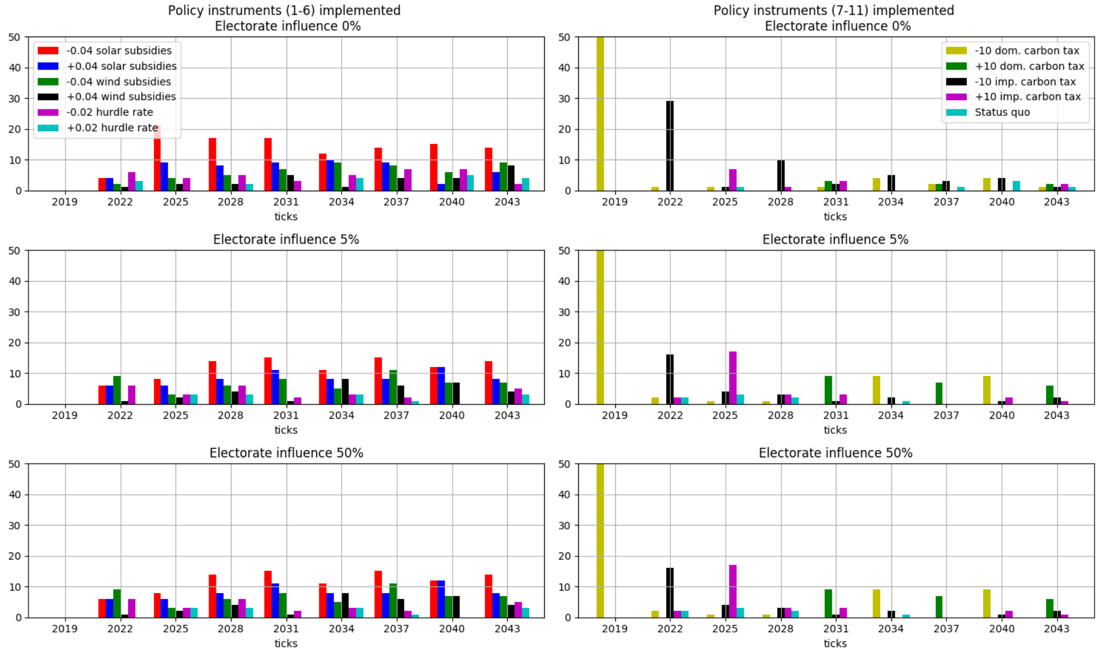

Figure D.3 – Cumulative policy instrument selection for a 0% electricity demand growth scenario for all fifty repetitions for each simulation.

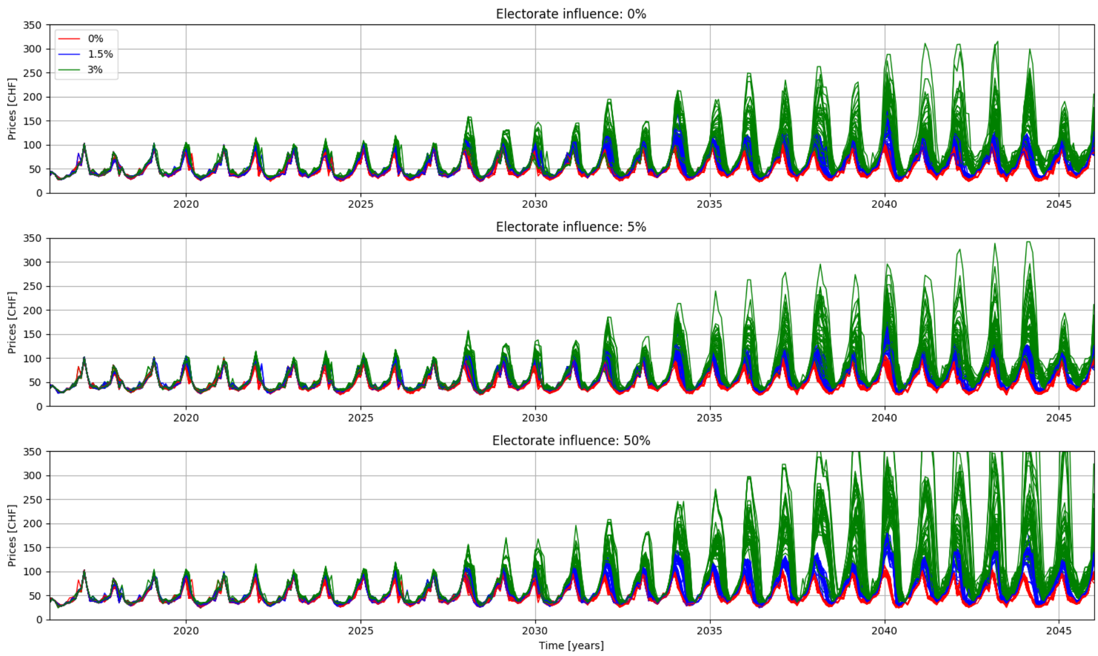

Figure D.4 – Monthly averaged electricity prices, split according to the three electorate influence rates and for the three demand growth scenarios.

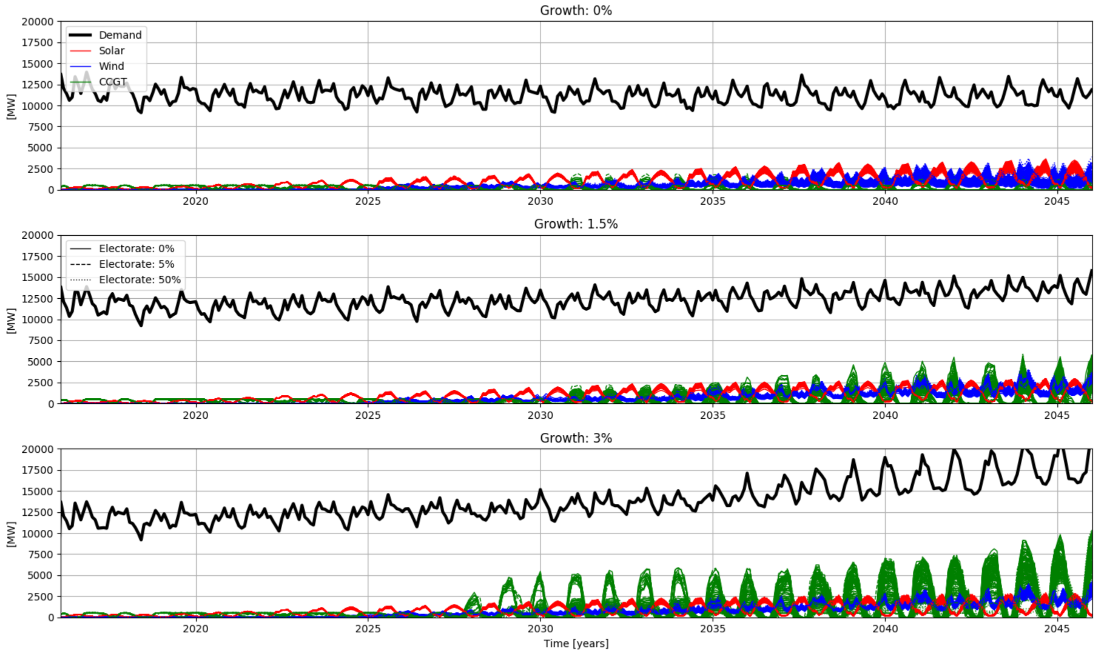

Figure D.5 – Overall demand with the electricity supplied by solar, wind and CCGT sources for all scenarios.

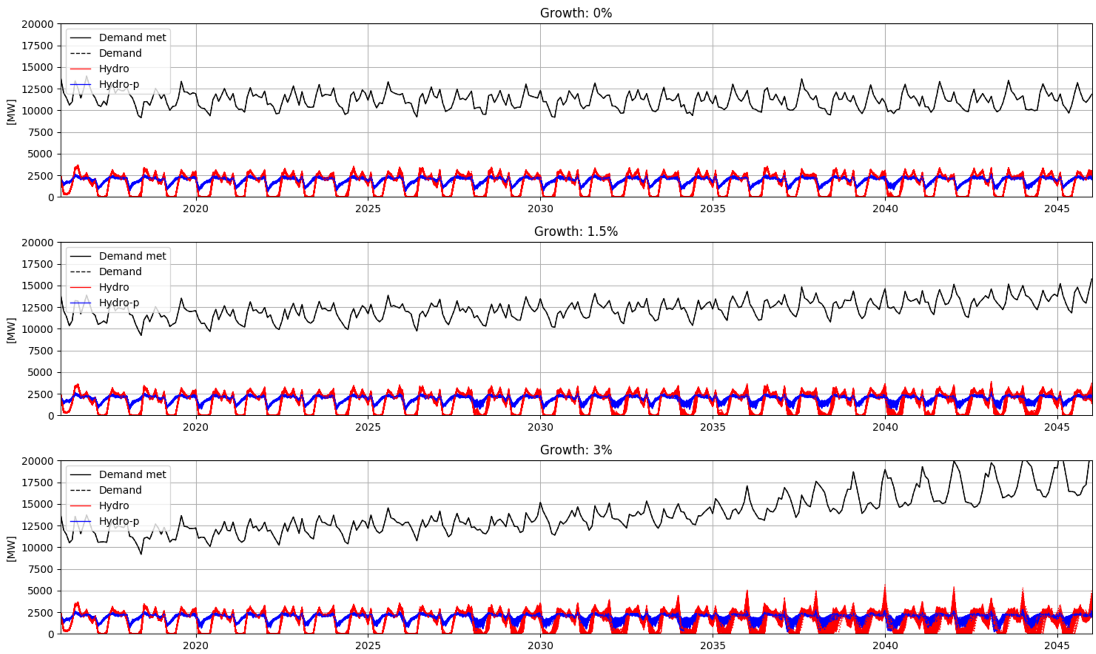

Figure D.6 – Overall demand with the hydro electricity supplied for all scenarios.

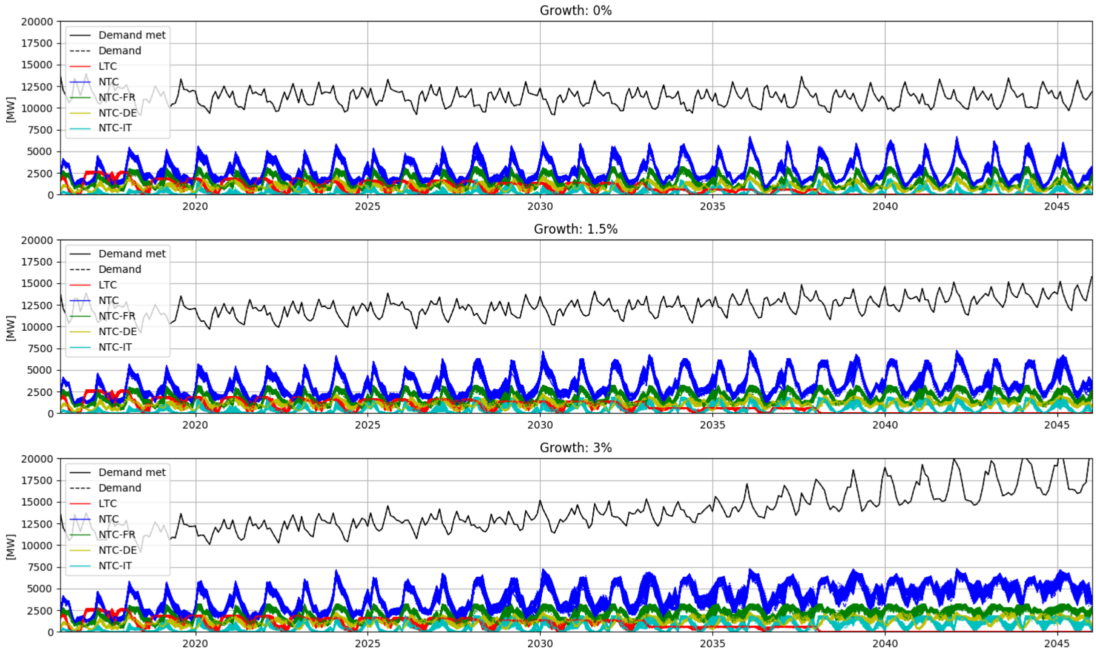

Figure D.7 – Overall demand with the electricty supplied by NTC and LTC sources for all scenarios.

E The ACF implementation predation model results

This appendix collects all of the graphs that were plotted using the results from the ACF implementation model with the predation model. The graphs show the results for all of the scenarios in a .PNG compressed format. The results can also be found in higher quality on GitHub[1].

[1]https://github.com/kleinrap/policyemergence_A/tree/master/3_PredationModel/0_ResultsGraphs

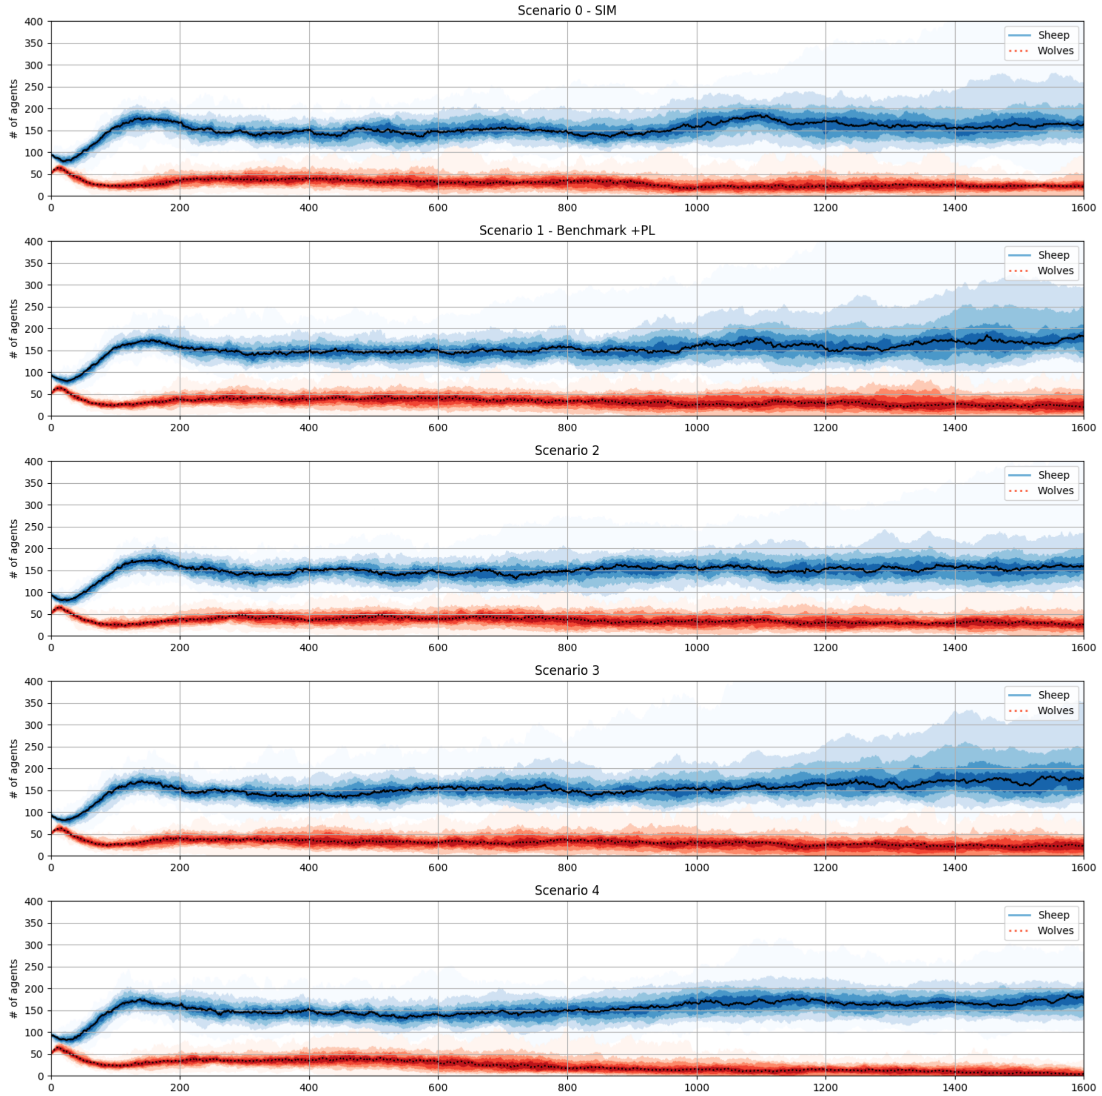

Figure E.1 – Predation model results for wolf and sheep counts for scenarios SM, 0, 1, 2, and 3 (+PL implementation).

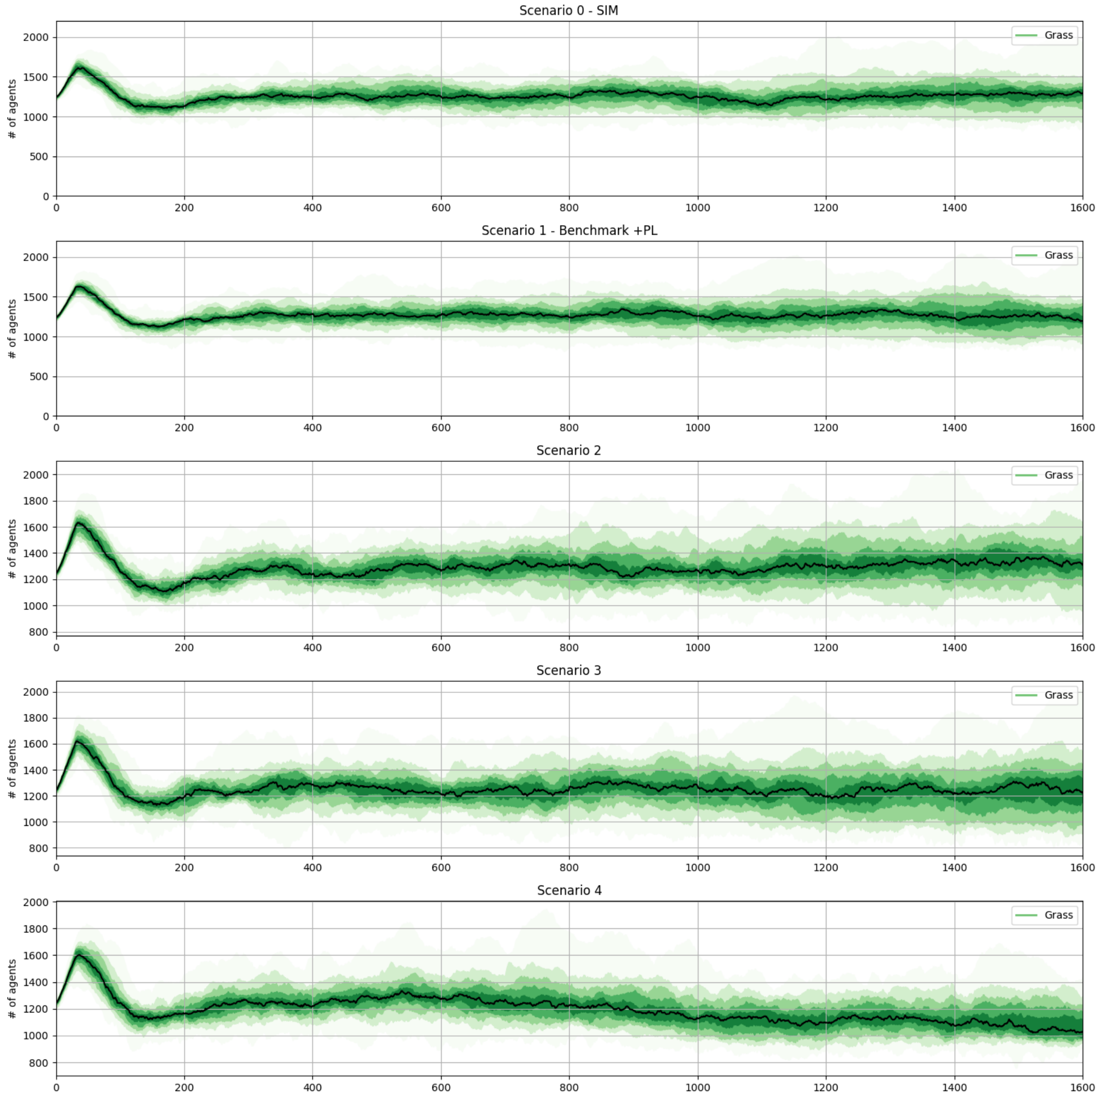

Figure E.2 – Predation model results for grass patch count for scenarios SM, 0, 1, 2, and 3 (+PL implementation).

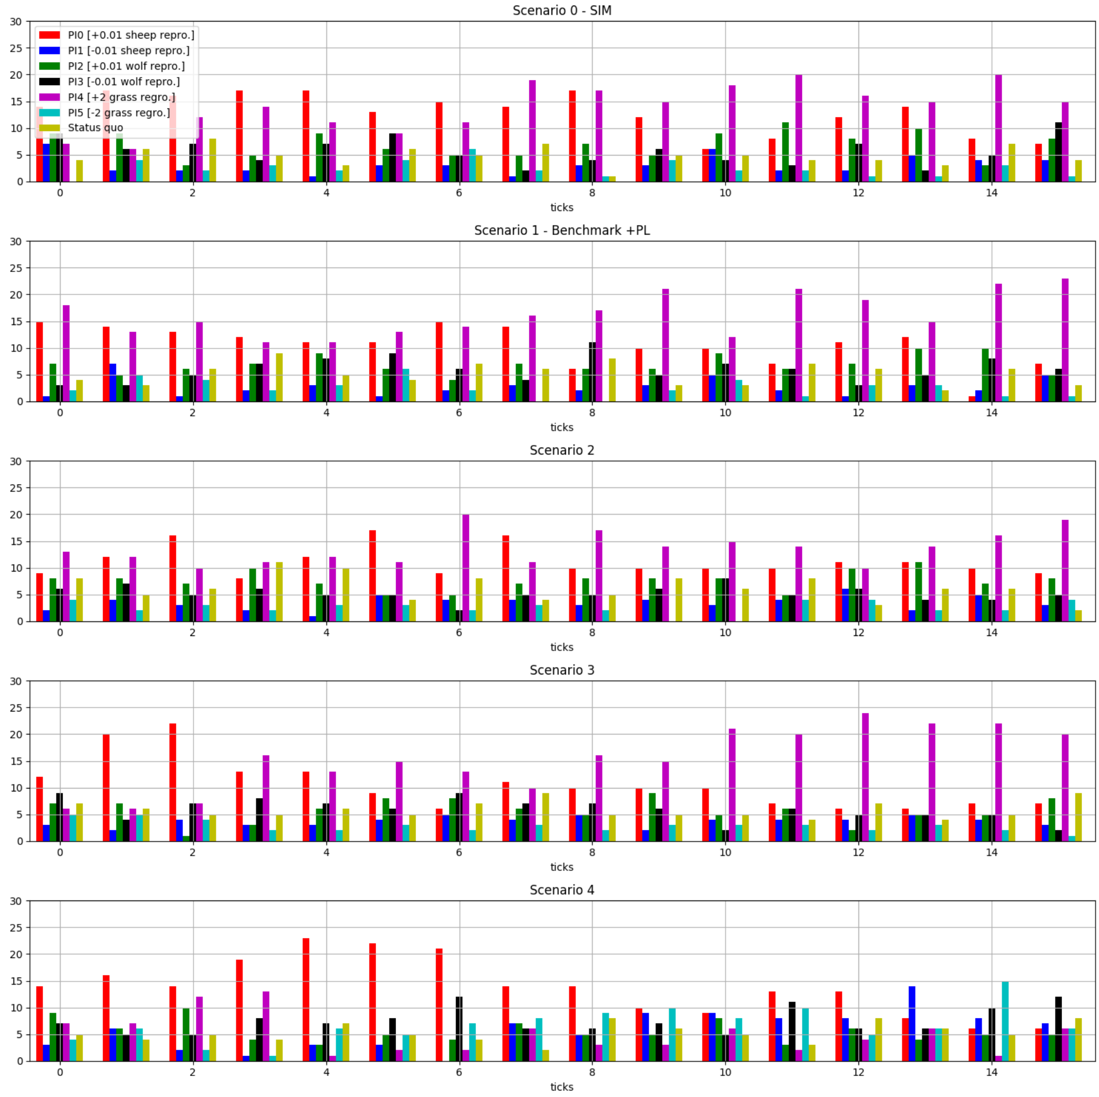

Figure E.3 – Policy instruments selected for scenarios SM, 0, 1, 2, and 3 (+PL implementation).

Figure E.4 – Policy core issues selected per affiliation for scenarios SM, 0, 1, 2, and 3 (+PL implementation).

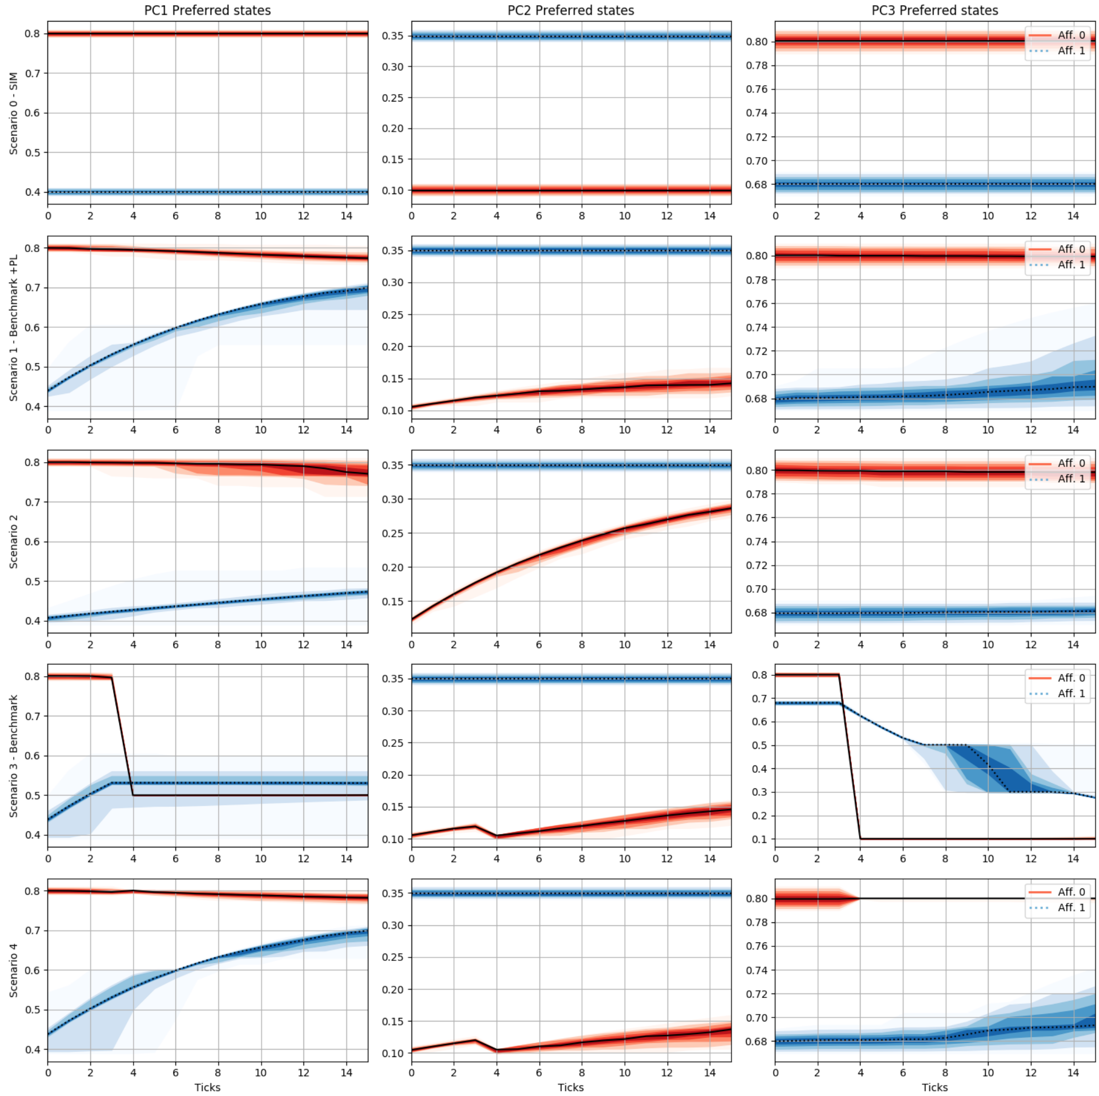

Figure E.5 – Policy core issue preferred states for scenarios SM, 0, 1, 2, and 3 (+PL implementation).

Figure E.6 – Secondary issue selected per affiliation for scenarios SM, 0, 1, 2, and 3 (+PL implementation).

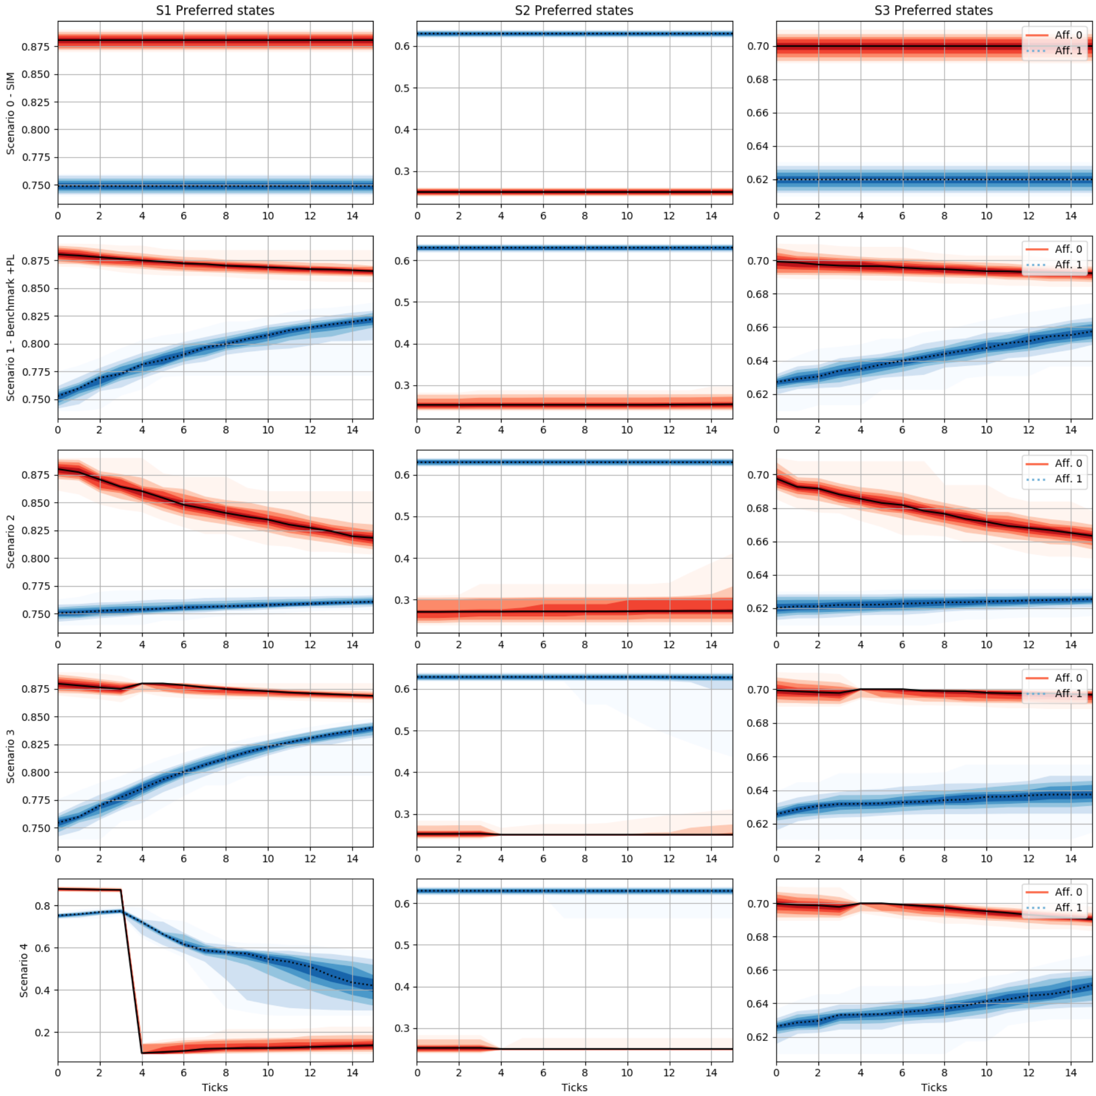

Figure E.7 – Secondary issue preferred states for scenarios SM, 0, 1, 2, and 3 (+PL implementation).

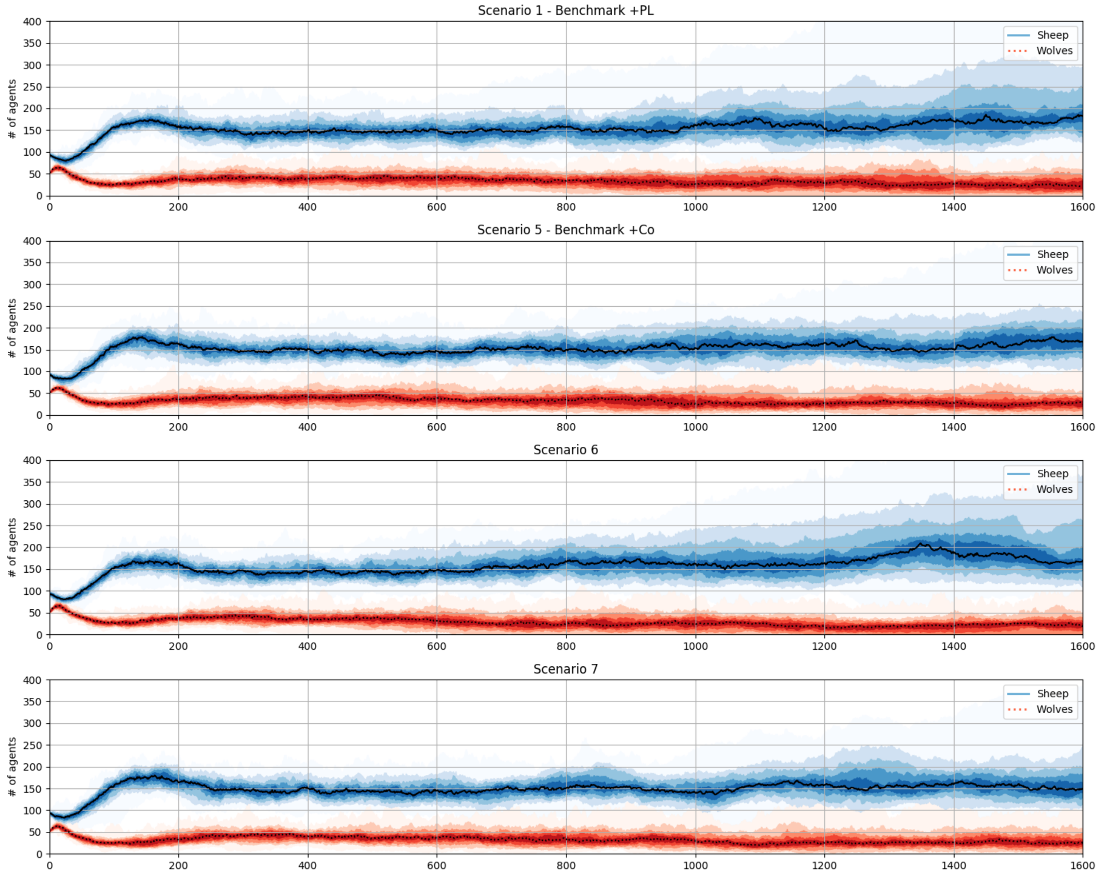

Figure E.8 – Predation model results for wolf and sheep counts for scenarios 0, 5, 6 and 7 (+Co implementation).

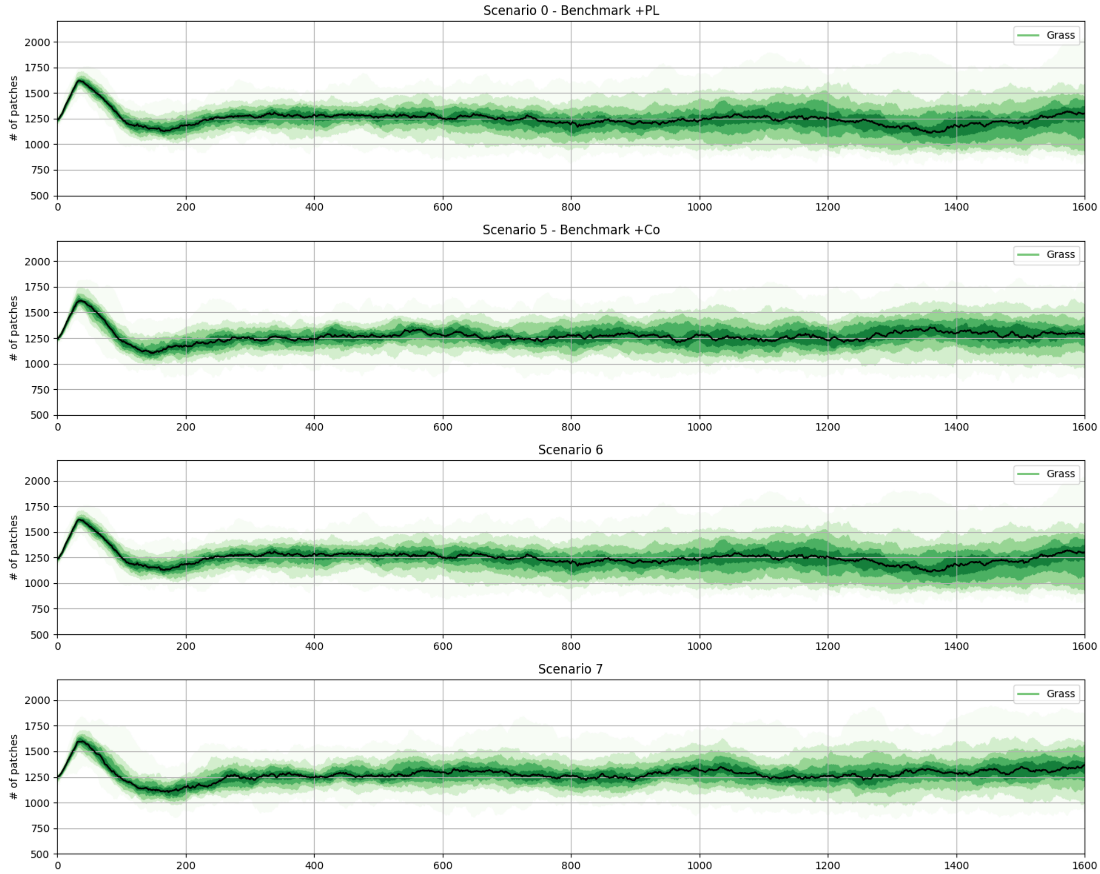

Figure E.9 – Predation model results for grass patch count for scenarios 0, 5, 6 and 7 (+Co implementation).

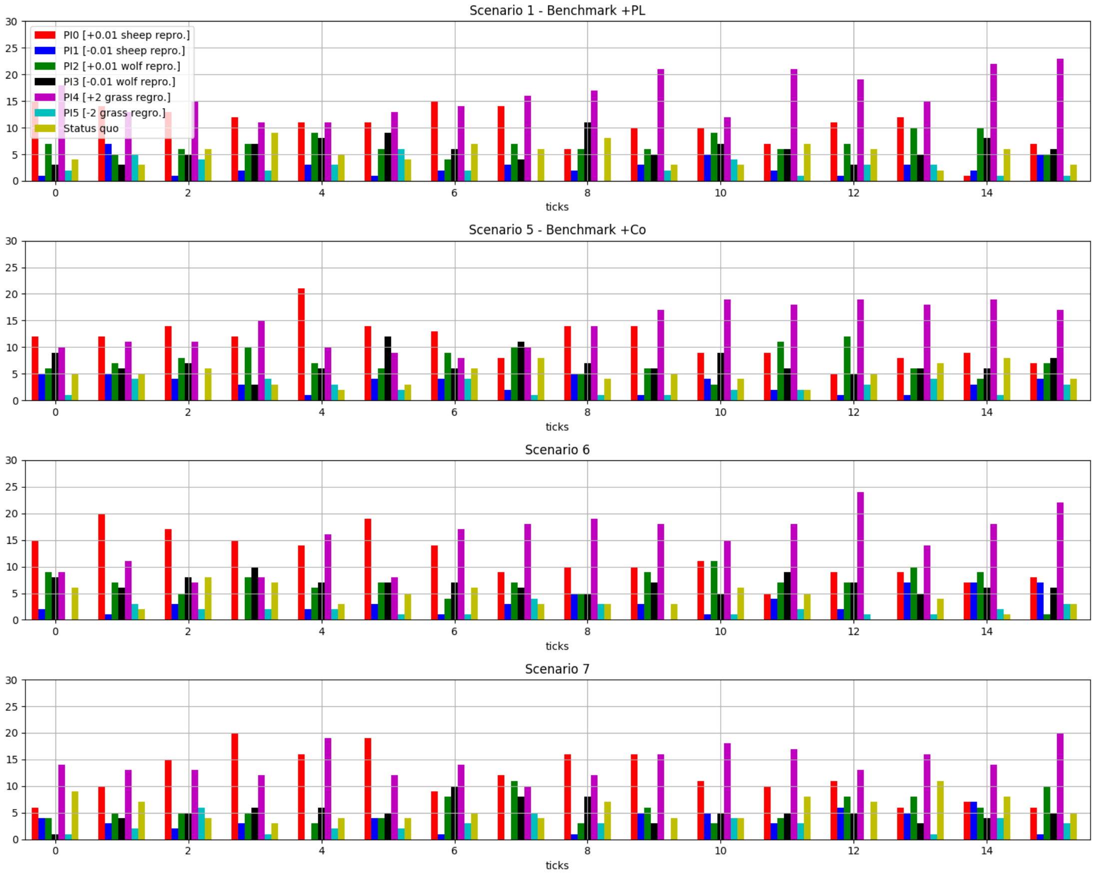

Figure E.10 – Policy instruments selected for scenarios 0, 5, 6 and 7 (+Co implementation).

Figure E.11 – Policy core issues selected per affiliation for scenarios 0, 5, 6 and 7 (+Co implementation).

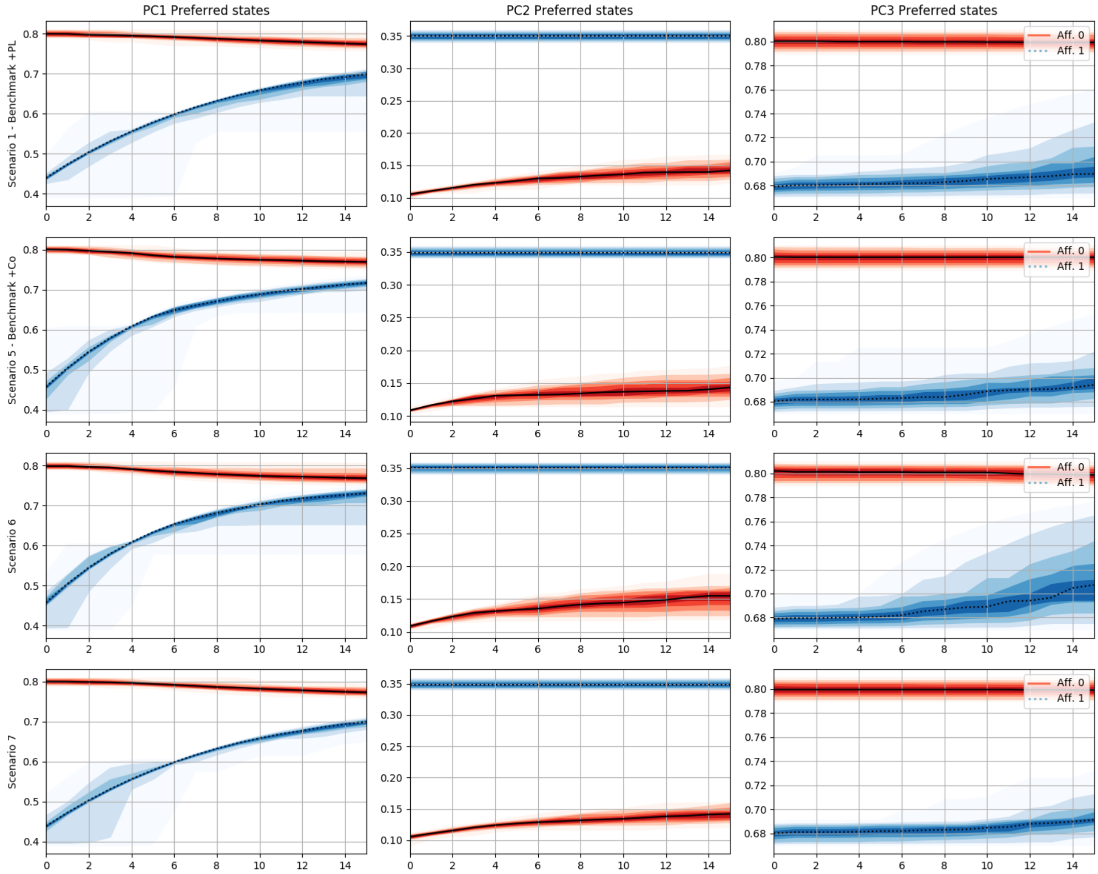

Figure E.12 – Policy core issue preferred states for scenarios 0, 5, 6 and 7 (+Co implementation).

Figure E.13 – Secondary issue selected per affiliation for scenarios 0, 5, 6 and 7 (+Co implementation).

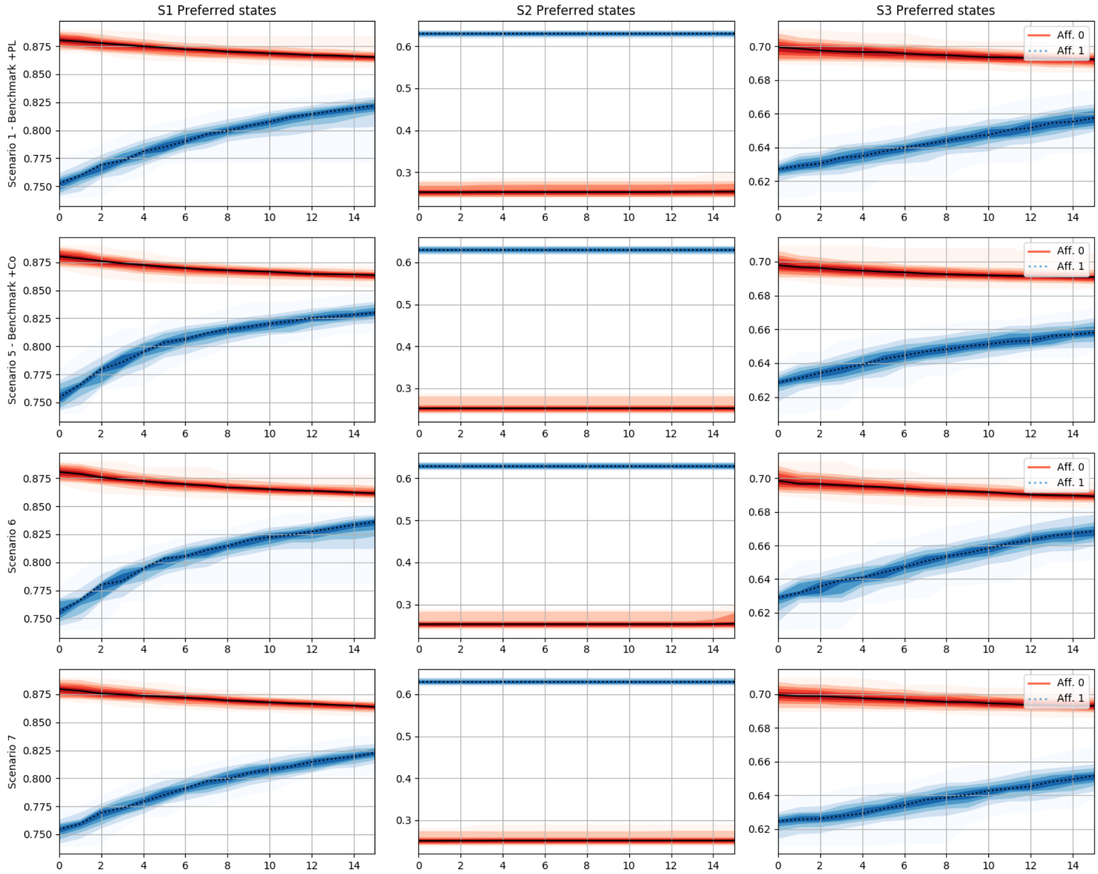

Figure E.14 – Secondary issue preferred states for scenarios 0, 5, 6 and 7 (+Co implementation).

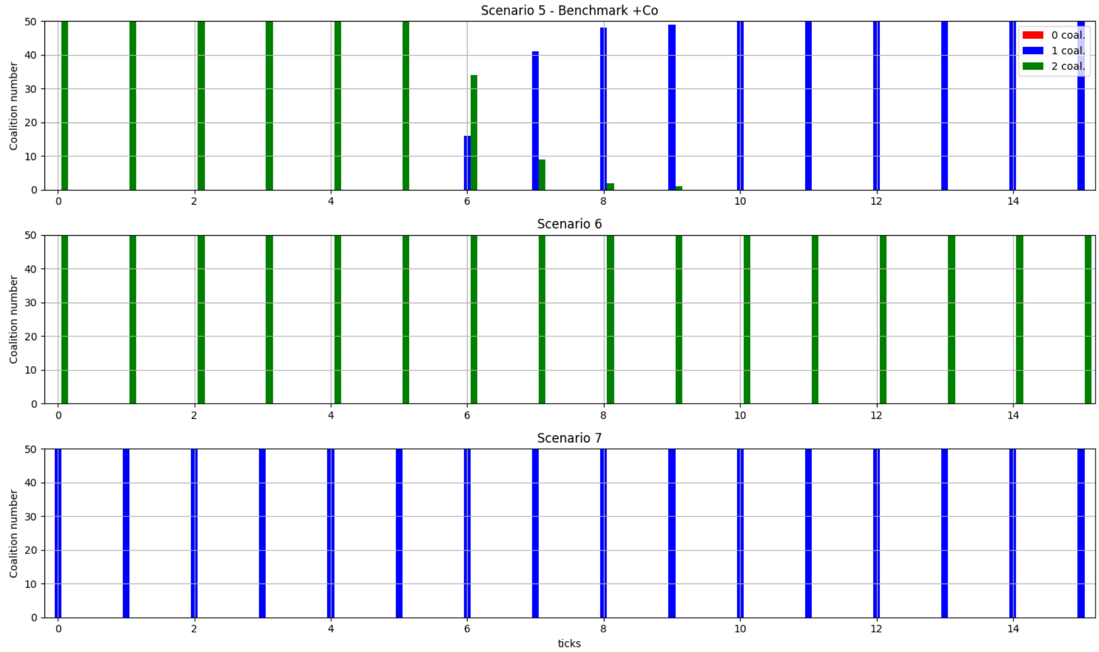

Figure E.15 – Number of coalitions for scenarios 0, 5, 6 and 7 (+Co implementation).

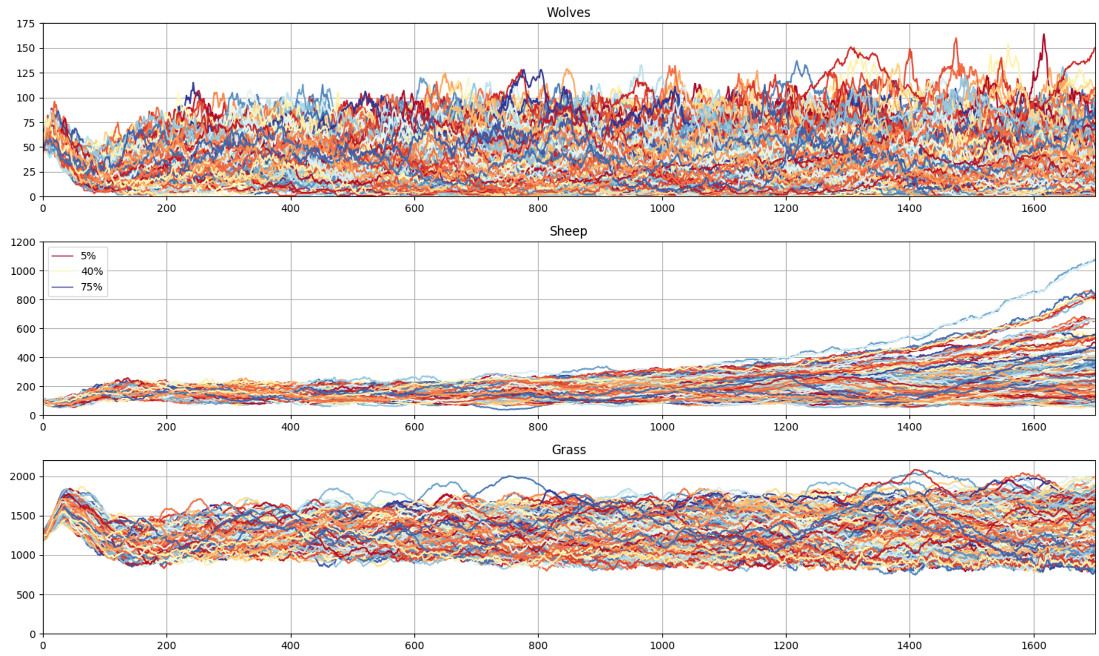

Figure E.16 – Predation model results for wolf, sheep and grass counts for scenario 8 with the colours representing the different levels of the shared resources parameter.

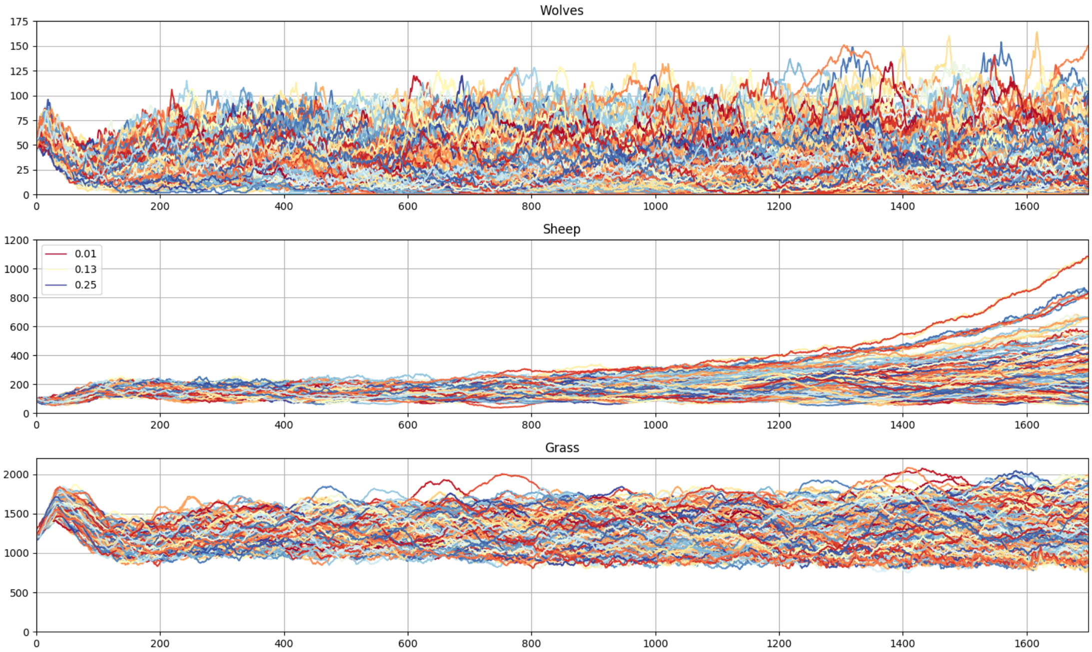

Figure E.17 – Predation model results for wolf, sheep and grass counts for scenarios 8 with the colours representing the different levels of the coalition creation threshold parameter.

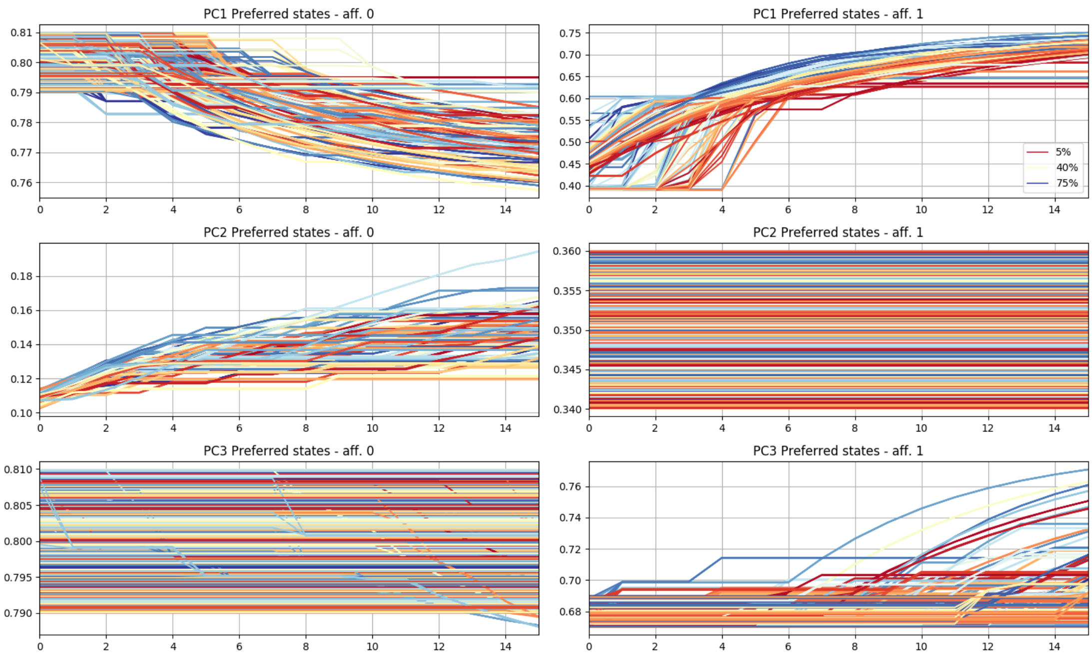

Figure E.18 – Policy core issue states for the scenario 8 with the colours representing the different levels of the shared resource parameter.

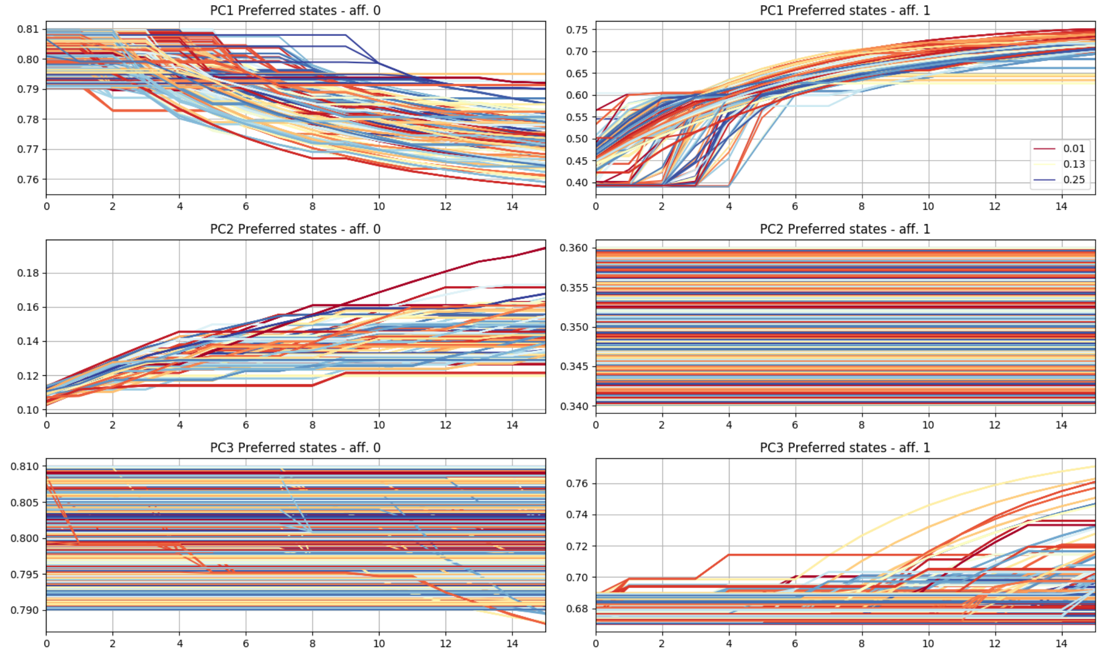

Figure E.19 – Policy core issue states for the scenario 8 with the colours representing the different levels of the coalition creation threshold parameter.

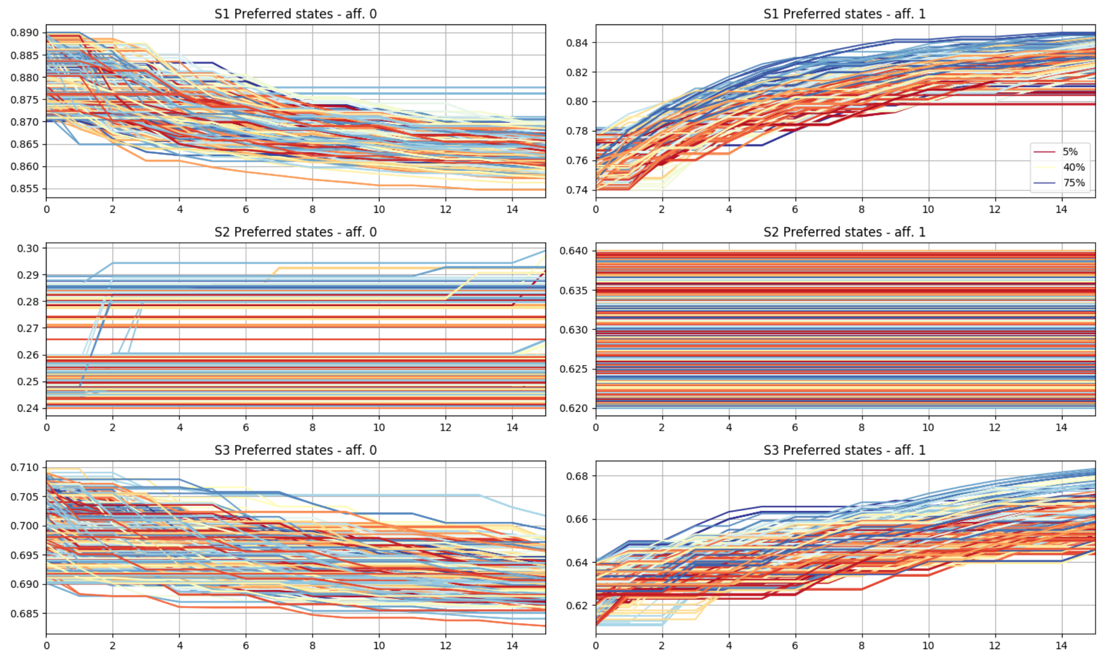

Figure E.20 – Secondary issue states for the scenario 8 with the colours representing the different levels of pooled resources from the agents joining the coalitions.

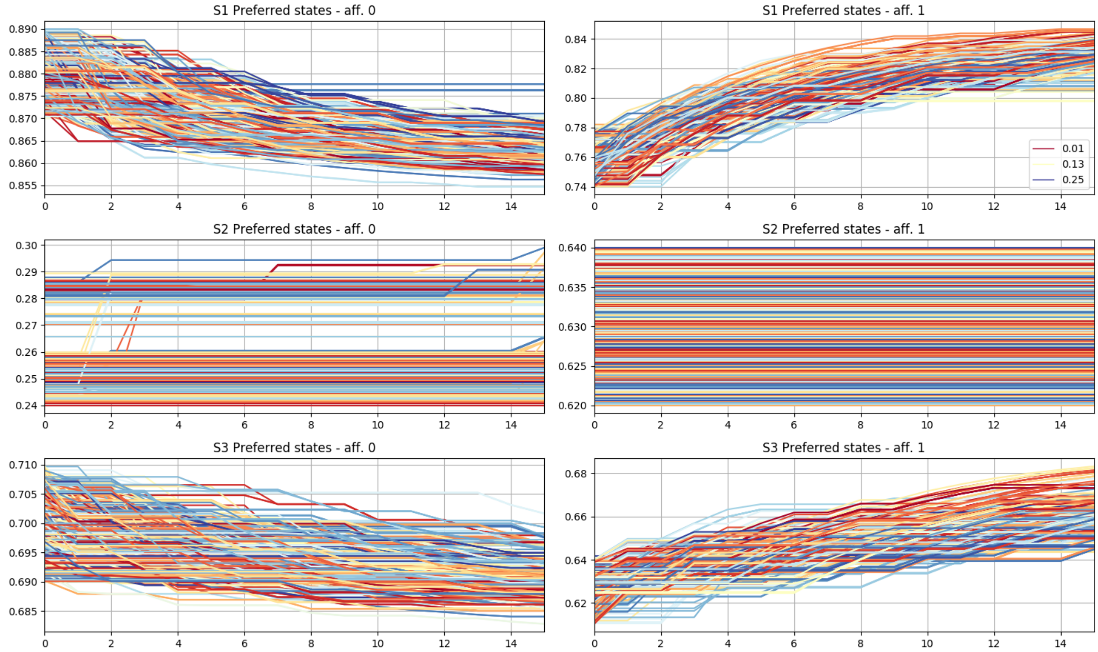

Figure E.21 – Secondary issue states for the scenario 8 with the colours representing the different levels of thresholds required to join coalitions.

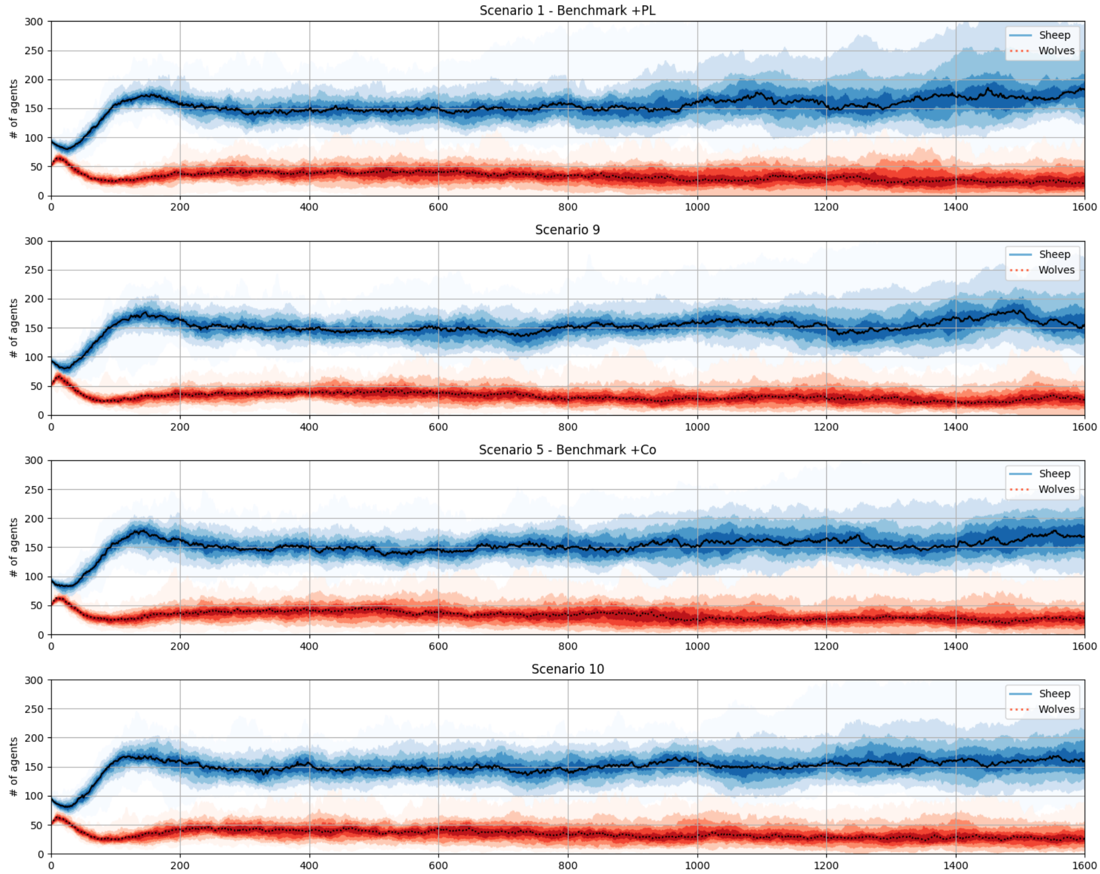

Figure E.22 – Predation model results for wolf and sheep counts for scenarios 0, 5, 9 and 10 (+PK implementation).

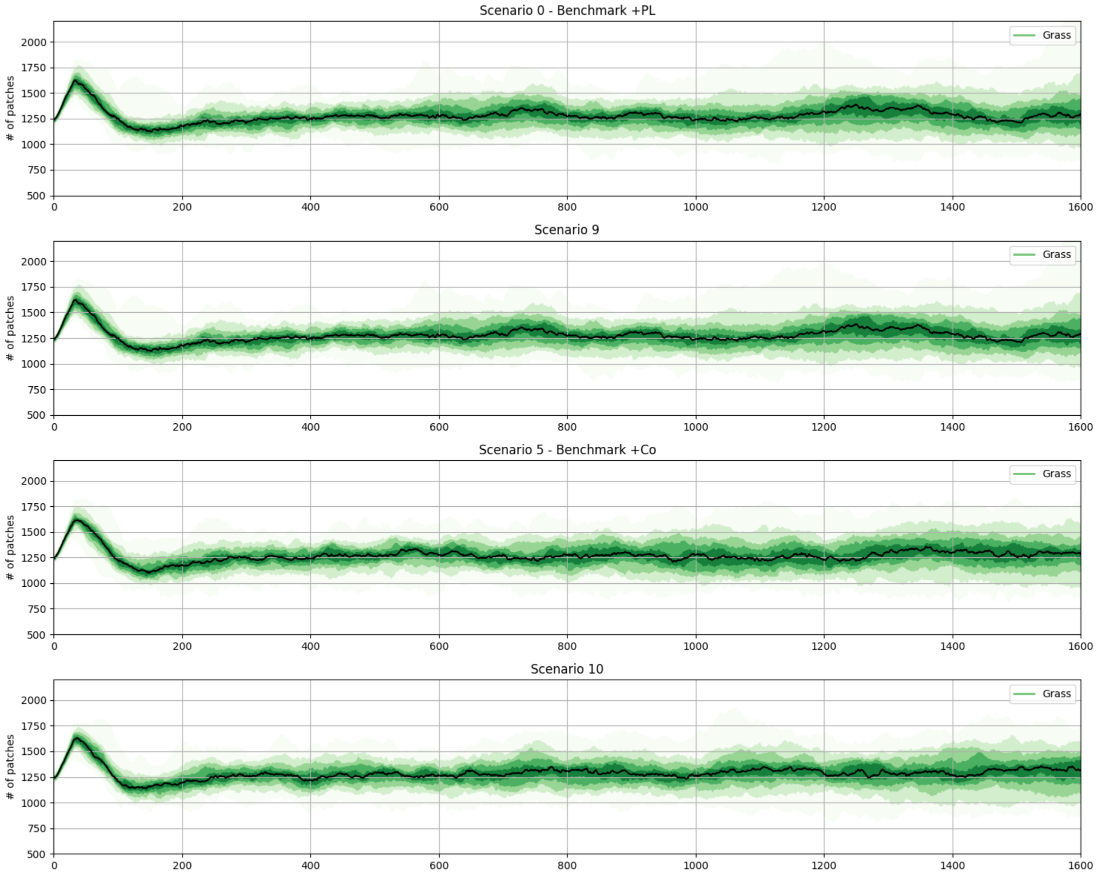

Figure E.23 – Predation model results for grass patch count for scenarios 0, 5, 9 and 10 (+PK implementation).

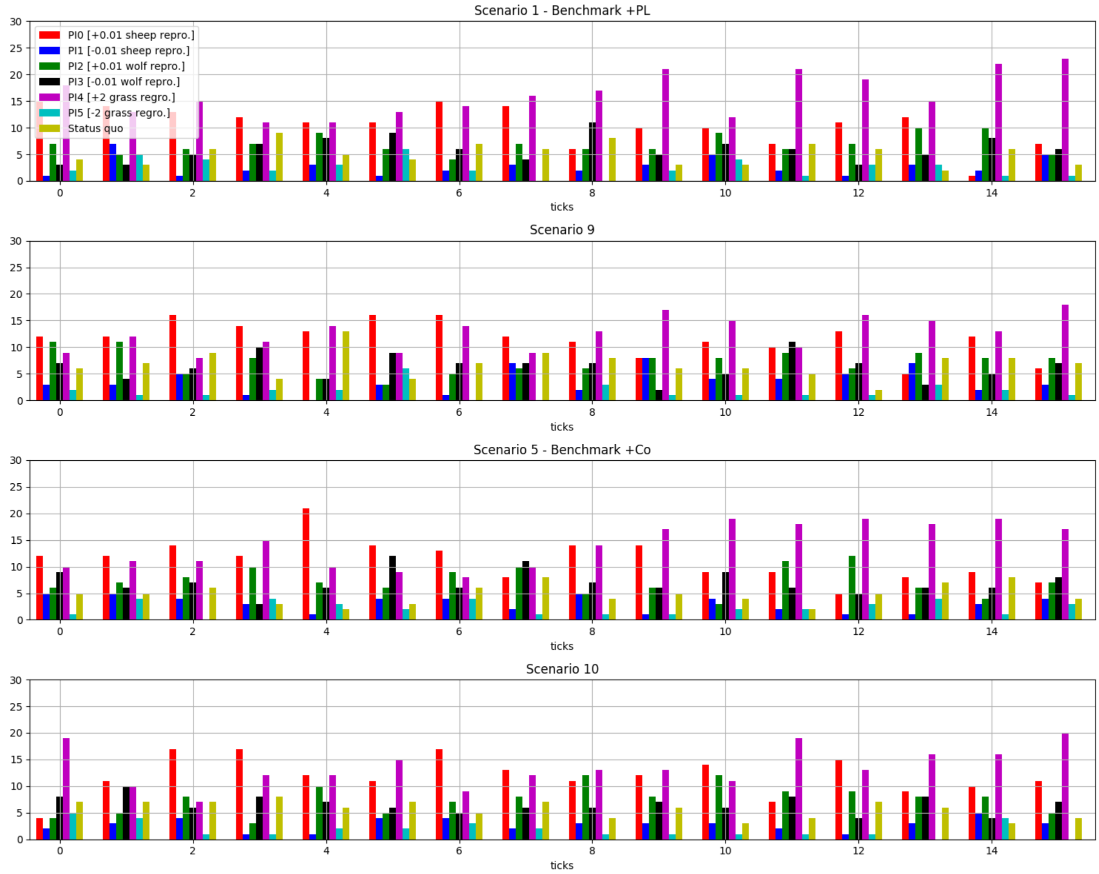

Figure E.24 – Policy instruments selected for scenarios 0, 5, 9 and 10 (+PK implementation).

Figure E.25 – Policy core issues selected per affiliation for scenarios 0, 5, 9 and 10 (+PK implementation).

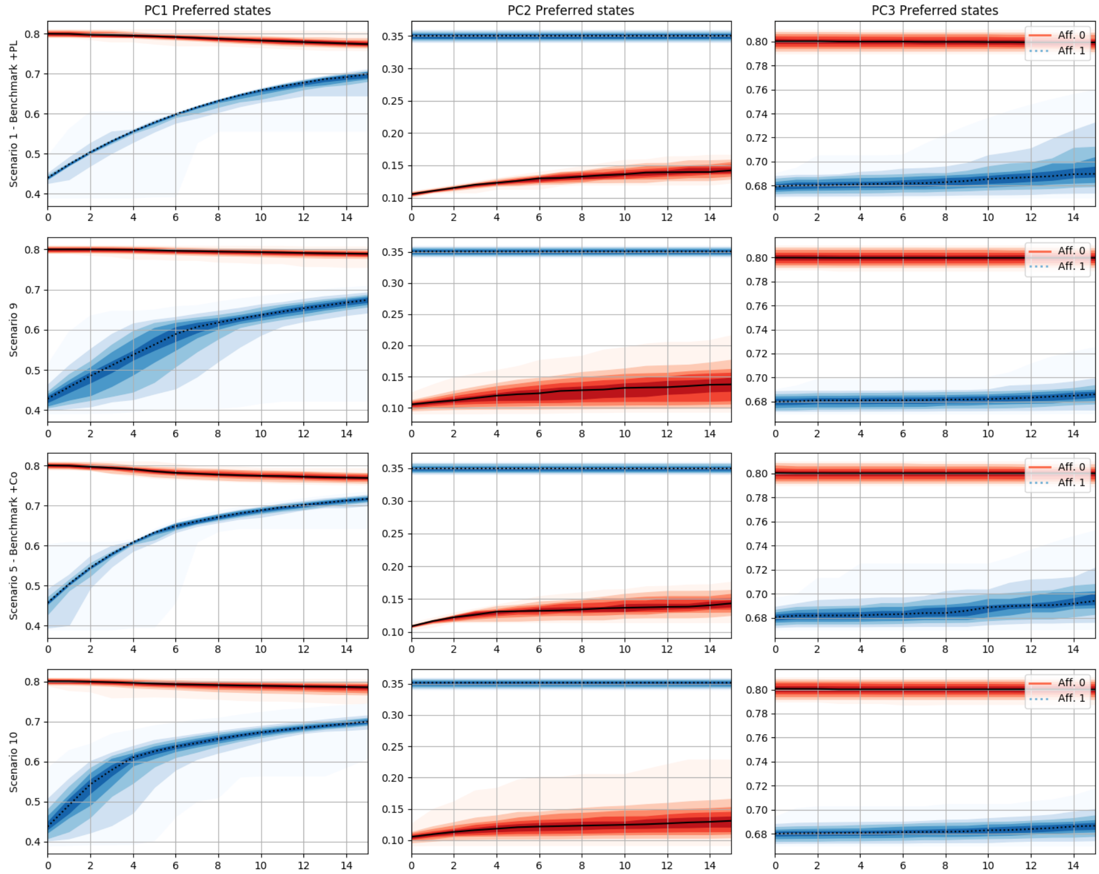

Figure E.26 – Policy core issue preferred states for scenarios 0, 5, 9 and 10 (+PK implementation).

Figure E.27 – Secondary issue selected per affiliation for scenarios 0, 5, 9 and 10 (+PK implementation).

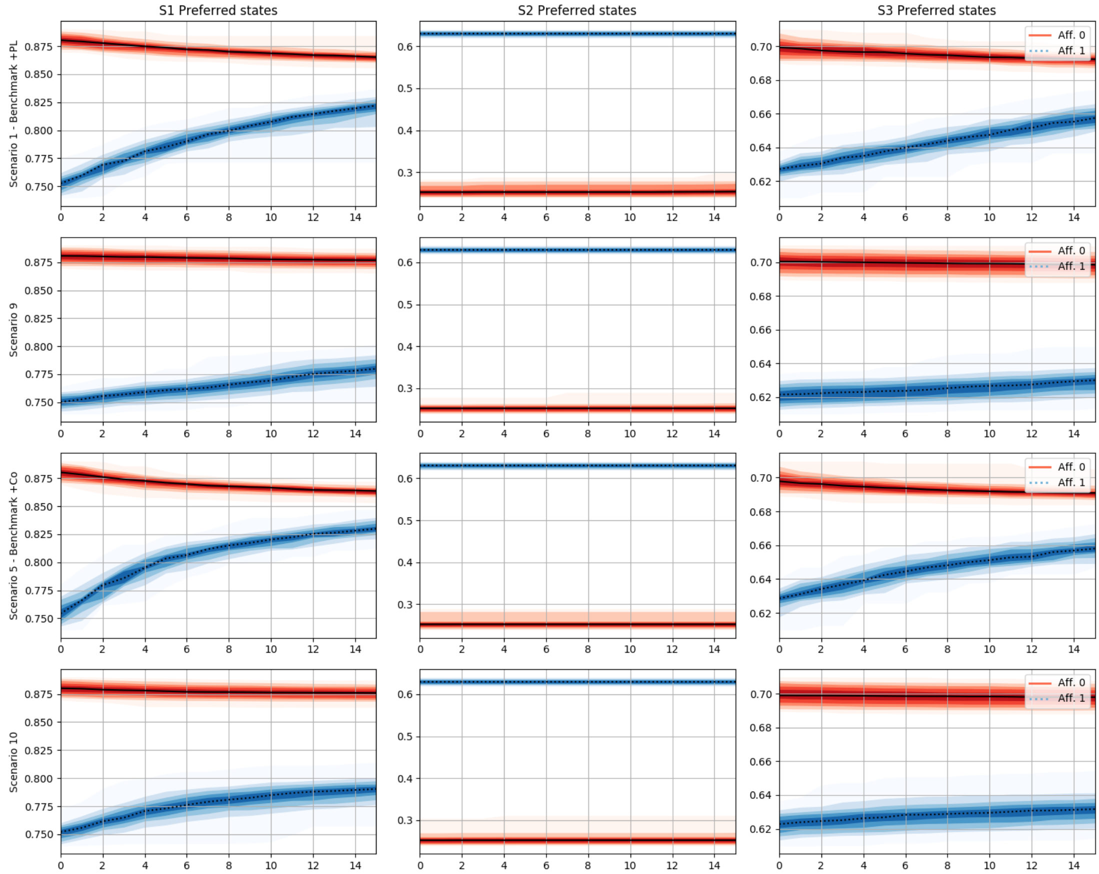

Figure E.28 – Secondary issue preferred states for scenarios 0, 5, 9 and 10 (+PK implementation).

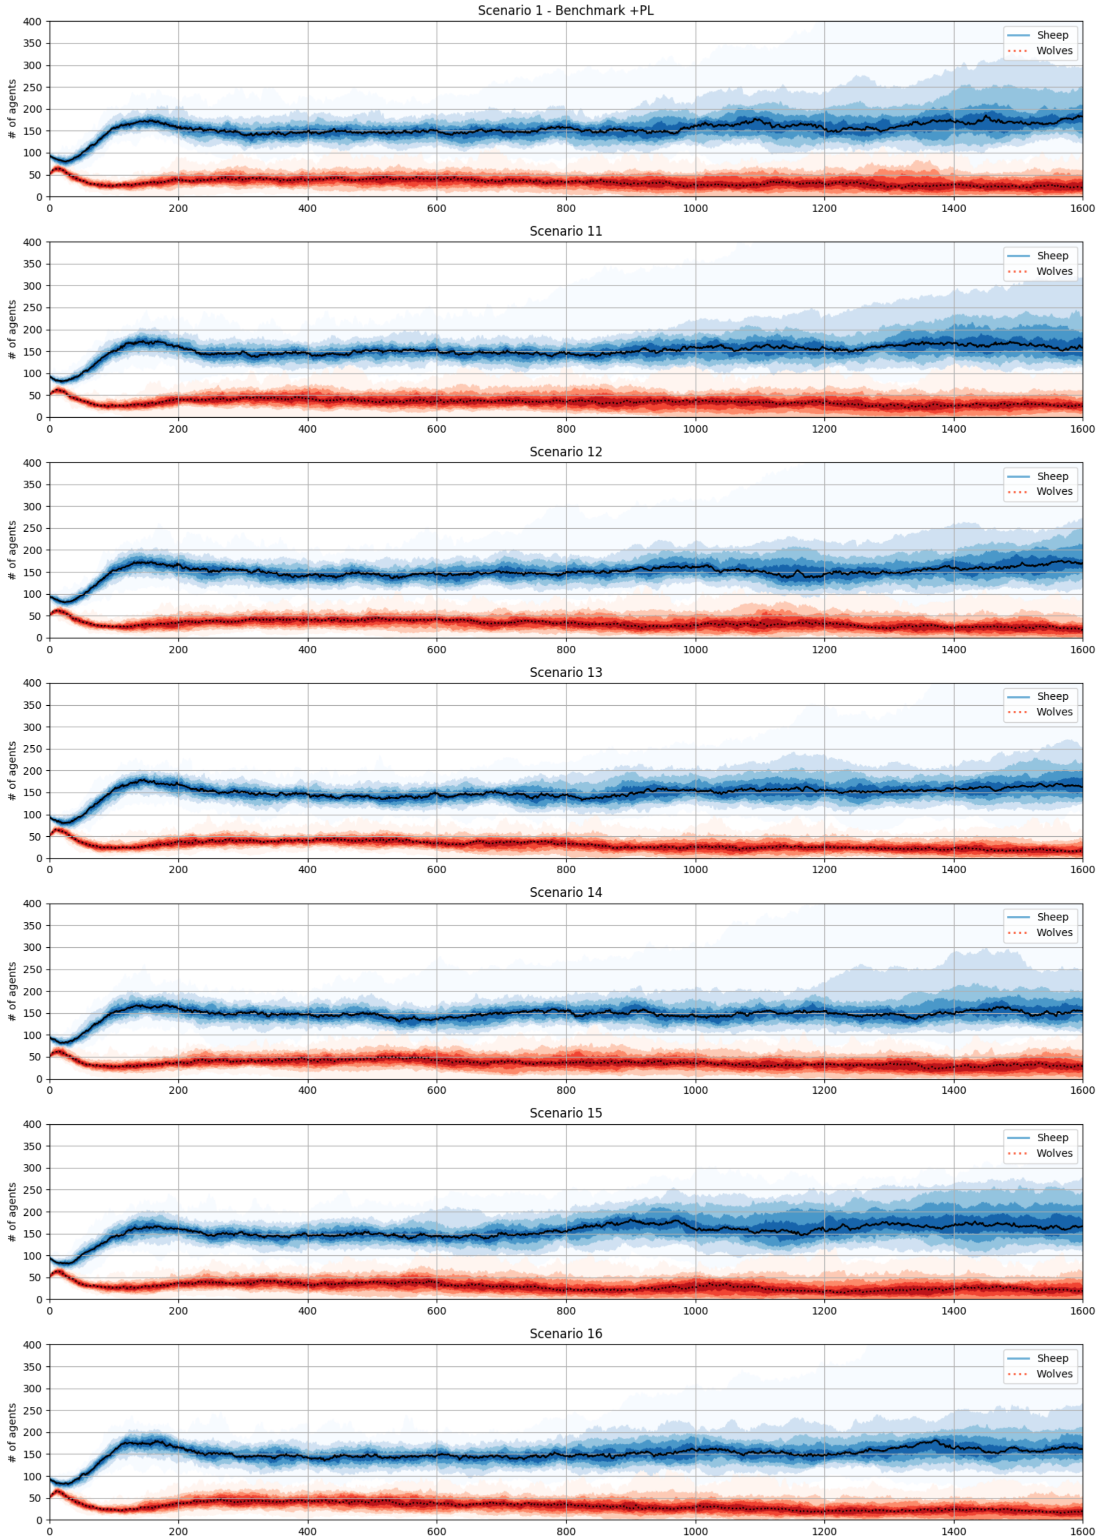

Figure E.29 – Predation model results for wolf and sheep counts for scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation).

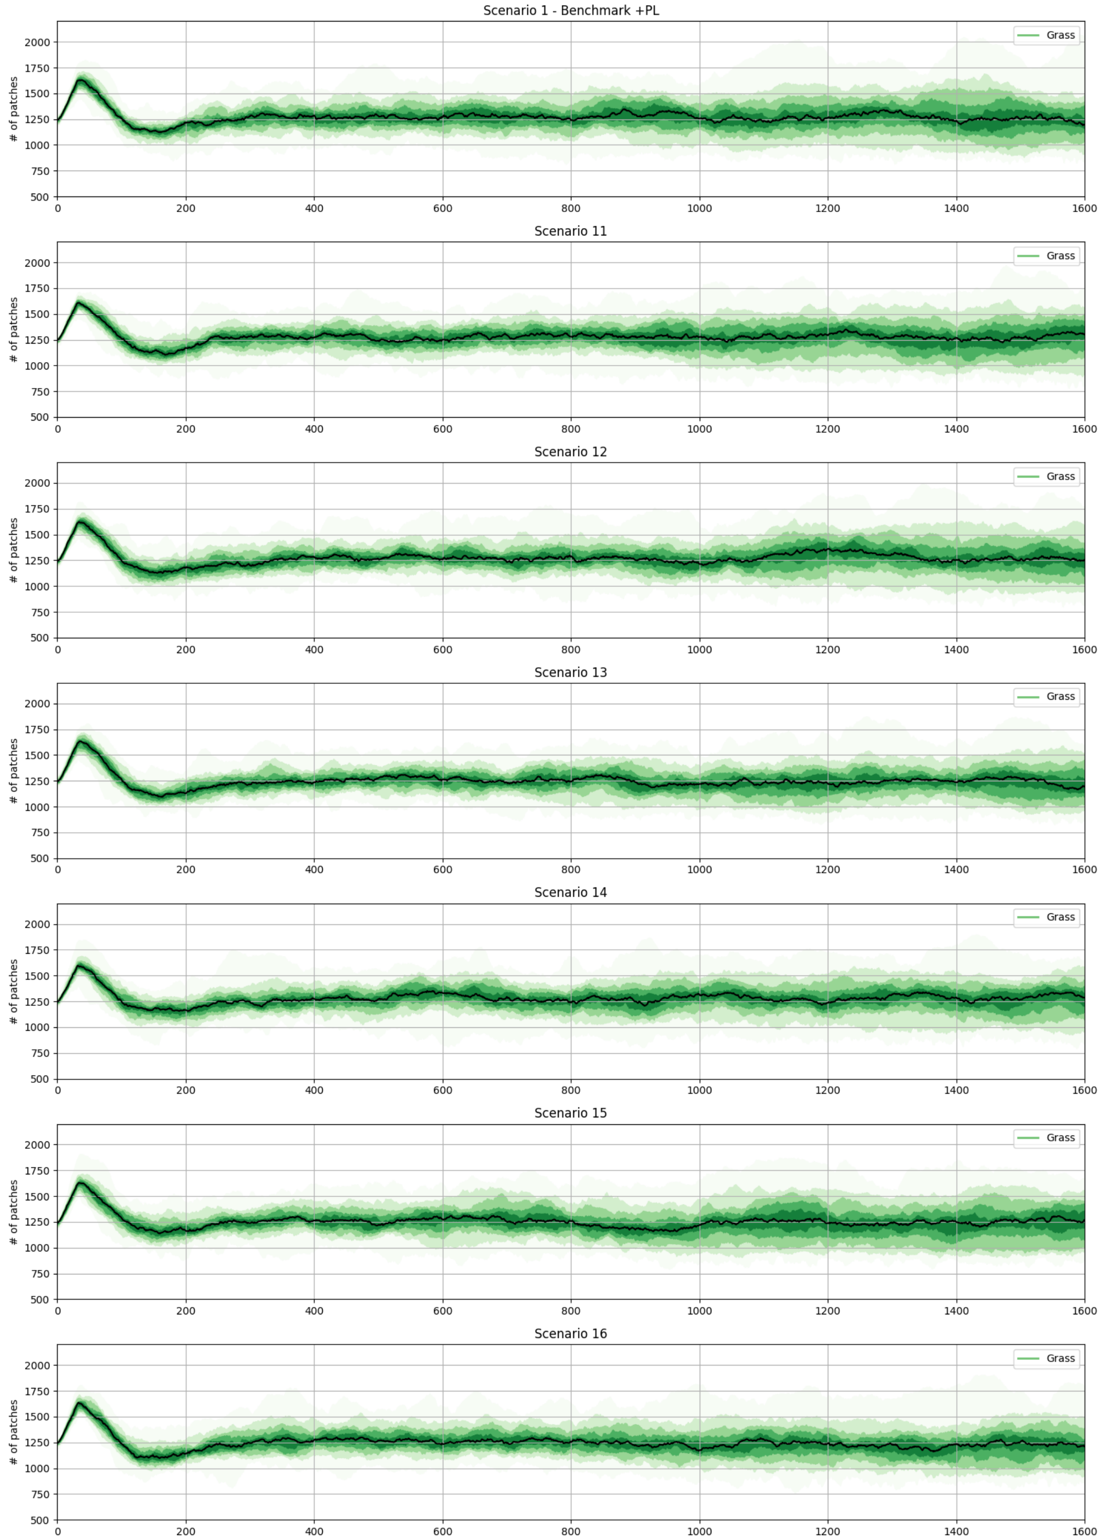

Figure E.30 – Predation model results for grass patch count for scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation).

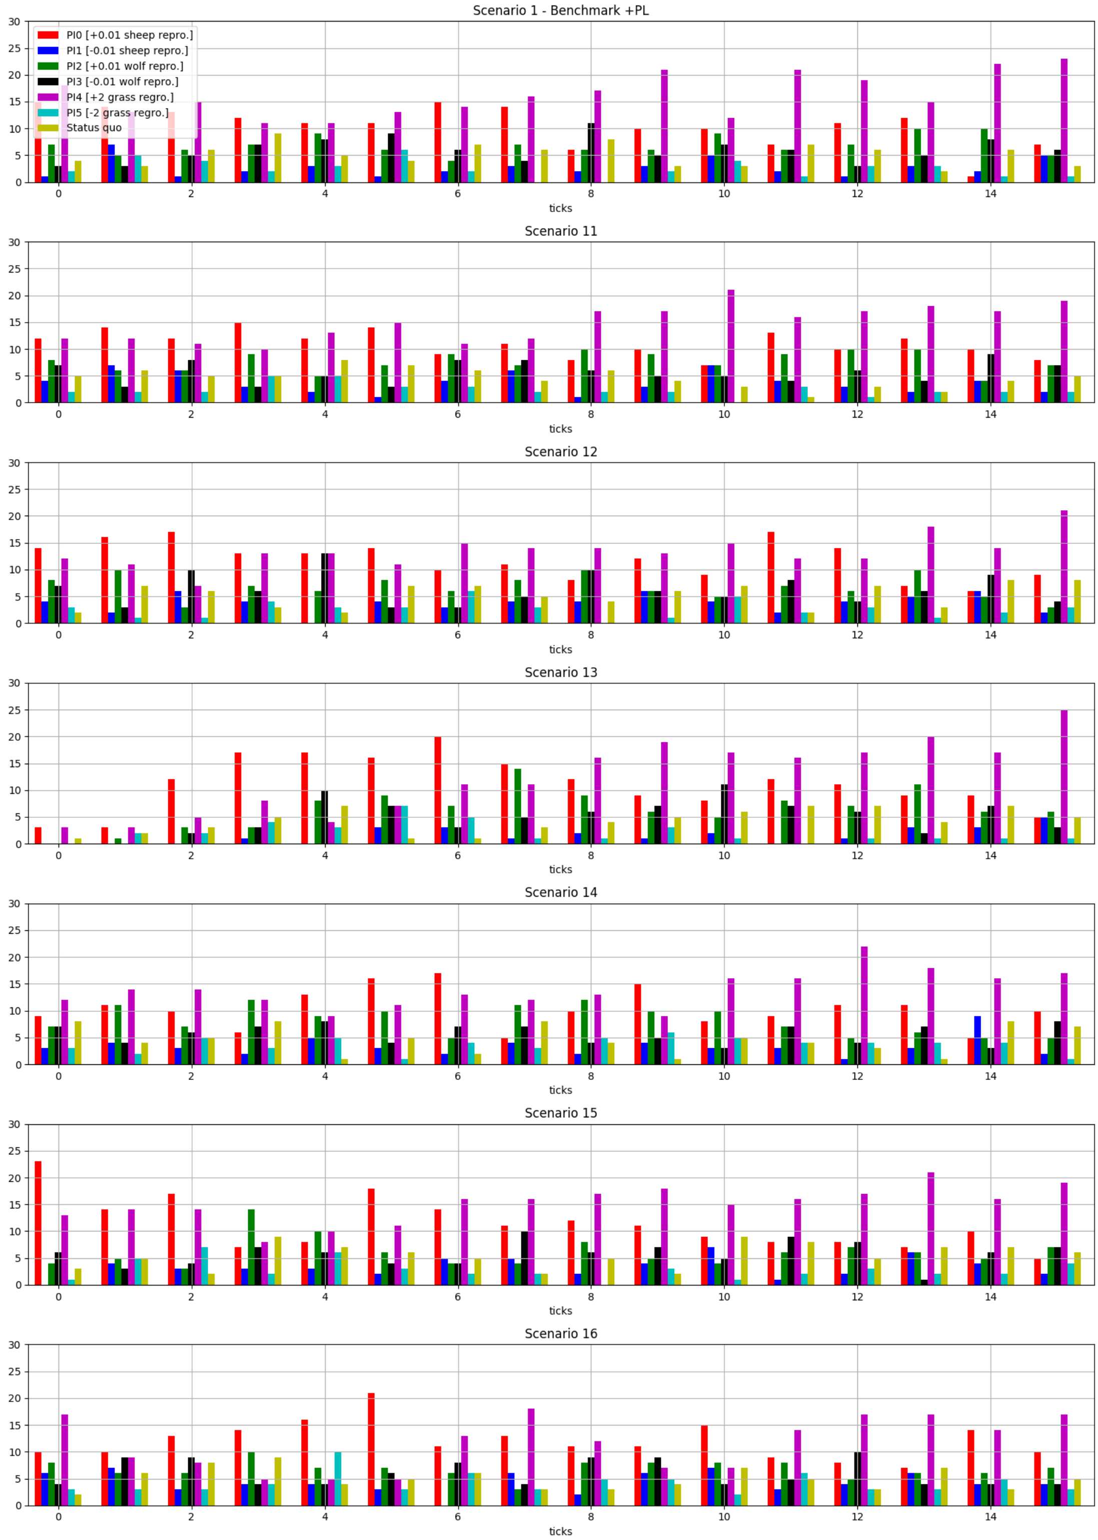

Figure E.31 – Policy instruments selected for scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation).

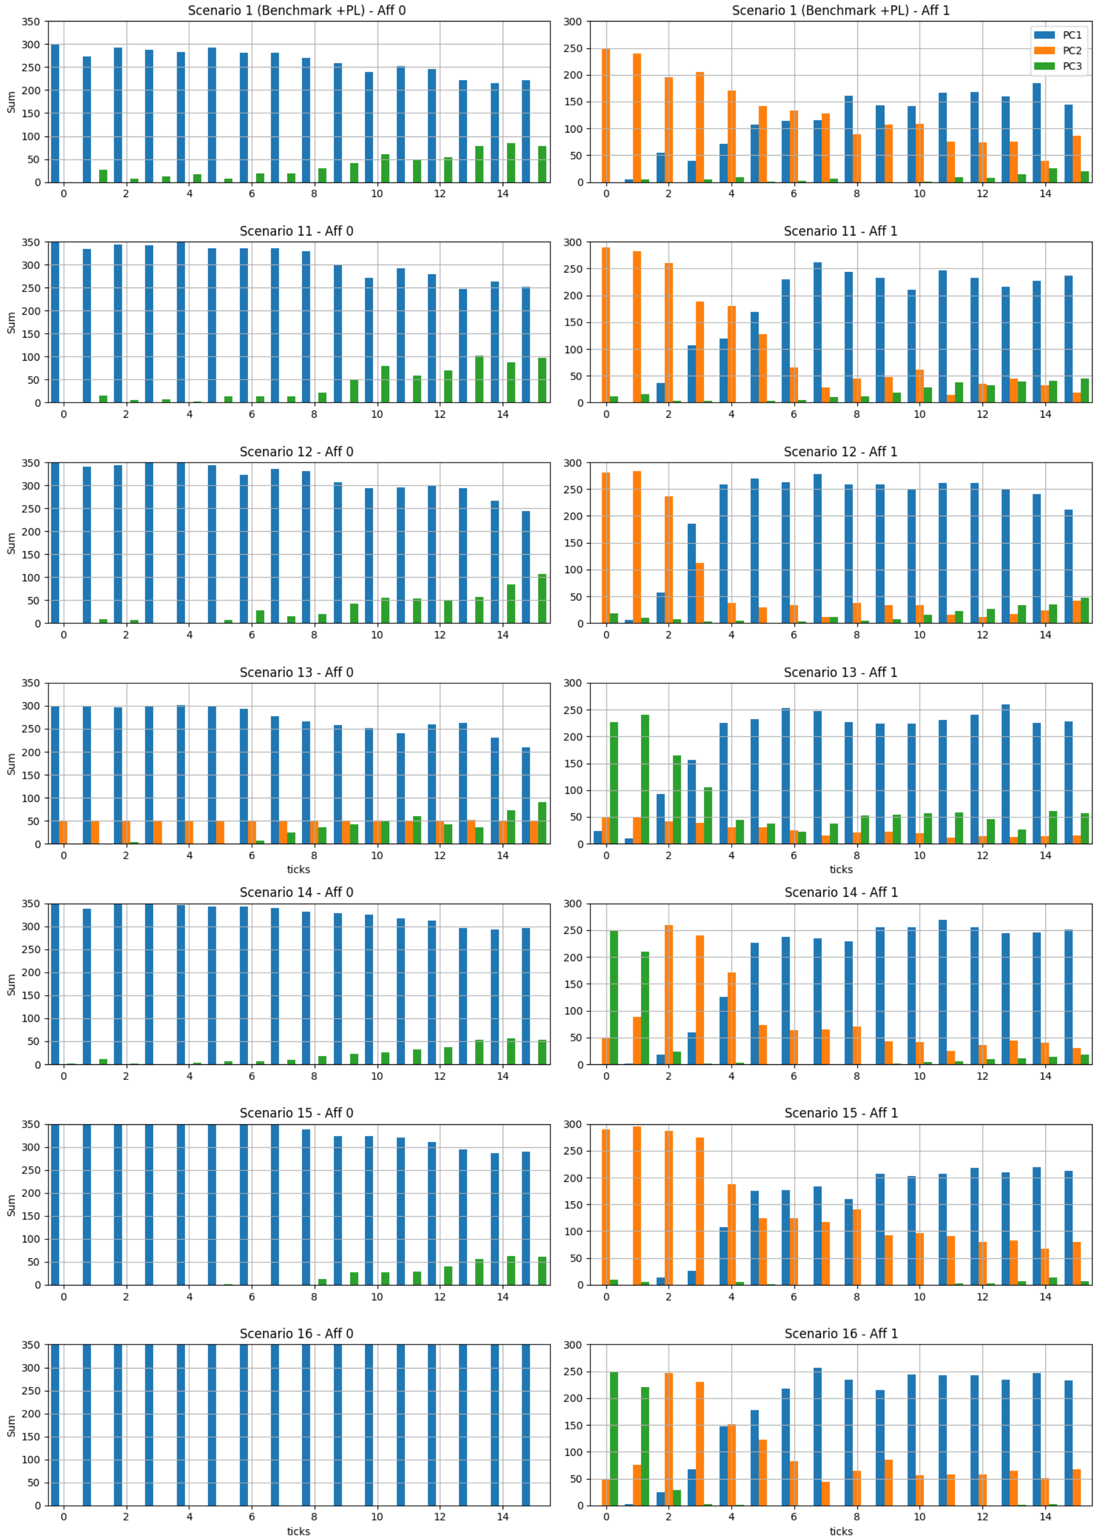

Figure E.32 – Policy core issues selected per affiliation for scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation).

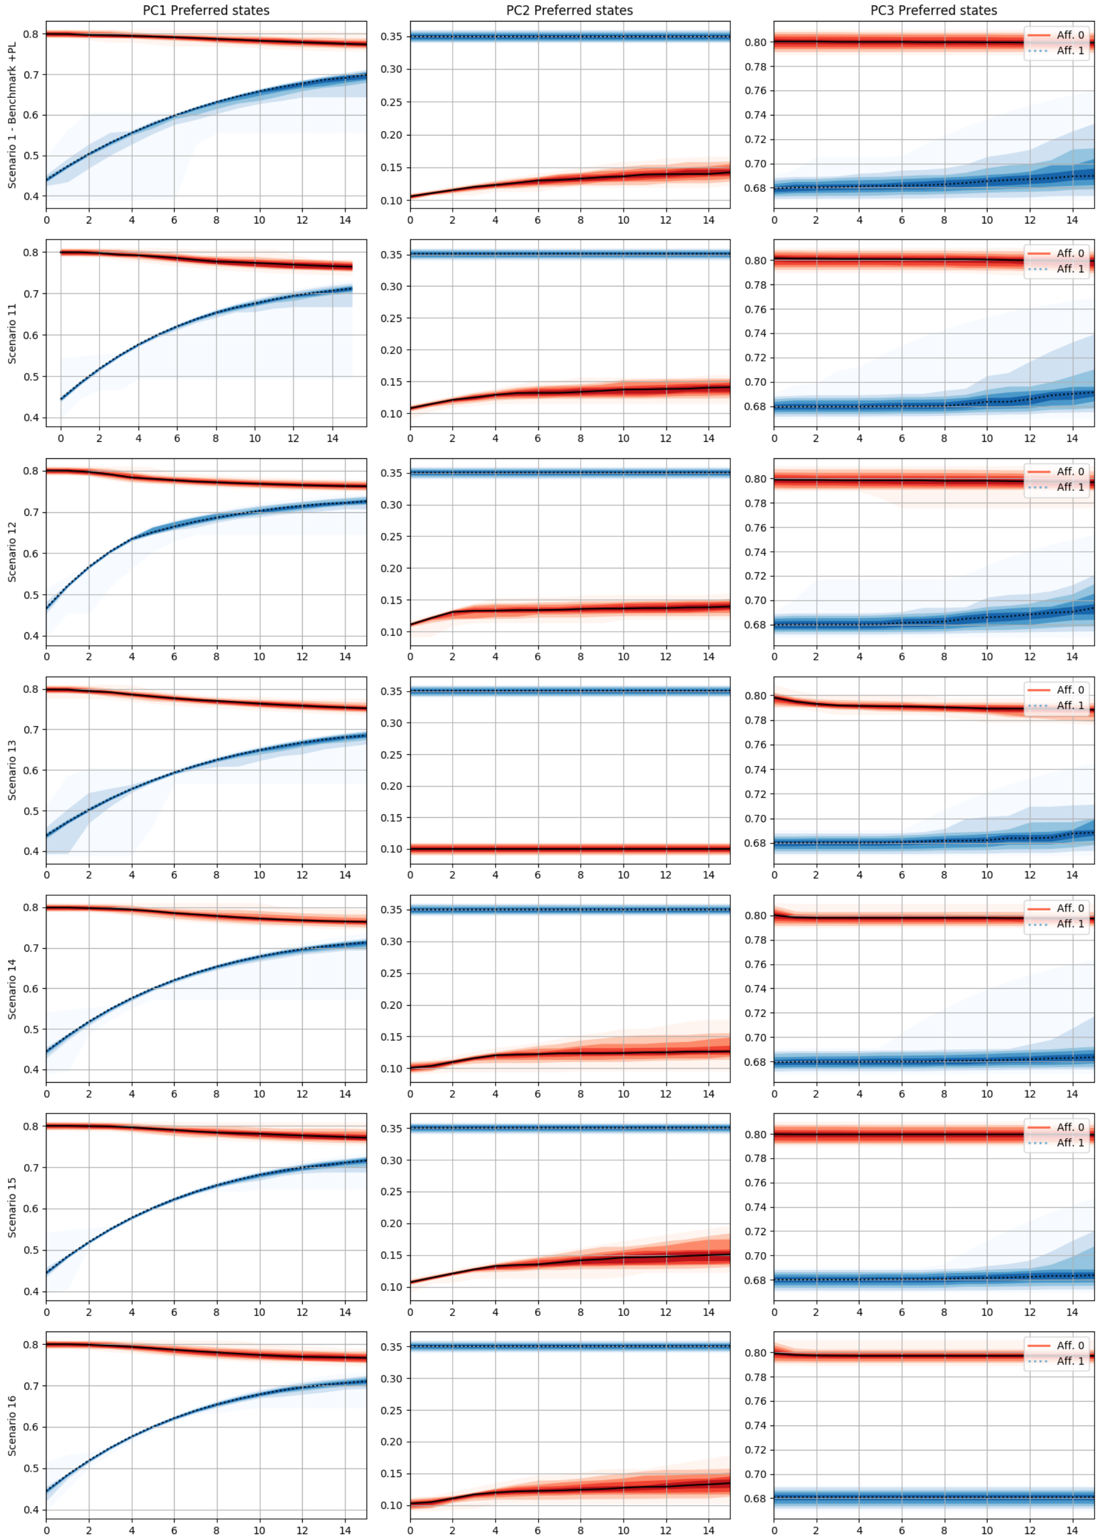

Figure E.33 – Policy core issue preferred states scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation).

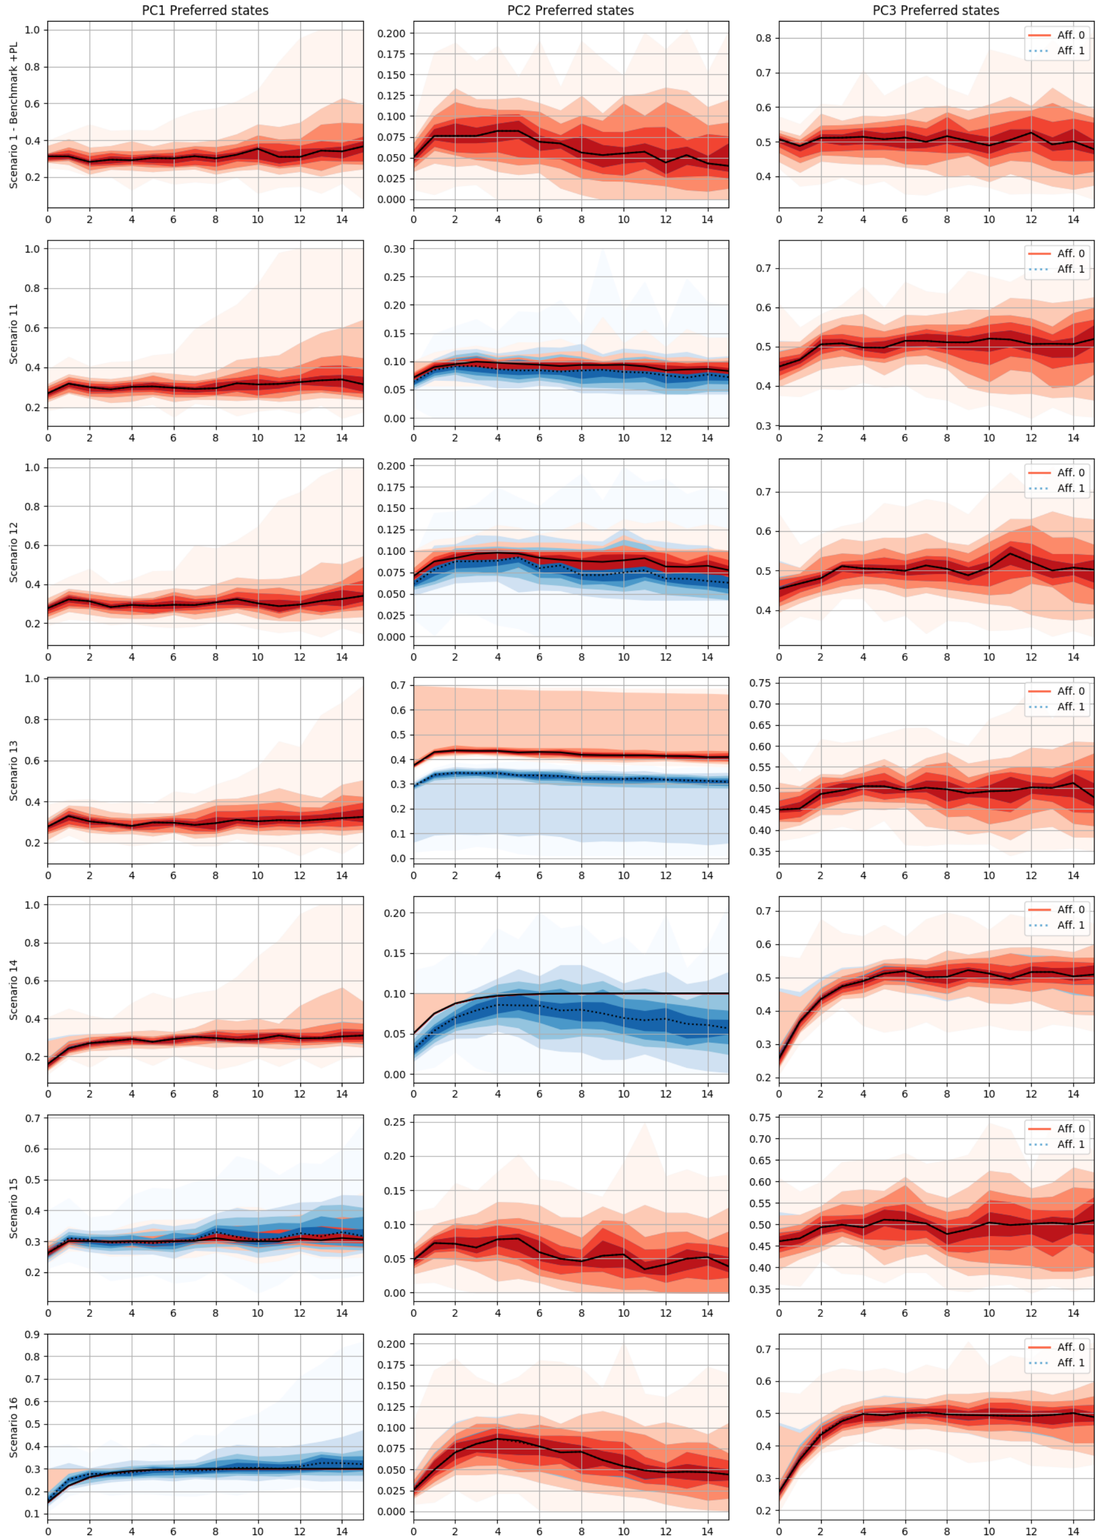

Figure E.34 – Policy core issue actual beliefs scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation).

Figure E.35 – Secondary issue selected per affiliation for scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation).

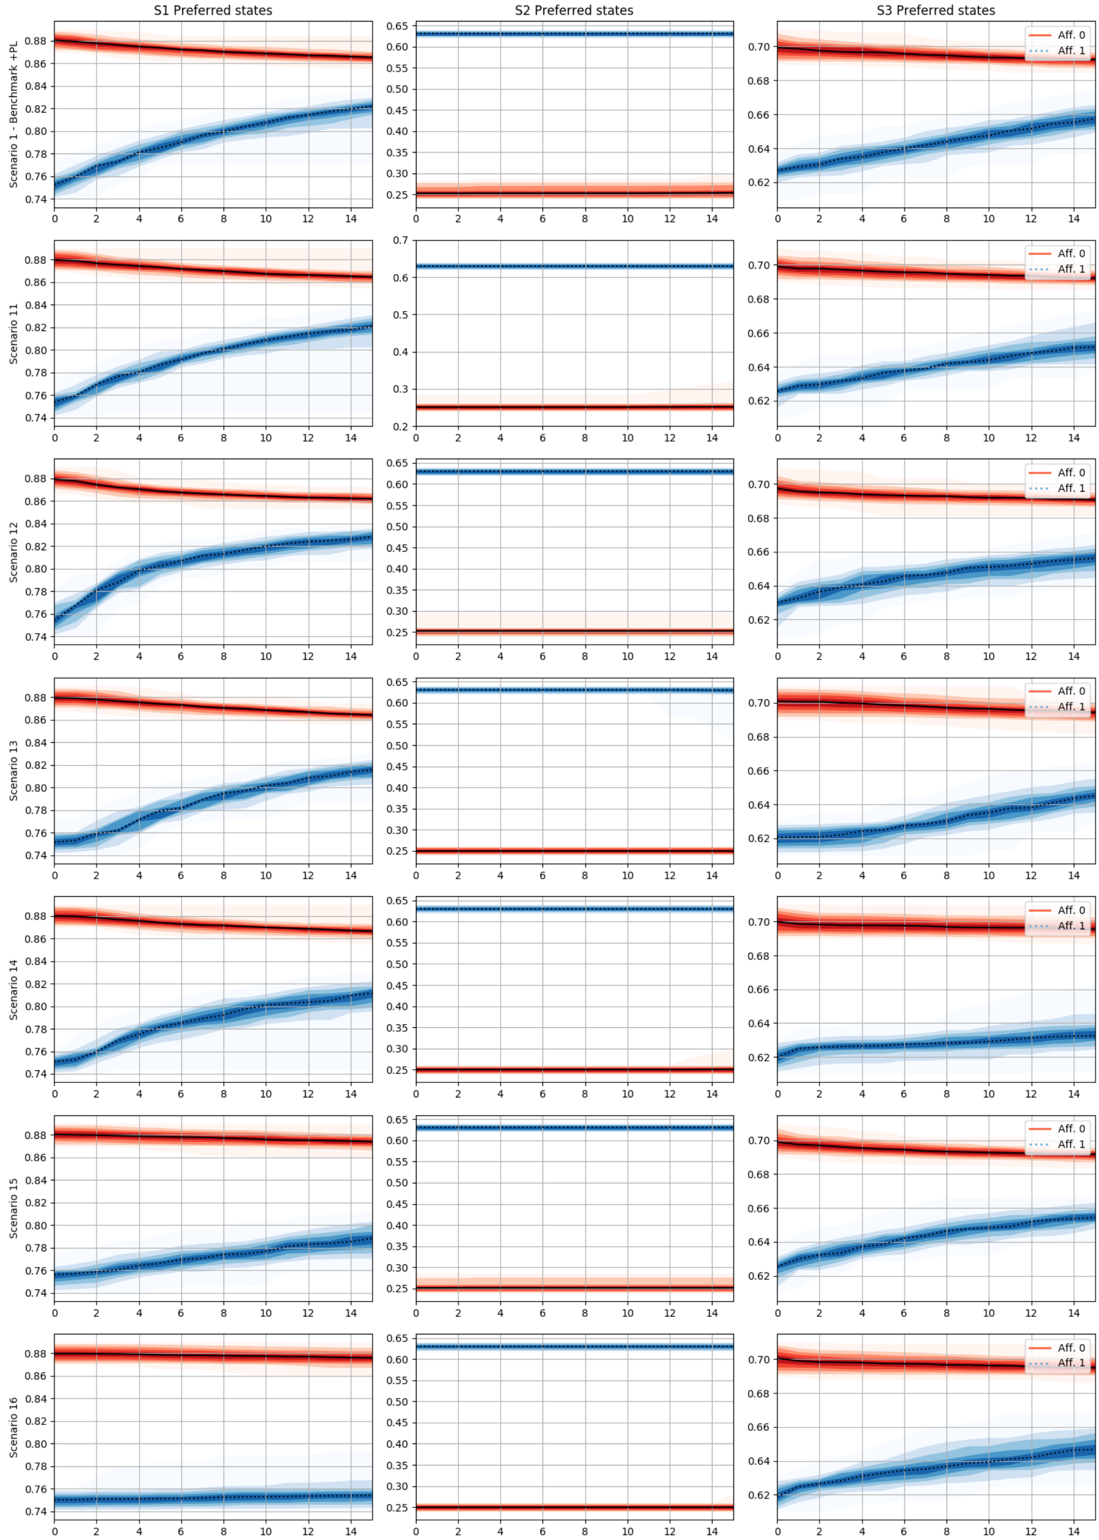

Figure E.36 – Secondary issue preferred states scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation).

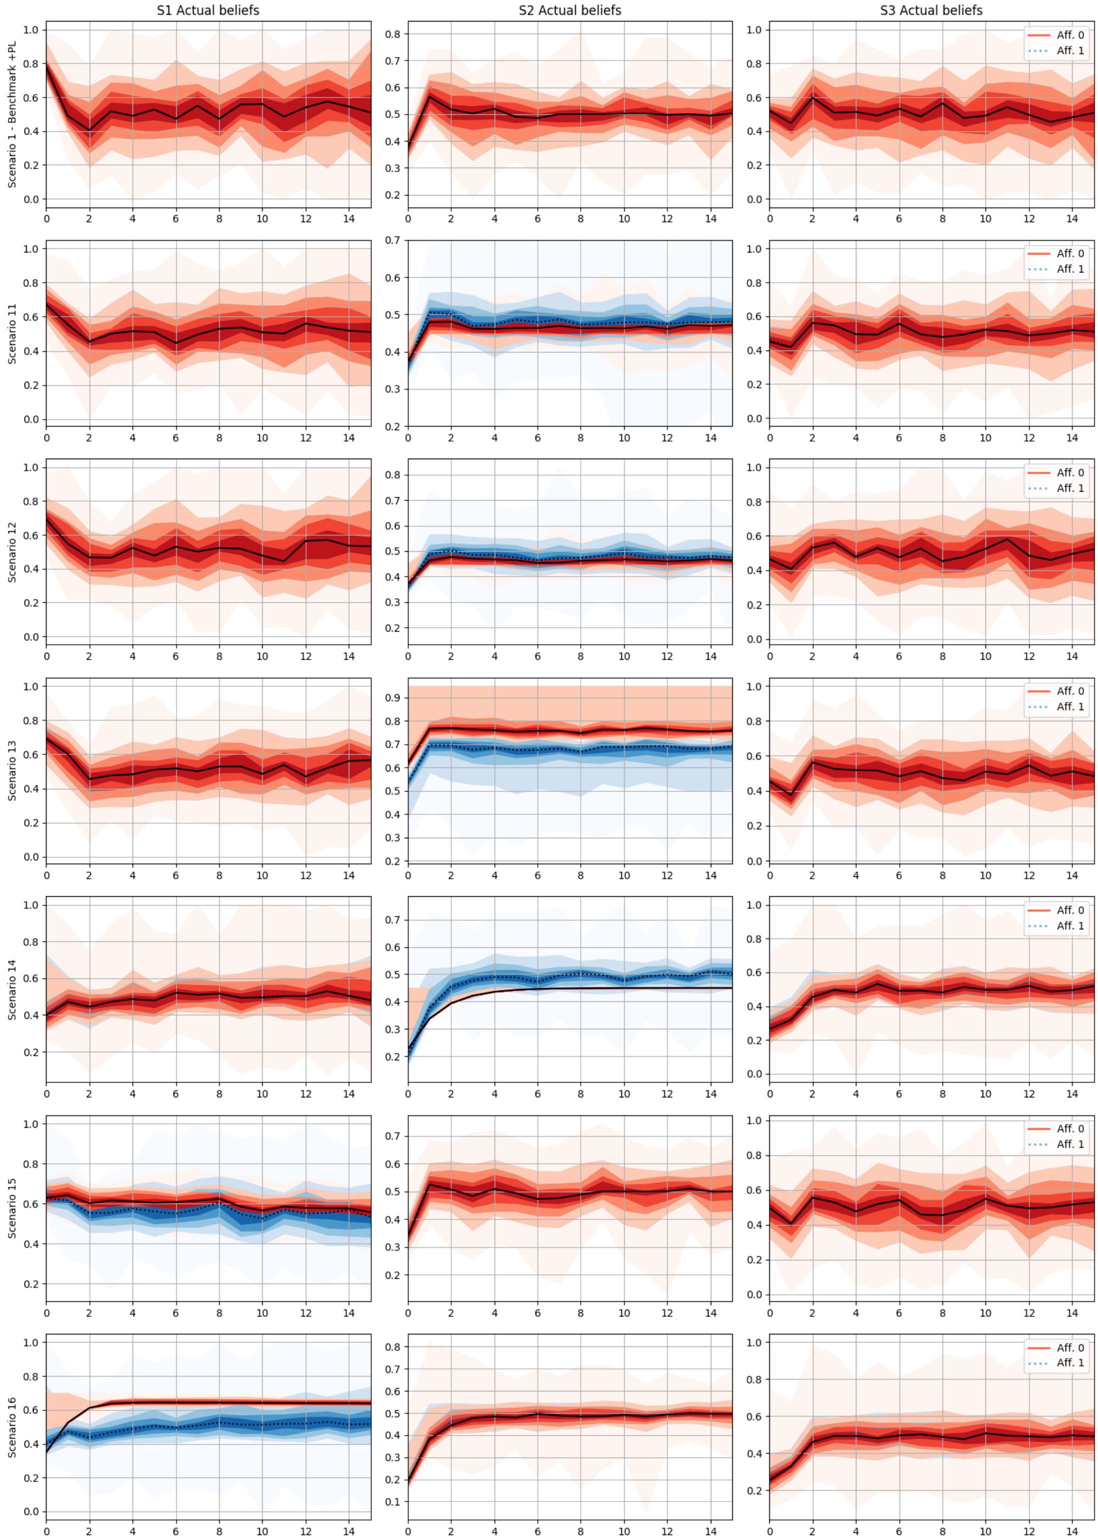

Figure E.37 – Secondary issue actual beliefs scenarios 0, 11, 12, 13, 14, 15 and 16 (+PI implementation).

F The ACF implementation electricity model results

This appendix collects all of the graphs that were plotted using the results from the ACF implementation model with the Swiss electricity market model. The graphs show the results for all of the scenarios in a .PNG compressed format. The results can also be found in higher quality on GitHub[1].

[1] https://github.com/kleinrap/policyemergence_A/tree/master/2_ElectricityModel/0_ResultsGraphs

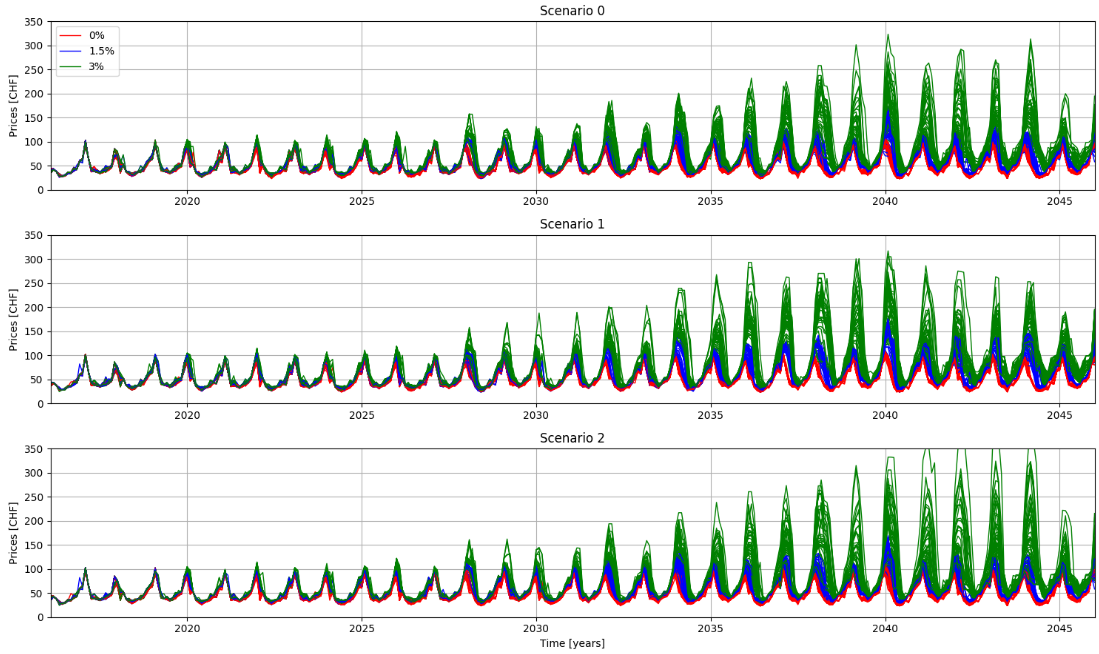

Figure F.1 – Monthly averaged electricity prices, split according to the three electorate influence rates and for the three demand growth scenarios.

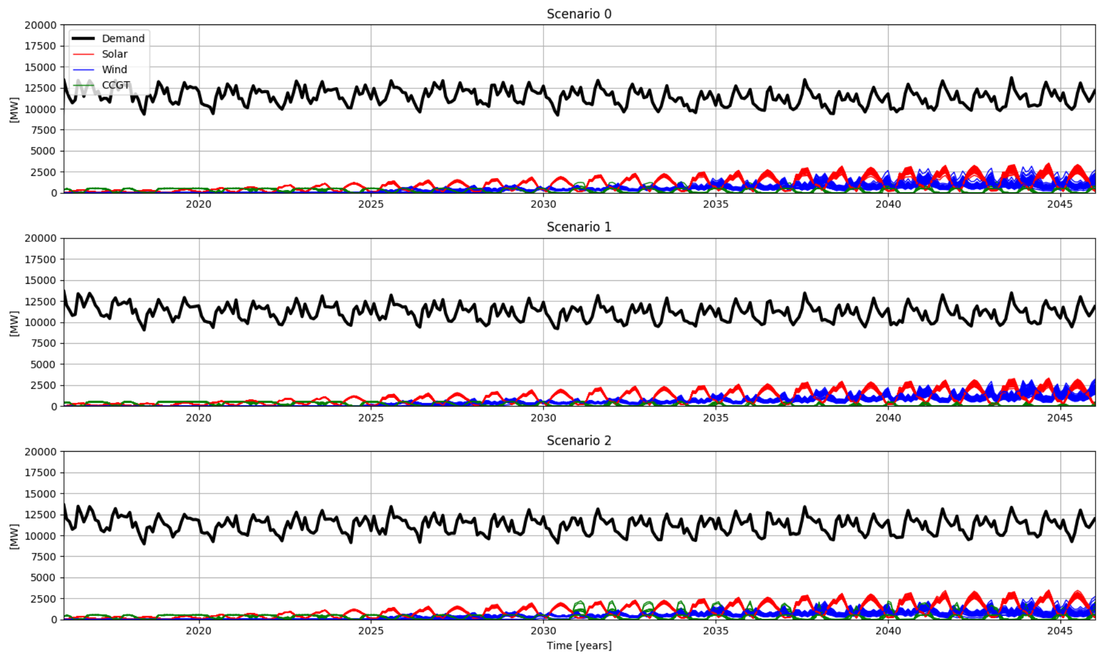

Figure F.2 – Overall demand with the electricity supplied by solar, wind and CCGT sources for all policy scenarios and for the 0% demand growth scenario.

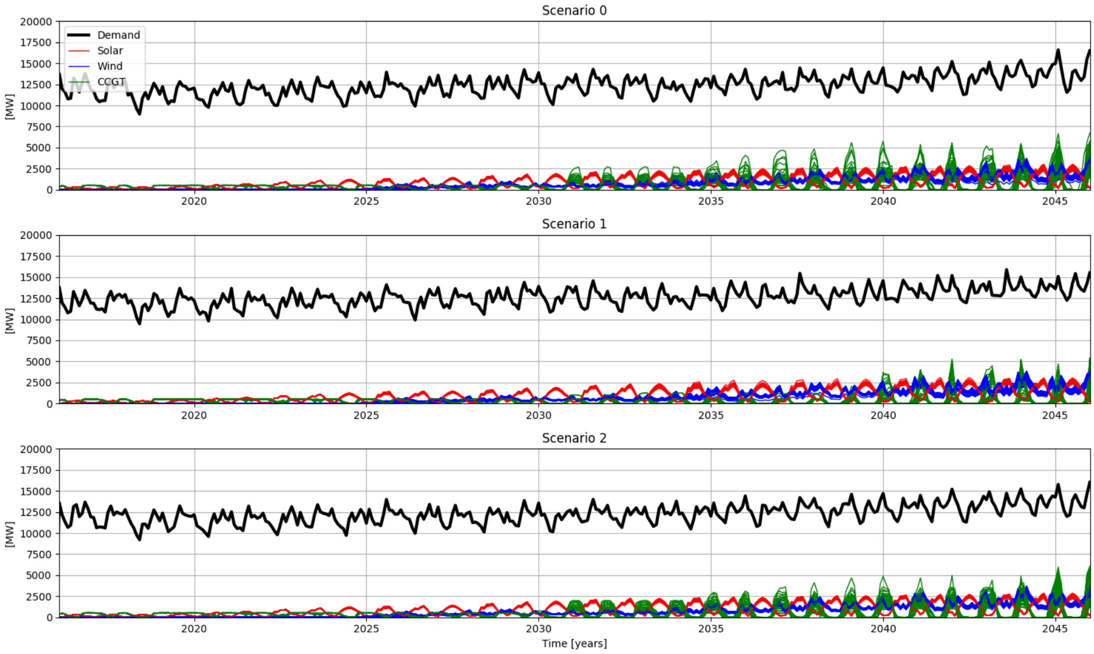

Figure F.3 – Overall demand with the electricity supplied by solar, wind and CCGT sources for all policy scenarios and for the 1.5% demand growth scenario.

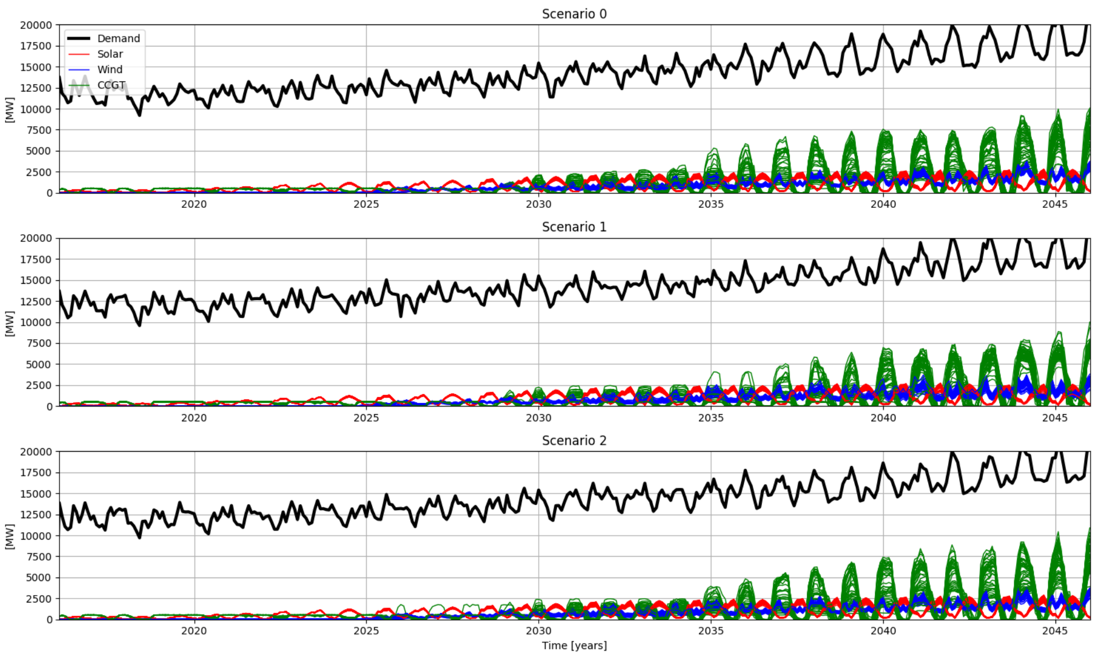

Figure F.4 – Overall demand with the electricity supplied by solar, wind and CCGT sources for all policy scenarios and for the 3% demand growth scenario.

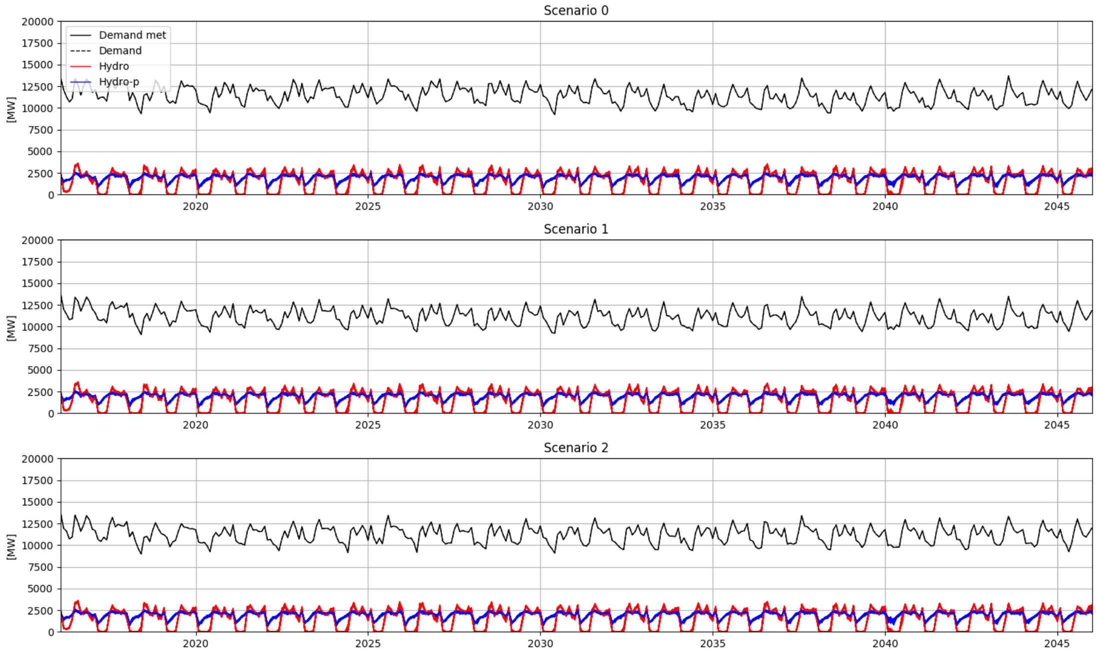

Figure F.5 – Overall demand with the hydro electricity supplied for all policy scenarios and for the 0% demand growth scenario.

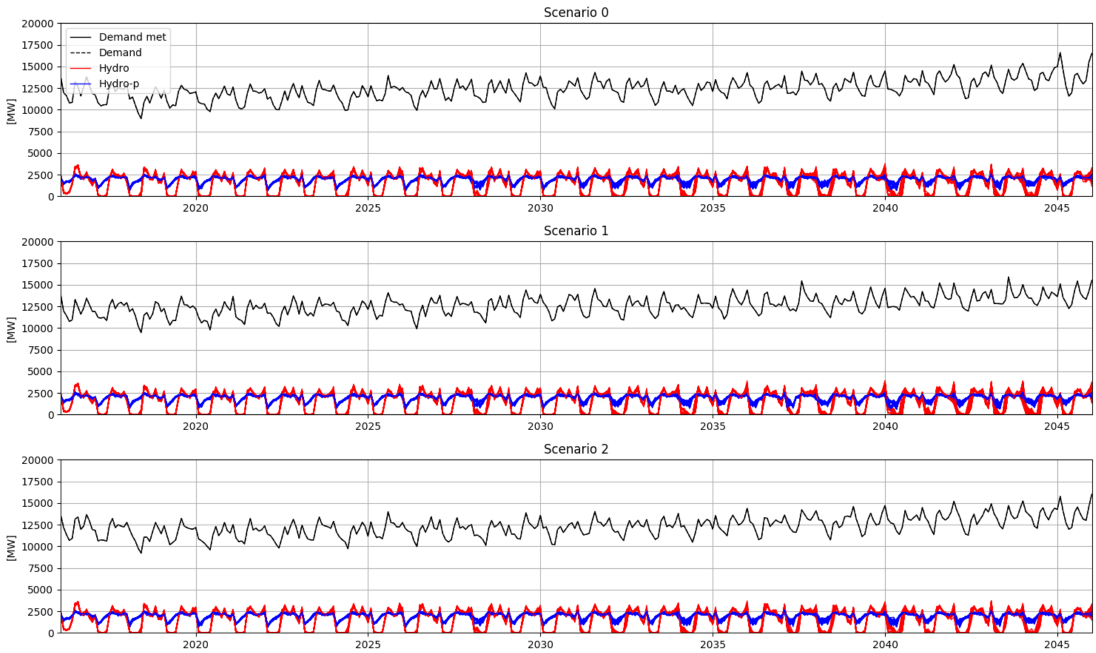

Figure F.6 – Overall demand with the hydro electricity supplied for all policy scenarios and for the 1.5% demand growth scenario.

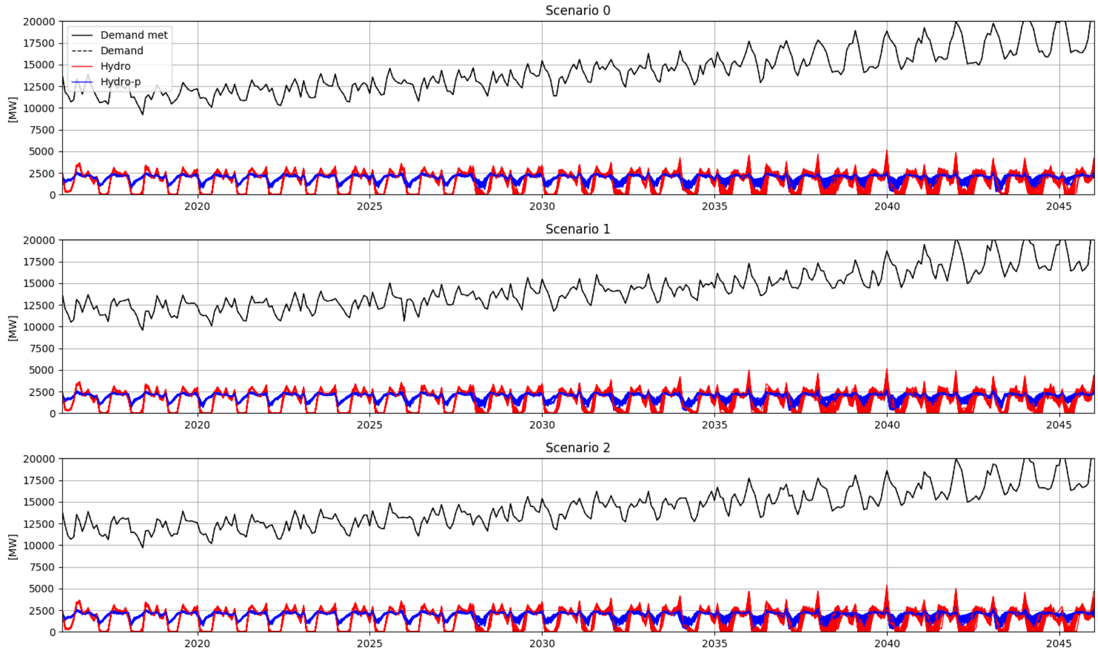

Figure F.7 – Overall demand with the hydro electricity supplied for all policy scenarios and for the 3% demand growth scenario.

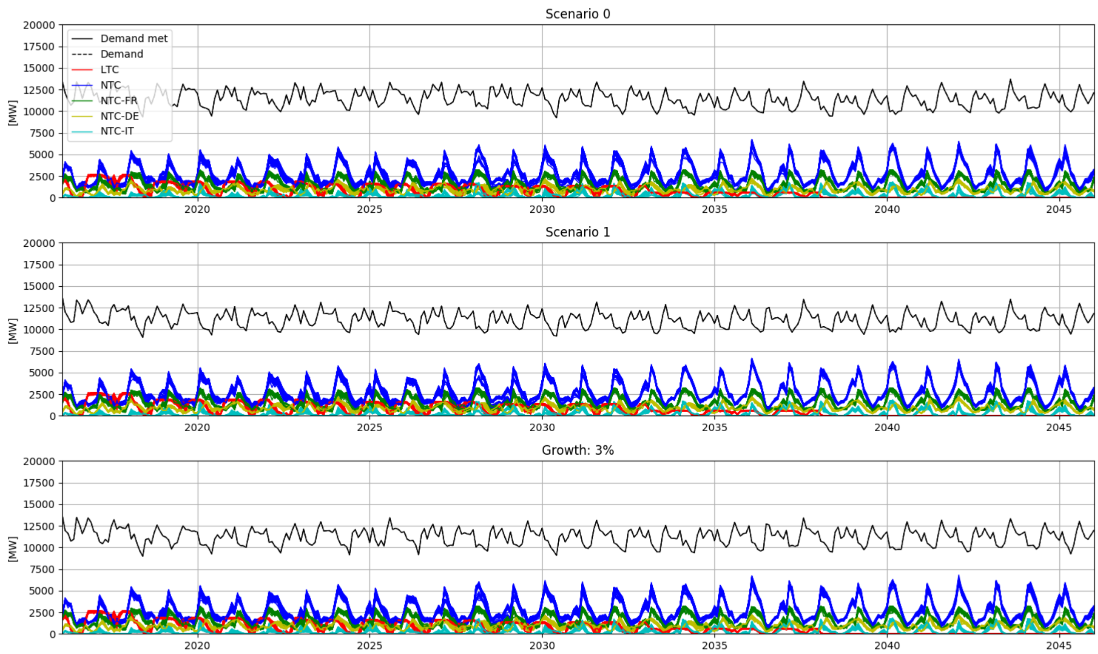

Figure F.8 – Overall demand with the electricty supplied by NTC and LTC sources for all policy scenarios and for the 0% demand growth scenario.

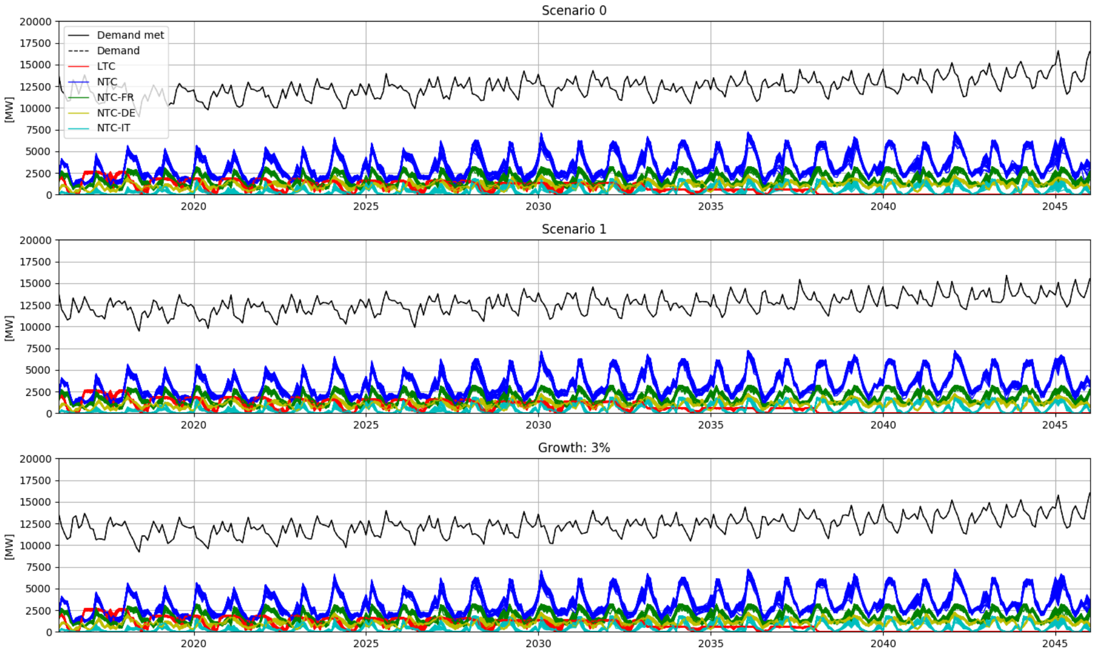

Figure F.9 – Overall demand with the electricty supplied by NTC and LTC sources for all policy scenarios and for the 1.5% demand growth scenario.

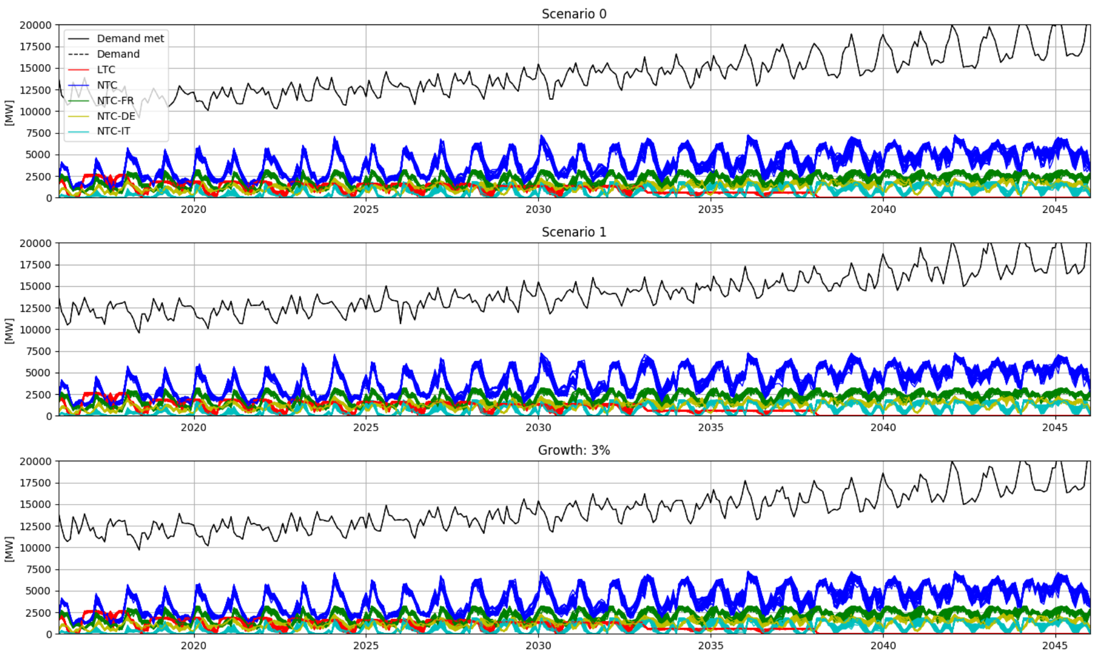

Figure F.10 – Overall demand with the electricty supplied by NTC and LTC sources for all policy scenarios and for the 3% demand growth scenario.

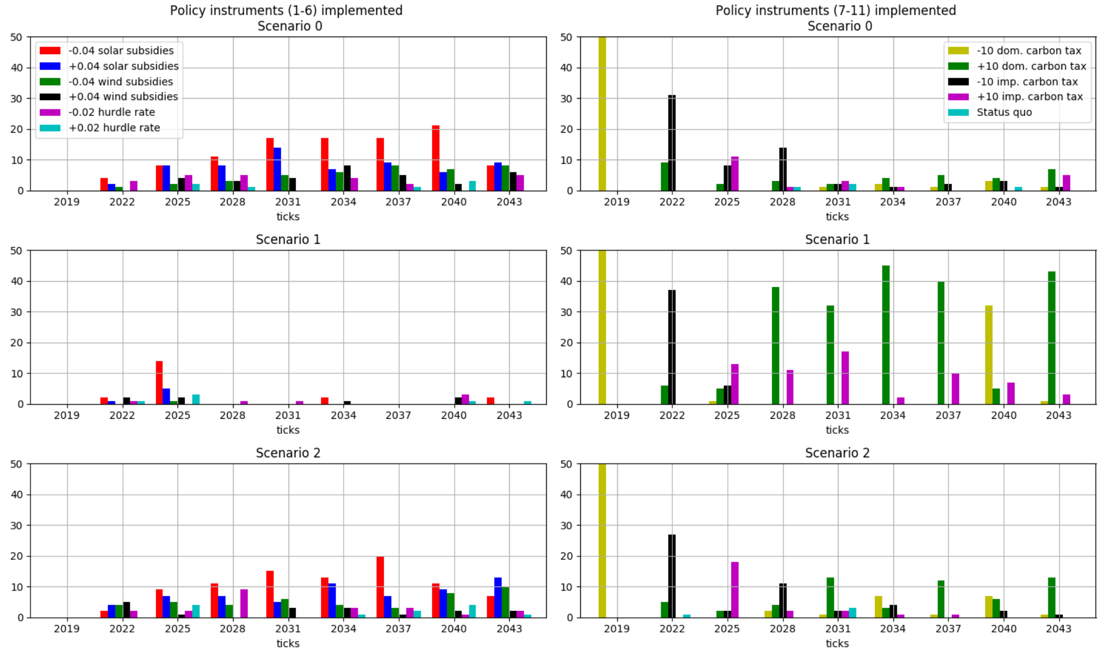

Figure F.11 – Cumulative policy instrument selection for each demand growth for all fifty repetitions for all policy scenarios and for the 0% demand growth scenario.

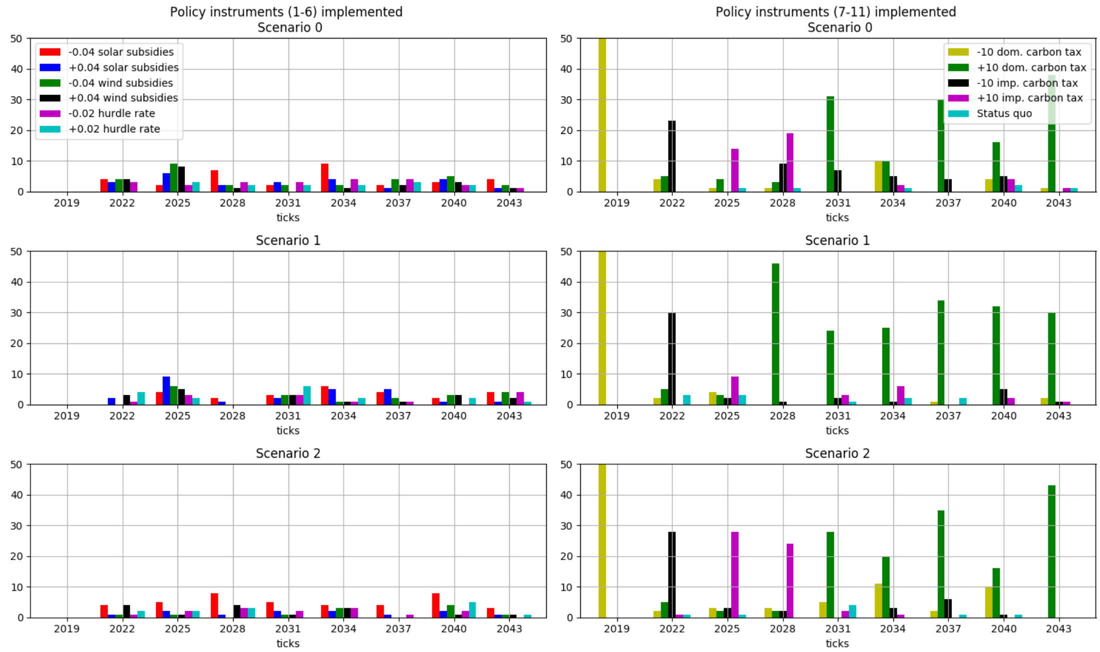

Figure F.12 – Cumulative policy instrument selection for each demand growth for all fifty repetitions for all policy scenarios and for the 1.5% demand growth scenario.

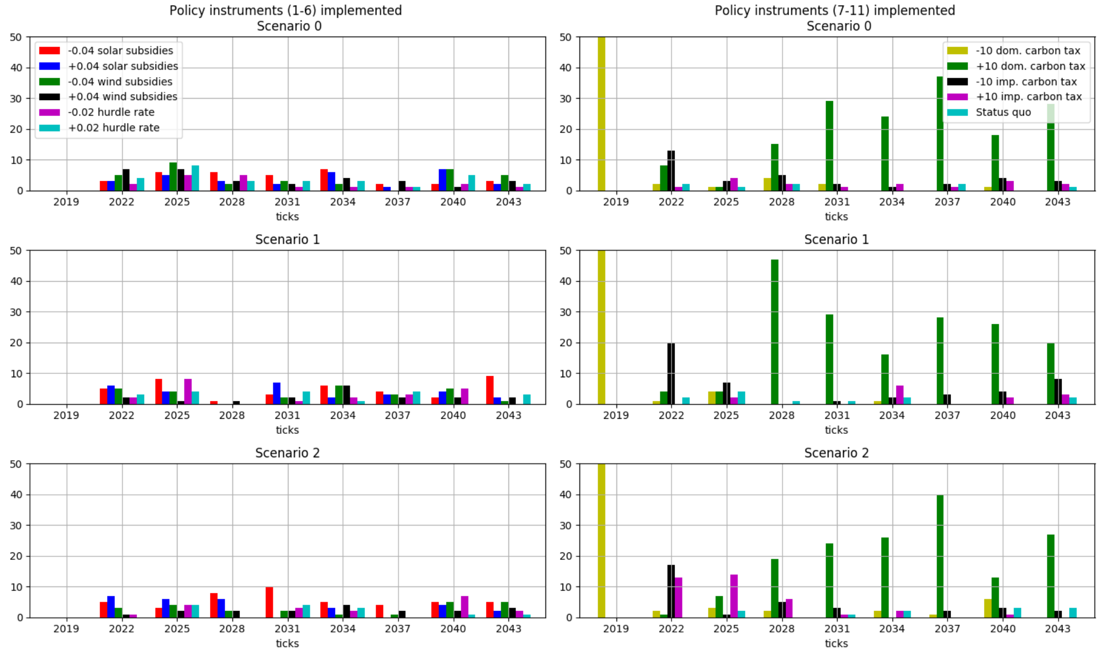

Figure F.13 – Cumulative policy instrument selection for each demand growth for all fifty repetitions for all policy scenarios and for the 3% demand growth scenario.

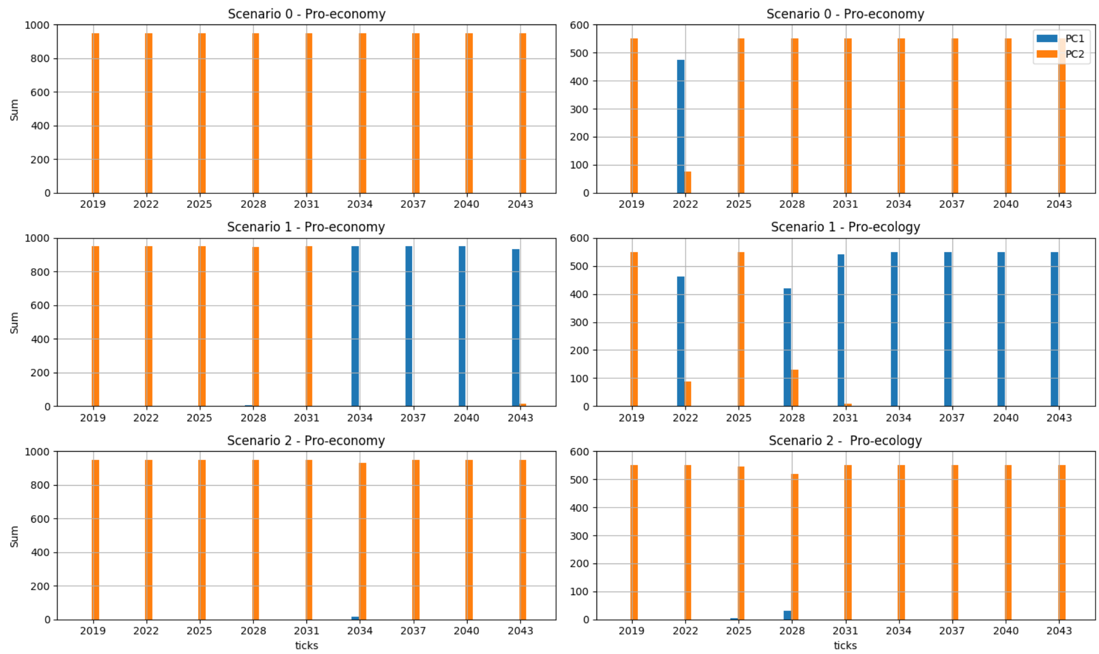

Figure F.14 – Policy core issues selected per affiliation for all policy scenarios and for the 0% demand growth scenario.

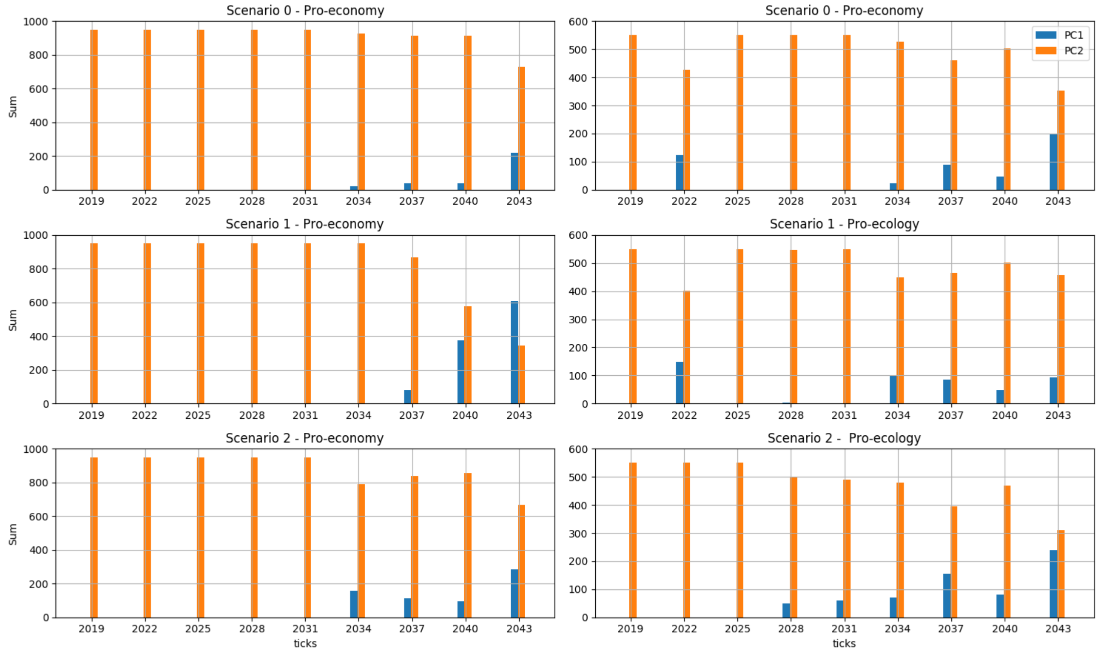

Figure F.15 – Policy core issues selected per affiliation for all policy scenarios and for the 1.5% demand growth scenario.

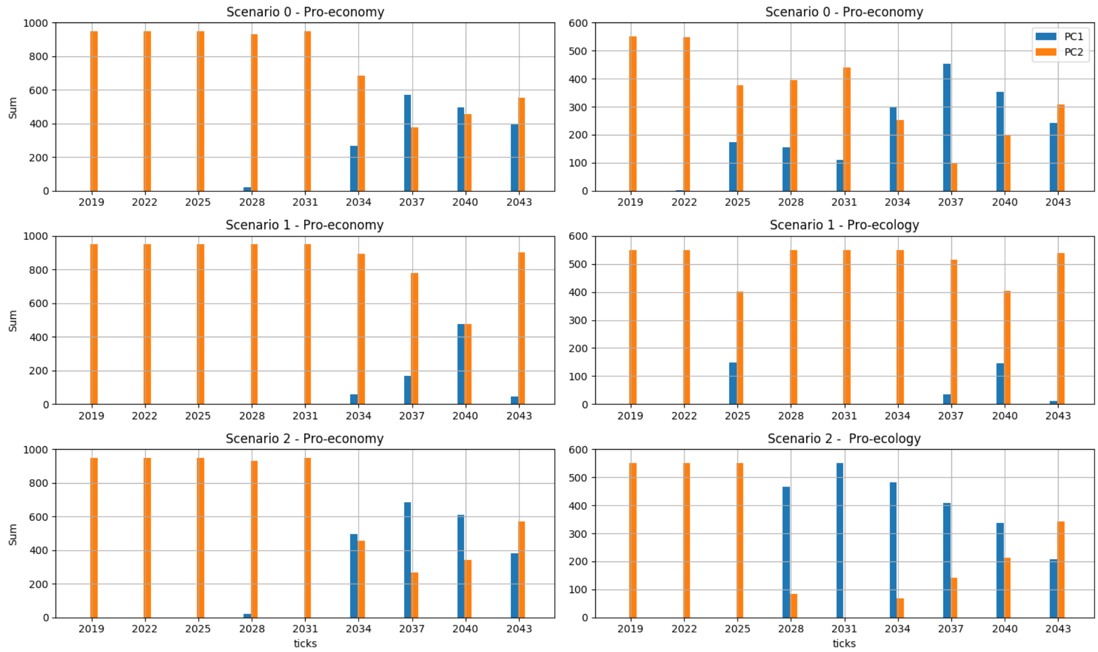

Figure F.16 – Policy core issues selected per affiliation for all policy scenarios and for the 1.5% demand growth scenario.

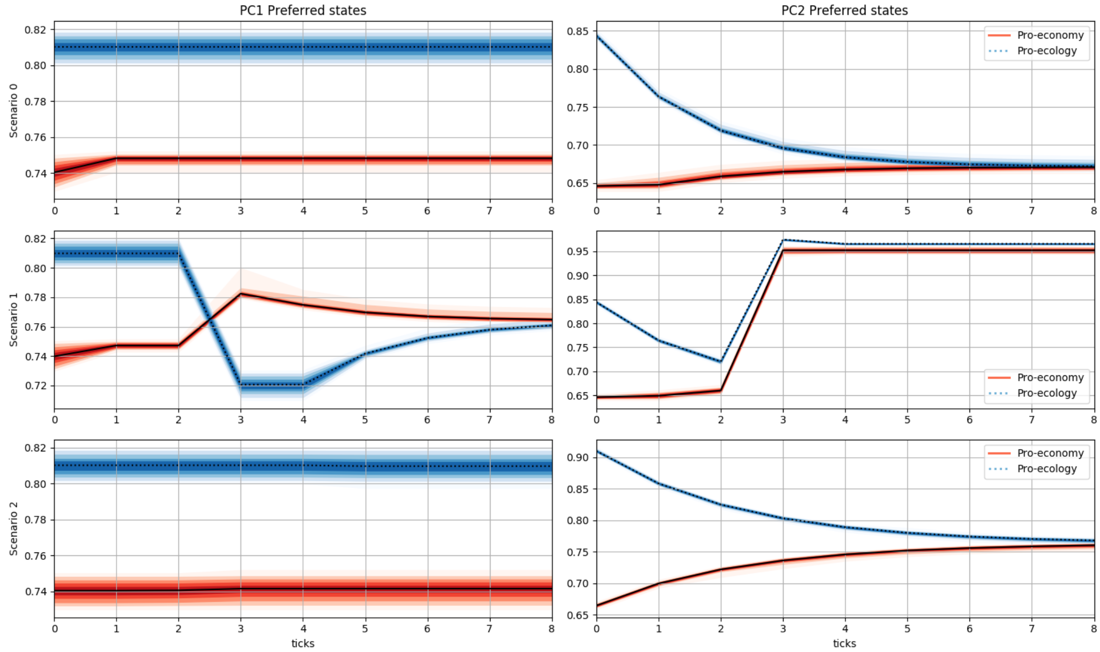

Figure F.17 – Policy core issue preferred states for all policy scenarios and for the 0% demand growth scenario.

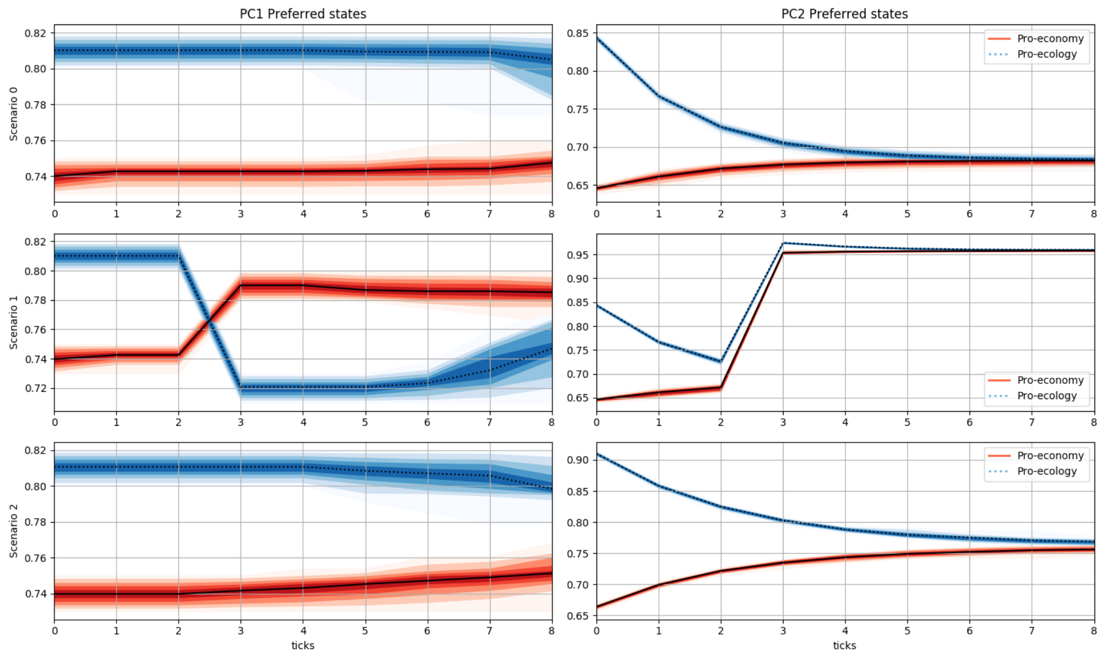

Figure F.18 – Policy core issue preferred states for all policy scenarios and for the 1.5% demand growth scenario.

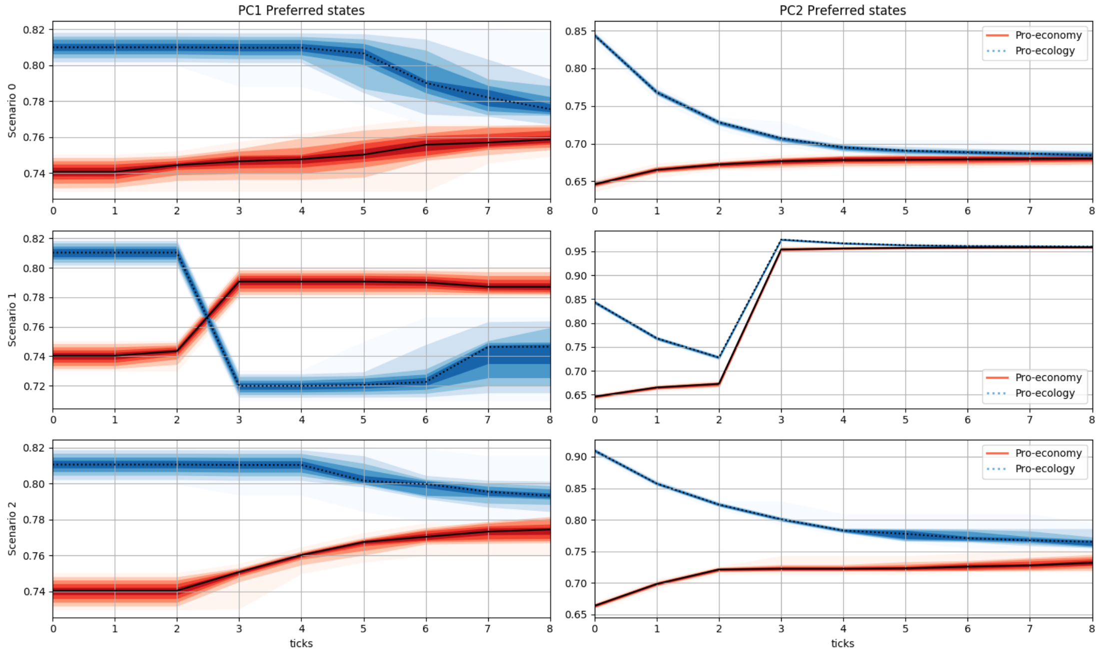

Figure F.19 – Policy core issue preferred states for all policy scenarios and for the 3% demand growth scenario.

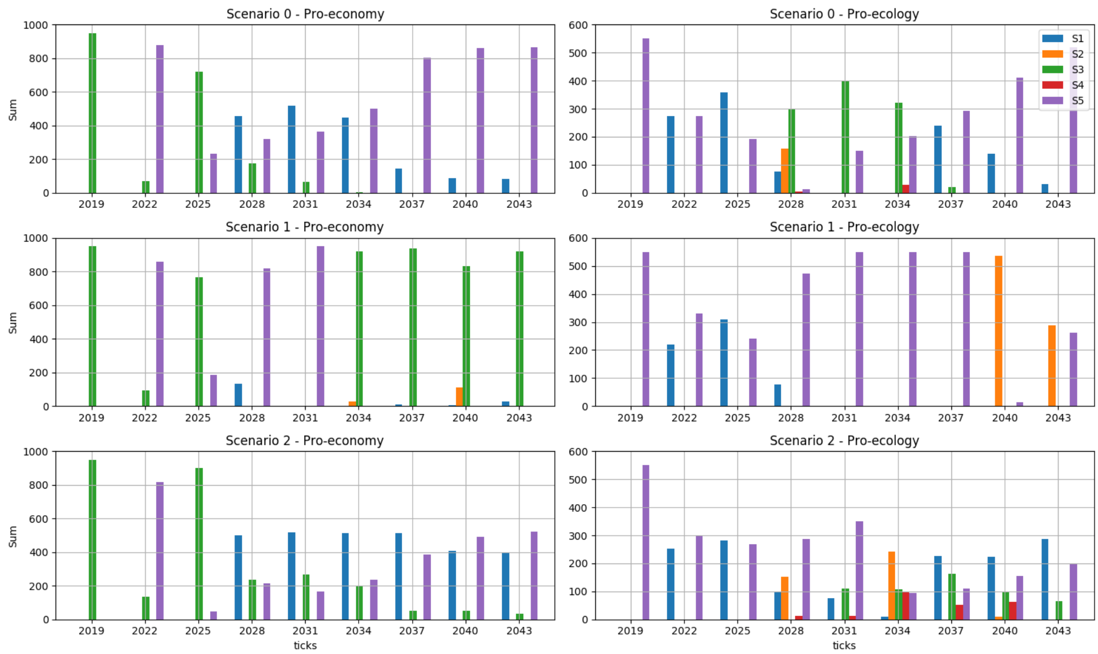

Figure F.20 – Secondary issue selected per affiliation for all policy scenarios and for the 0% demand growth scenario.

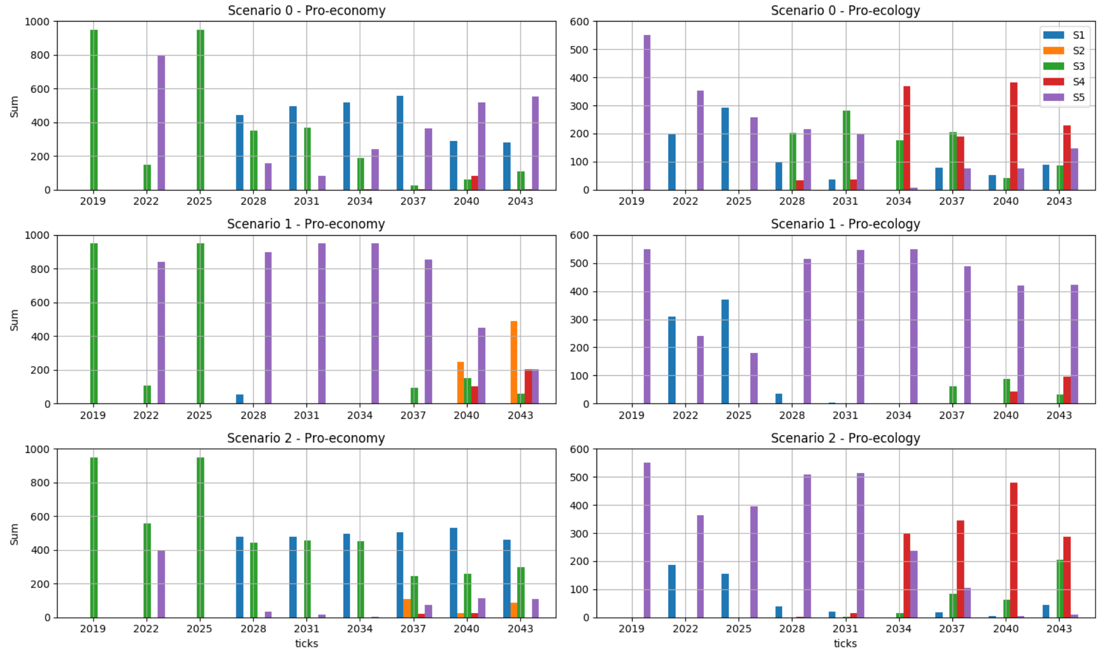

Figure F.21 – Secondary issue selected per affiliation for all policy scenarios and for the 1.5% demand growth scenario.

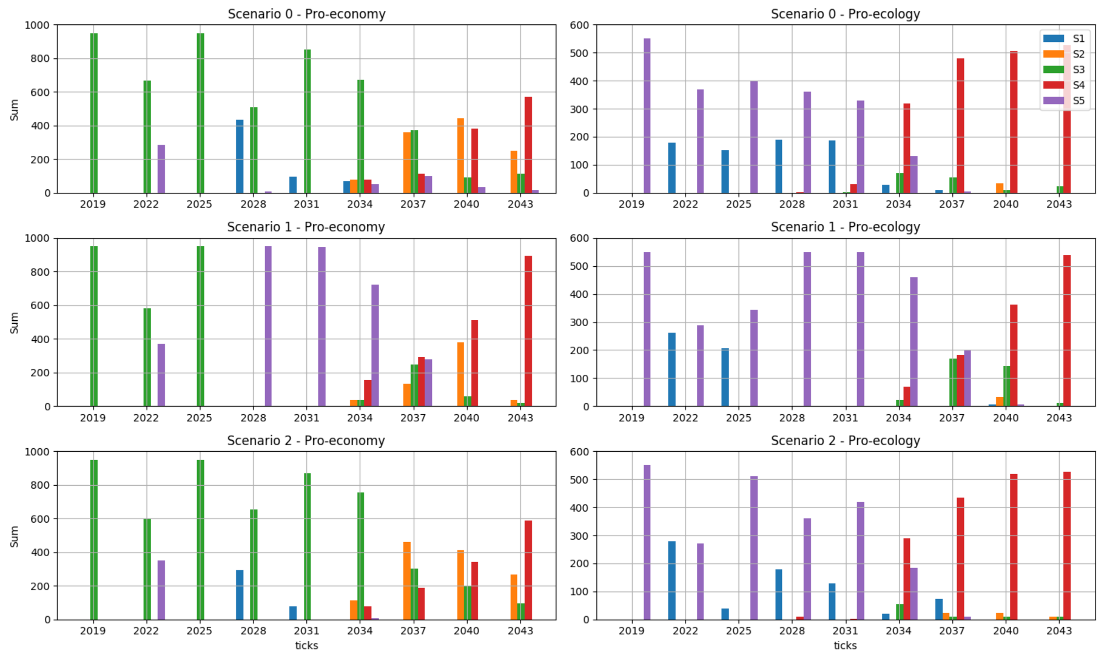

Figure F.22 – Secondary issue selected per affiliation for all policy scenarios and for the 3% demand growth scenario.

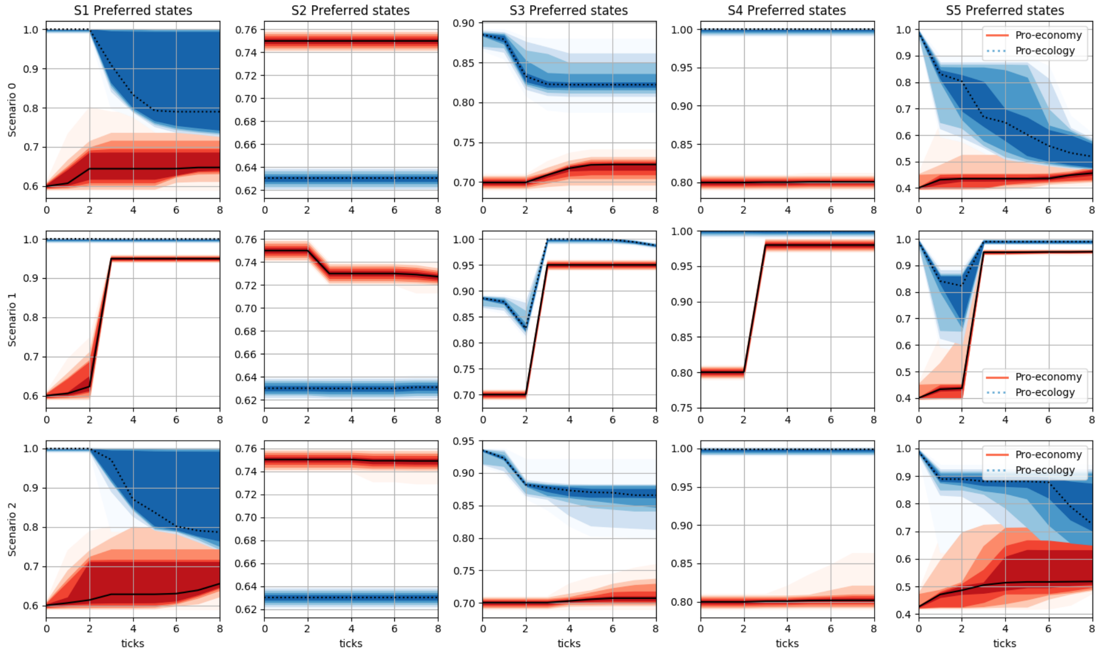

Figure F.23 – Secondary issue preferred states for all policy scenarios and for the 0% demand growth scenario.

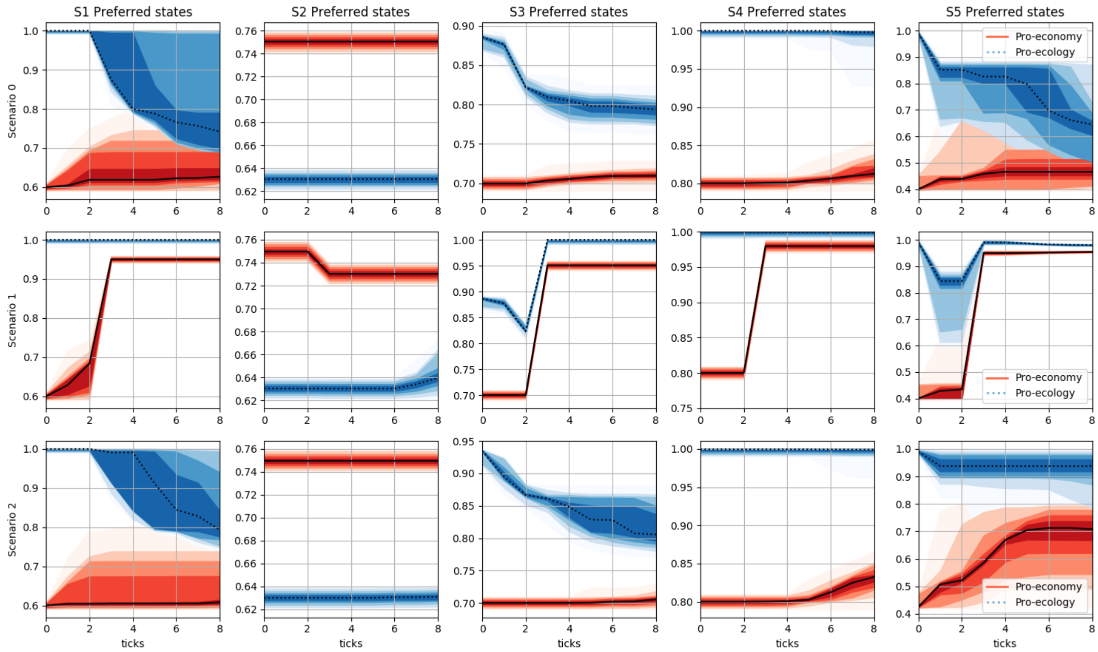

Figure F.24 – Secondary issue preferred states for all policy scenarios and for the 1.5% demand growth scenario.

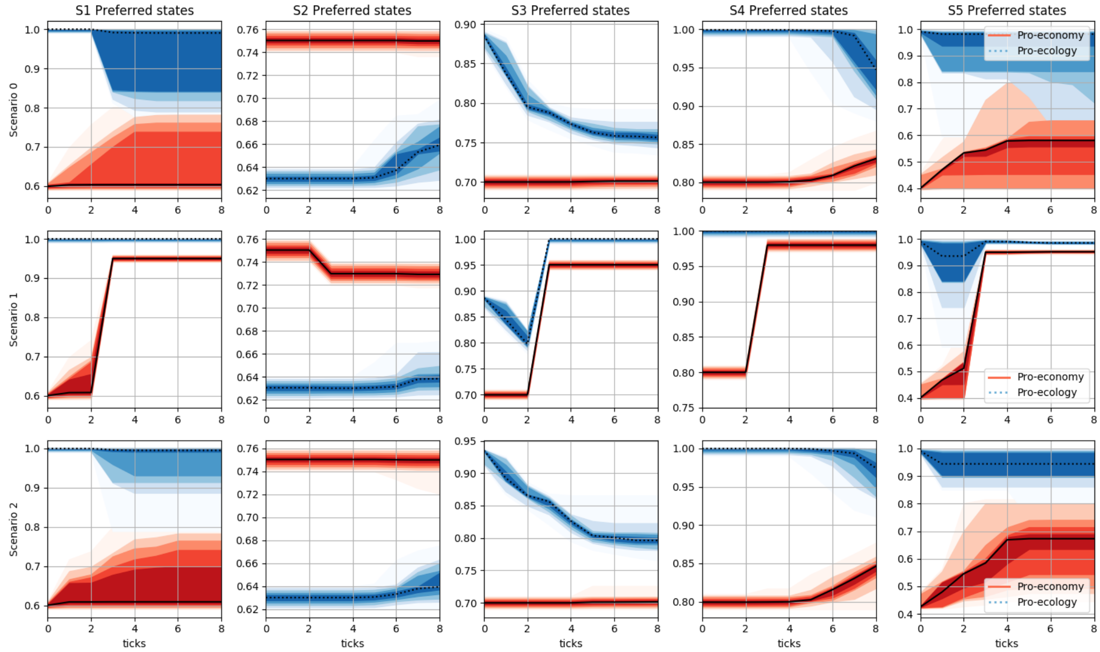

Figure F.25 – Secondary issue preferred states for all policy scenarios and for the 3% demand growth scenario.

Bibliography

Achterkamp, M. and Akkerman, A. (2003). Identifying latent conflict in collective bargaining. *Rationality and society*, 15(1):15–43.

Adam, S. and Kriesi, H. (2007). The network approach. *Theories of the policy process*, 2:129–154.

Adams, D. (2004). Usable knowledge in public policy. *Australian Journal of Public Administration*, 63(1):29–42.

Akkermans, H. (2001). Emergent supply networks: System dynamics simulation of adaptive supply agents. In *System Sciences, 2001. Proceedings of the 34th Annual Hawaii International Conference on*, pages 11–pp. IEEE.

Albright, E. A. (2011). Policy change and learning in response to extreme flood events in hungary: an advocacy coalition approach. *Policy studies journal*, 39(3):485–511.

Alink, F., Boin, A., and t'Hart, P. (2001). Institutional crises and reforms in policy sectors: the case of asylum policy in europe. *Journal of European Public Policy*, 8(2):286–306.

Arceneaux, K., Johnson, M., Lindstädt, R., and Vander Wielen, R. J. (2016). The influence of news media on political elites: Investigating strategic responsiveness in congress. *American Journal of Political Science*, 60(1):5–29.

Axtell, R. L., Epstein, J. M., Dean, J. S., Gumerman, G. J., Swedlund, A. C., Harburger, J., Chakravarty, S., Hammond, R., Parker, J., and Parker, M. (2002). Population growth and collapse in a multiagent model of the kayenta anasazi in long house valley. *Proceedings of the National Academy of Sciences*, 99(suppl 3):7275–7279.

Baier, V. E., March, J. G., and Saetren, H. (1986). Implementation and ambiguity. *Scandinavian Journal of Management Studies*, 2(3-4):197–212.

Bailey, M. (2003). The other side of the coin: A computational analysis of publicly financed campaigns. In Kollman, K., Miller, J. H., and Page, S. E., editors, *Computational Models in Political Economy*, chapter 6, pages 159–186. MIT Press.

Bandelow, N. C., Vogeler, C. S., Hornung, J., Kuhlmann, J., and Heidrich, S. (2017). Learning as a necessary but not sufficient condition for major health policy change: A qualitative comparative analysis combining acf and msf. *Journal of Comparative Policy Analysis: Research and Practice*, pages 1–16.

Bankes, S. (2011). The use of complexity for policy exploration. *The SAGE handbook of complexity and management*, pages 570–589.

Bankes, S. C. (2002). Tools and techniques for developing policies for complex and uncertain systems. *Proceedings of the National Academy of Sciences*, 99(suppl 3):7263–7266.

Baron, D. P. and Herron, M. C. (2003). A dynamic model of multidimensional collective choice. In Kollman, K., Miller, J. H., and Page, S. E., editors, *Computational models in political economy*, chapter 1, pages 13–48. MIT Press.

Barreteau, O., Bots, P., Daniell, K., Etienne, M., Perez, P., Barnaud, C., Bazile, D., Becu, N., Castella, J.-C., Daré, W., et al. (2013). Participatory approaches. In *Simulating social complexity*, pages 197–234. Springer.

Barton, C. M., Ullah, I. I., and Bergin, S. (2010). Land use, water and mediterranean landscapes: modelling long-term dynamics of complex socio-ecological systems. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 368(1931):5275–5297.

Baumgartner, F. R. and Jones, B. D. (1993). *Agendas and instability in American politics*. University of Chicago Press, 2nd edition edition.

Baumgartner, F. R., Jones, B. D., and Mortensen, P. B. (2014). Punctuated equilibrium theory: Explaining stability and change in public policymaking. In Sabatier, P. A. and Weible, C. M., editors, *Theories of the policy process*, pages 59–103. Westview Press Boulder, CO, third edition edition.

Bendor, J., Diermeier, D., and Ting, M. (2003). Recovering behavioralism: Adaptively rational strategic behavior with endogenous aspirations. In Kollman, K., Miller, J. H., and Page, S. E., editors, *Computational political economy*, chapter 8, pages 213–274. MIT Press.

Benoit, K. and Laver, M. (2006). *Party Policy in Modern Democracies*. Routledge.

Berardo, R. and Scholz, J. T. (2010). Self-organizing policy networks: risk, partner selection, and cooperation in estuaries. *American Journal of Political Science*, 54(3):632–649.

Berger, T. (2001). Agent-based spatial models applied to agriculture: a simulation tool for technology diffusion, resource use changes and policy analysis. *Agricultural economics*, 25(2-3):245–260.

Bhavnani, R., Miodownik, D., and Nart, J. (2008). Rescape: An agent-based model for modeling resources, ethnicity and conflict. *Journal of Artificial Societies and Social Simulation*, 11(2).

Birkland, T. A. and Lawrence, R. G. (2009). Media framing and policy change after columbine. *American Behavioral Scientist*, 52(10):1405–1425.

Black, D. et al. (1958). *The theory of committees and elections*. Springer.

Bloodgood, E. and Clough, E. (2008). Ngos in international affairs: An agent-based model and the boomerang effect. In *International Studies Association Annual Conference*, pages 26–28.

Boahene, K., Snijders, T. A., and Folmer, H. (1999). An integrated socioeconomic analysis of innovation adoption: the case of hybrid cocoa in ghana. *Journal of Policy Modeling*, 21(2):167–184.

Bots, P. W. G. (2013). Designing the policy analysis process. In Thissen, W. A. H. and Walker, W. E., editors, *Public Policy Analysis*, volume 179, pages 103–132. Springer.

Boyce, M. S. (1992). Wolf recovery for yellowstone national park: A simulation model. In *Wildlife 2001: Populations*, pages 123–138. Springer.

Brailsford, S. C., Desai, S. M., and Viana, J. (2010). Towards the holy grail: combining system dynamics and discrete-event simulation in healthcare. In *Proceedings of the Winter Simulation Conference*, pages 2293–2303. Winter Simulation Conference.

Brugger, H. I. (2017). *Local Policy Networks: Actors, Challenges and Cooperation in the Case of the German Energy Transition*. PhD thesis, Universitaet Konstanz.

Bueno De Mesquita, B. (1984). Forecasting policy decisions: an expected utility approach to post-khomeini iran. *PS: Political Science & Politics*, 17(2):226–236.

Bueno De Mesquita, B. (2006). Game theory, political economy, and the evolving study of war and peace. *American Political Science Review*, 100(4):637–642.

Bueno de Mesquita, B. (2011). A new model for predicting policy choices: Preliminary tests. *Conflict Management and Peace Science*, 28(1):65–87.

Bueno De Mesquita, B., Newman, D., and Rabushka, A. (1985). *Forecasting Political Events - The Future of Hong*. Yale University Press, New Haven and London.

Cairney, P. (2006). Venue shift following devolution: when reserved meets devolved in scotland. *Regional & Federal Studies*, 16(4):429–445.

Cairney, P. (2013). Standing on the shoulders of giants: How do we combine the insights of multiple theories in public policy studies? *Policy Studies Journal*, 41(1):1–21.

Cairney, P. (2015). How can policy theory have an impact on policymaking? the role of theory-led academic–practitioner discussions. *Teaching Public Administration*, 33(1):22–39.

Cairney, P. and Weible, C. M. (2015). Comparing and contrasting peter hall's paradigms and ideas with the advocacy coalition framework. In *Policy paradigms in theory and practice*, pages 83–99. Springer.

Calanni, J. C., Siddiki, S. N., Weible, C. M., and Leach, W. D. (2014). Explaining coordination in collaborative partnerships and clarifying the scope of the belief homophily hypothesis. *Journal of Public Administration Research and Theory*, 25(3):901–927.

Calder, M., Craig, C., Culley, D., de Cani, R., Donnelly, C. A., Douglas, R., Edmonds, B., Gascoigne, J., Gilbert, N., Hargrove, C., et al. (2018). Computational modelling for decision-making: where, why, what, who and how. *Royal Society open science*, 5(6):172096.

Cederman, L.-E. (1997). *Emergent actors in world politics: how states and nations develop and dissolve*, volume 2. Princeton University Press.

Cederman, L.-E. (2003). Modeling the size of wars: from billiard balls to sandpiles. *American Political science review*, 97(1):135–150.

Chahal, K. and Eldabi, T. (2008). Applicability of hybrid simulation to different modes of governance in uk healthcare. In *Simulation Conference, 2008. WSC 2008. Winter*, pages 1469–1477. IEEE.

Chang, M.-H. and Harrington Jr, J. E. (2006). Agent-based models of organizations. *Handbook of computational economics*, 2:1273–1337.

Chmieliauskas, A., Chappin, E. J., Davis, C. B., Nikolic, I., and Dijkema, G. P. (2012). New methods for analysis of systems-of-systems and policy: the power of systems theory, crowd sourcing and data management. In *System of Systems*. InTech.

Choi, T. and Robertson, P. J. (2014). Deliberation and decision in collaborative governance: A simulation of approaches to mitigate power imbalance. *Journal of Public Administration Research and Theory*, 24(2):495–518.

Cioffi-Revilla, C. (2009). Simplicity and reality in computational modeling of politics. *Computational and Mathematical Organization Theory*, 15(1):26–46.

Cioffi-Revilla, C. and Rouleau, M. (2010). Mason rebeland: An agent-based model of politics, environment, and insurgency. *International Studies Review*, 12(1):31–52.

Cohen, M. D., March, J. G., and Olsen, J. P. (1972). A garbage can model of organizational choice. *Administrative science quarterly*, pages 1–25.

Cohen, N. (2016). Policy entrepreneurs and agenda setting. In Zahariadis, N., editor, *Handbook of public policy agenda setting*, chapter 10, pages 180–199. Edward Elgar Publishing.

De Marchi, S. (2003). A computational model of voter sophistication, ideology, and candidate position taking. In Kollman, K., Miller, J. H., and Page, S. E., editors, *Computational models in political economy*, chapter 5, pages 143–158. MIT Press Cambridge, MA.

De Marchi, S. and Page, S. E. (2014). Agent-based models. *Annual review of political science*, 17:1–20.

De Vries, L., Chappin, E. J., and Richstein, J. C. (2013). Emlab-generation-an experimentation environment for electricity policy analysis. *TU Delft*.

Dela Santa, E. (2013). The politics of implementing philippine tourism policy: A policy network and advocacy coalition framework approach. *Asia Pacific Journal of Tourism Research*, 18(8):913–933.

Dewan, T. and Hortala-Vallve, R. (2011). The three as of government formation: Appointment, allocation, and assignment. *American Journal of Political Science*, 55(3):610–627.

Dunlop, C. A. and Radaelli, C. M. (2018). The lessons of policy learning: types, triggers, hindrances and pathologies. *Policy & Politics*, 46(2):255–272.

Durant, R. F. and Diehl, P. F. (1989). Agendas, alternatives, and public policy: Lessons from the us foreign policy arena. *Journal of Public Policy*, 9(2):179–205.

Earnest, D. C. (2008). Coordination in large numbers: An agent-based model of international negotiations. *International Studies Quarterly*, 52(2):363–382.

Edmonds, B., Le Page, C., Bithell, M. C.-B. E., Grimm, V., Meyer, R., Montanola-Sale, C., Ormerod, P., and Squazzoni, F. (2019). Different modelling purposes. *Journal of Artificial Societies and Social Simulation*, 22(3).

Edmondson, D. L., Kern, F., and Rogge, K. S. (2018). The co-evolution of policy mixes and socio-technical systems: Towards a conceptual framework of policy mix feedback in sustainability transitions. *Research Policy*.

Elgin, D. J. and Weible, C. M. (2013). A stakeholder analysis of colorado climate and energy issues using policy analytical capacity and the advocacy coalition framework. *Review of Policy Research*, 30(1):114–133.

Elmelund-Præstekær, C. and Wien, C. (2008). What's the fuss about? the interplay of media hypes and politics. *The International Journal of Press/Politics*, 13(3):247–266.

Epstein, J. M. (2008). Why model? *Journal of Artificial Societies and Social Simulation*, 11(4):12.

Fischer, M., Ingold, K., Sciarini, P., and Varone, F. (2016). Dealing with bad guys: actor- and process-level determinants of the "devil shift" in policy making. *Journal of public policy*, 36(2):309–334.

Fowler, J. H. and Laver, M. (2008). A tournament of party decision rules. *Journal of Conflict Resolution*, 52(1):68–92.

Gailmard, S. and Patty, J. W. (2007). Slackers and zealots: Civil service, policy discretion, and bureaucratic expertise. *American Journal of Political Science*, 51(4):873–889.

Geller, A. (2011). The use of complexity-based models in international relations: a technical overview and discussion of prospects and challenges. *Cambridge Review of International Affairs*, 24(01):63–80.

Ghorbani, A., Bots, P., Dignum, V., and Dijkema, G. (2013). Maia: a framework for developing agent-based social simulations. *Journal of Artificial Societies and Social Simulation*, 16(2):9.

Ghorbani, A., Dechesne, F., Dignum, V., and Jonker, C. (2014). Enhancing abm into an inevitable tool for policy analysis. *Policy and Complex Systems*, 1(1).

Grimm, V., Berger, U., DeAngelis, D. L., Polhill, J. G., Giske, J., and Railsback, S. F. (2010). The odd protocol: a review and first update. *Ecological modelling*, 221(23):2760–2768.

Guetzkow, H., F., C., Alger, R. A., Brody, R. C., Noel, R. C., and Snyder, editors (1963). *Simulation in International Relations: Developments for Research and Teaching.*

Gulati, G. J., Hadlock, C. R., and Gainsborough, J. F. (2011). Vodys: an agent-based model for exploring campaign dynamics. *Social science computer review*, 29(2):250–272.

Haasnoot, M., Kwakkel, J. H., Walker, W. E., and ter Maat, J. (2013). Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world. *Global environmental change*, 23(2):485–498.

Haasnoot, M., Middelkoop, H., Offermans, A., Van Beek, E., and Van Deursen, W. P. (2012). Exploring pathways for sustainable water management in river deltas in a changing environment. *Climatic Change*, 115(3-4):795–819.

Haasnoot, M., Middelkoop, H., Van Beek, E., and Van Deursen, W. (2011). A method to develop sustainable water management strategies for an uncertain future. *Sustainable Development*, 19(6):369–381.

Hall, P. A. (1993). Policy paradigms, social learning, and the state: the case of economic policymaking in britain. *Comparative politics*, pages 275–296.

Hallegatte, S., Shah, A., Lempert, R., Brown, C., and Gill, S. (2012). Investment decision making under deep uncertainty-application to climate change. Technical report, The World Bank.

Hamarat, C., Kwakkel, J. H., and Pruyt, E. (2013). Adaptive robust design under deep uncertainty. *Technological Forecasting and Social Change*, 80(3):408–418.

Hansen, P., Liu, X., and Morrison, G. M. (2019). Agent-based modelling and socio-technical energy transitions: A systematic literature review. *Energy Research & Social Science*, 49:41–52.

Henry, A. D. (2011). Ideology, power, and the structure of policy networks. *Policy Studies Journal*, 39(3):361–383.

Henry, A. D. and Brugger, H. I. (2018). Agent-based explorations of environmental consumption in segregated networks. In García-Díaz, C. and Olaya, C., editors, *Social systems engineering: The design of complexity*, pages 199–214. Wiley Online Library.

Henry, A. D. and Dietz, T. (2011). Information, networks, and the complexity of trust in commons governance. *International Journal of the Commons*, 5(2):188–212.

Henry, A. D. and Dietz, T. (2012). Understanding environmental cognition. *Organization & Environment*, 25(3):238–258.

Herron, K. G. and Jenkins-Smith, H. C. (2014). Public perspectives on nuclear security. *Risk, Hazards & Crisis in Public Policy*, 5(2):109–133.

Herweg, N., Huß, C., and Zohlnhöfer, R. (2015). Straightening the three streams: Theorising extensions of the multiple streams framework. *European Journal of Political Research*, 54(3):435–449.

Hethcote, H. W., Wang, W., Han, L., and Ma, Z. (2004). A predator–prey model with infected prey. *Theoretical population biology*, 66(3):259–268.

Hofferbert, R. (1974). *The study of public policy*. Bobbs-Merrill.

Hogwood, B. W. and Peters, B. G. (1982). The dynamics of policy change: Policy succession. *Policy Sciences*, 14(3):225–245.

Holtz, G., Alkemade, F., de Haan, F., Köhler, J., Trutnevyte, E., Luthe, T., Halbe, J., Papachristos, G., Chappin, E., Kwakkel, J., et al. (2015). Prospects of modelling societal transitions: Position paper of an emerging community. *Environmental Innovation and Societal Transitions*, 17:41–58.

Howlett, M. (2018). Moving policy implementation theory forward: A multiple streams/critical juncture approach. *Public Policy and Administration*.

Howlett, M., McConnell, A., and Perl, A. (2015). Streams and stages: Reconciling kingdon and policy process theory. *European Journal of Political Research*, 54(3):419–434.

Howlett, M., McConnell, A., and Perl, A. (2016a). Moving policy theory forward: connecting multiple stream and advocacy coalition frameworks to policy cycle models of analysis. *Australian Journal of Public Administration*, 76(1):65–79.

Howlett, M., McConnell, A., and Perl, A. (2016b). Weaving the fabric of public policies: comparing and integrating contemporary frameworks for the study of policy processes. *Journal of Comparative Policy Analysis: Research and Practice*, 18(3):273–289.

Huitema, D. and Meijerink, S. (2010). *Water policy entrepreneurs: A Research Companion to Water Transitions around the Globe*. Edward Elgar.

Ingold, K., Fischer, M., and Cairney, P. (2017). Drivers for policy agreement in nascent subsystems: An application of the advocacy coalition framework to fracking policy in switzerland and the uk. *Policy studies journal*, 45(3):442–463.

Ingold, K. and Varone, F. (2011). Treating policy brokers seriously: evidence from the climate policy. *Journal of Public Administration Research and Theory*, page mur035.

Jackson, J. E. (2003). A computational theory of electoral competition. In Kollman, K., Miller, J. H., and Page, S. E., editors, *Computational Models in Political Economy*, chapter 4, pages 109–142. MIT Press Cambridge, MA.

Jager, W. and Edmonds, B. (2015). Policy making and modelling in a complex world. In *Policy Practice and Digital Science*, pages 57–73. Springer.

Jann, W. and Wegrich, K. (2007). Theories of the policy cycle. In Fischer, F., Miller, G. J., and Sidney, M. S., editors, *Handbook of Public Policy Analysis: Theproes, Politics, and Methods*, pages 43–62. Taylor & Francis, CRC Press.

Jenkins-Smith, H., Silva, C. L., Gupta, K., and Ripberger, J. T. (2014). Belief system continuity and change in policy advocacy coalitions: Using cultural theory to specify belief systems, coalitions, and sources of change. *Policy Studies Journal*, 42(4):484–508.

Jenkins-Smith, H. C. and Sabatier, P. A. (1999). The advocacy coalition framework: An assessment. *Theories of the policy process*, 118:117–166.

Jenkins-Smith, H. C., St. Clair, G. K., and Woods, B. (1991). Explaining change in policy subsystems: Analysis of coalition stability and defection over time. *American Journal of Political Science*, 35(4):851–80.

John, P. (2018). Theories of policy change and variation reconsidered: a prospectus for the political economy of public policy. *Policy Sciences*, 51(1):1–16.

Johnson, P. E. (1999). Simulation modeling in political science. *American Behavioral Scientist*, 42(10):1509–1530.

Jones, M. D., Peterson, H. L., Pierce, J. J., Herweg, N., Bernal, A., Lamberta Raney, H., and Zahariadis, N. (2016). A river runs through it: A multiple streams meta-review. *Policy Studies Journal*, 44(1):13–36.

Kammermann, L. and Strotz, C. (2014). Akteure und koalitionen in der schweizer energiepolitik nach fukushima. Master's thesis, University of Bern.

Karslen, R., Papachristos, G., and Rehmatulla, N. (2019). An agent-based model of climate-energy policies to promote wind propulsion technology in shipping. *Environmental Innovation and Societal Transitions*, 31:33–53.

Kim, D.-H. and Juhn, J.-H. (1997). System dynamics as a modeling platform for multi-agent systems. In *The 15th International Conference of the System Dynamics Society*, pages 1–7.

Kingdon, J. W. (1984). *Agendas, alternatives, and public policies*. Little Brown.

Kingdon, J. W. (2014). *Agendas, alternatives, and public policies*. Pearson News International Edition, second edition edition.

Klein, R. and Finger, M. (2020). The long term impact of the electorate on the swiss electricity market transition. In Kachi, A. and Hettich, P., editors, *Swiss Energy Governance*.

Klein, R., Meeuwsen, S., and Slinger, J. (2016). Strategies in an uncertain world: A systems dynamics analysis of different flood protection strategies.

Köhler, J., De Haan, F., Holtz, G., Kubeczko, K., Moallemi, E., Papachristos, G., and Chappin, E. (2018). Modelling sustainability transitions: An assessment of approaches and challenges. *Journal of Artificial Societies and Social Simulation*, 21(1).

Kollman, K., Miller, J. H., and Page, S. E. (2003). *Computational models in political economy*. MIT Press.

Kollman, K. and Page, S. E. (2005). Computational methods and models of politics. In Tesfatsion, L. and Judd, K., editors, *The Handbook of Computational Economics*, pages 1433–1463. Elsevier.

Kriesi, H. and Jegen, M. (2001). The swiss energy policy elite: The actor constellation of a policy domain in transition. *European Journal of Political Research*, 39(2):251–287.

Kukkonen, A., Ylä-Anttila, T., and Broadbent, J. (2017). Advocacy coalitions, beliefs and climate change policy in the united states. *Public Administration*, 95(3):713–729.

Kwadijk, J. C., Haasnoot, M., Mulder, J. P., Hoogvliet, M. M., Jeuken, A. B., van der Krogt, R. A., van Oostrom, N. G., Schelfhout, H. A., van Velzen, E. H., van Waveren, H., et al. (2010). Using adaptation tipping points to prepare for climate change and sea level rise: a case study in the netherlands. *Wiley interdisciplinary reviews: climate change*, 1(5):729–740.

Kwakkel, J. H., Haasnoot, M., and Walker, W. E. (2015). Developing dynamic adaptive policy pathways: a computer-assisted approach for developing adaptive strategies for a deeply uncertain world. *Climatic Change*, 132(3):373–386.

Kwakkel, J. H. and Pruyt, E. (2015). Using system dynamics for grand challenges: the esdma approach. *Systems Research and Behavioral Science*, 32(3):358–375.

Kwakkel, J. H., Walker, W. E., and Marchau, V. (2010). Adaptive airport strategic planning. *European Journal of Transport and Infrastructure Research*, 10(3).

Laslier, J.-F. and Goktuna, B. O. (2016). Opportunist politicians and the evolution of electoral competition. *Journal of Evolutionary Economics*, 26(2):381–406.

Lättilä, L., Hilletofth, P., and Lin, B. (2010). Hybrid simulation models–when, why, how? *Expert Systems with Applications*, 37(12):7969–7975.

Laver, M. (2005). Policy and the dynamics of political competition. *American Political Science Review*, 99(2):263–281.

Laver, M. and Benoit, K. (2003). The evolution of party systems between elections. *American Journal of Political Science*, 47(2):215–233.

Laver, M. and Schilperoord, M. (2007). Spatial models of political competition with endogenous political parties. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 362(1485):1711–1721.

Laver, M. and Sergenti, E. (2011). *Party competition: An agent-based model.* Princeton University Press.

Leach, W. D. and Sabatier, P. A. (2005). To trust an adversary: Integrating rational and psychological models of collaborative policymaking. *American Political Science Review*, 99(4):491–503.

Lempert, R., Scheffran, J., and Sprinz, D. F. (2009). Methods for long-term environmental policy challenges. *Global Environmental Politics*, 9(3):106–133.

Lempert, R. J. (2002). Agent-based modeling as organizational and public policy simulators. *Proceedings of the National Academy of Sciences*, 99(suppl 3):7195–7196.

Li, J. and Xiao, R. (2017). Agent-based modelling approach for multidimensional opinion polarization in collective behaviour. *Journal of Artificial Societies and Social Simulation*, 20(2).

Lin, J.-S. (2008). Learning and party competition. *Swiss Political Science Review*, 14(1):131–155.

Lindblom, C. E. (1959). The science of "muddling through". *Public administration review*, pages 79–88.

Lubell, M., Schneider, M., Scholz, J. T., and Mete, M. (2002). Watershed partnerships and the emergence of collective action institutions. *American Journal of Political Science*, pages 148–163.

Markard, J., Suter, M., and Ingold, K. (2016). Socio-technical transitions and policy change – advocacy coalitions in swiss energy policy. *Environmental Innovation and Societal Transitions*, 18:215 – 237.

Martinez-Moyano, I. J., Sallach, D. L., Bragen, M. J., and Thimmapuram, P. R. (2007). Design for a multilayer model of financial stability: Exploring the integration of system dynamics and agent-based models.

Matti, S. and Sandström, A. (2011). The rationale determining advocacy coalitions: Examining coordination networks and corresponding beliefs. *Policy Studies Journal*, 39(3):385–410.

Mayer, G. R., Sarjoughian, H. S., Allen, E. K., Falconer, S., and Barton, M. (2006). Simulation modeling for human community and agricultural landuse. *SIMULATION SERIES*, 38(2):65.

Mazarr, M. J. (2007). The iraq war and agenda setting. *Foreign Policy Analysis*, 3(1):1–23.

McBeth, M. K., Jones, M. D., and Shanahan, E. A. (2007). The narrative policy framework. In Sabatier, P. A., editor, *Theories of the policy process*, chapter 7, pages 220–261. Westview Press.

McCullough, D. R. and Barrett, R. H. (1992). *Wildlife 2001: populations*. Springer Science & Business Media.

McGinnis, M. D. (2011). An introduction to iad and the language of the ostrom workshop: a simple guide to a complex framework. *Policy Studies Journal*, 39(1):169–183.

Mettler, S. and SoRelle, M. (2014). Policy feedback theory. *Theories of the policy process*, pages 151–181.

Michael, D. R. and Chen, S. L. (2005). *Serious games: Games that educate, train, and inform*. Muska & Lipman/Premier-Trade.

Mintrom, M. and Luetjens, J. (2017). Policy entrepreneurs and problem framing: The case of climate change. *Environment and Planning C: Politics and Space*, 35(8):1362–1377.

Mintrom, M. and Norman, P. (2009). Policy entrepreneurship and policy change. *Policy Studies Journal*, 37(4):649–667.

Mintrom, M. and Vergari, S. (1996). Advocacy coalitions, policy entrepreneurs, and policy change. *Policy studies journal*, 24(3):420–434.

Muis, J. (2010). Simulating political stability and change in the netherlands (1998-2002): An agent-based model of party competition with media effects empirically tested. *Journal of Artificial Societies and Social Simulation*, 13(2).

Mustapa, S. I. and Bekhet, H. A. (2016). Analysis of co2 emissions reduction in the malaysian transportation sector: An optimisation approach. *Energy Policy*, 89:171–183.

Mutz, D. C. and Martin, P. S. (2001). Facilitating communication across lines of political difference: The role of mass media. *American political science review*, pages 97–114.

Nakamura, R. T. (1987). The textbook policy process and implementation research. *Review of Policy Research*, 7(1):142–154.

Nava Guerrero, G., Schwarz, P., and Slinger, J. (2016). A recent overview of the integration of system dynamics and agent-based modelling and simulation. In *34th international conference of the system dynamics society*. System Dynamics Society.

Nikolic, I., Lukszo, Z., Chappin, E., Warnier, M., Kwakkel, J., Bots, P., and Brazier, F. (2019). Guide for good modelling practice in policy support. Technical report, TU Delft.

Nohrstedt, D. (2009). Do advocacy coalitions matter? crisis and change in swedish nuclear energy policy. *Journal of public administration research and theory*, 20(2):309–333.

Nohrstedt, D. (2011). Shifting resources and venues producing policy change in contested subsystems: a case study of swedish signals intelligence policy. *Policy Studies Journal*, 39(3):461–484.

Nowlin, M. C. (2011). Theories of the policy process: State of the research and emerging trends. *Policy Studies Journal*, 39:41–60.

Nuno, D., Nuno, F., and Agostinho, C. R. (2017). Verifying and validating simulations. In Edmonds, B. and Meyer, R., editors, *Simulating Social Complexity*, chapter 9, pages 39–58. Springer, 2nd edition edition.

Oborn, E., Barrett, M., and Exworthy, M. (2011). Policy entrepreneurship in the development of public sector strategy: The case of london health reform. *Public Administration*, 89(2):325–344.

Offermans, A. (2016). The perspectives method; towards socially robust river management. In *Sustainable Development Research at ICIS: Taking stock and looking ahaed*, chapter 25, pages 287–298. Datawyse/Universitaire Pers Maastricht.

Onggo, B. S. (2014). Elements of a hybrid simulation model: A case study of the blood supply chain in low-and middle-income countries. In *Simulation Conference (WSC), 2014 Winter*, pages 1597–1607. IEEE.

Ostrom, E. (2007). Institutional rational choice: An assessment of the institutional analysis and development framework.

Pahl-Wostl, C. (2002). Participative and stakeholder-based policy design, evaluation and modeling processes. *Integrated assessment*, 3(1):3–14.

Papachristos, G. (2014). Towards multi-system sociotechnical transitions: why simulate. *Technology Analysis & Strategic Management*, 26(9):1037–1055.

Papachristos, G. (2017). Diversity in technology competition: The link between platforms and sociotechnical transitions. *Renewable and Sustainable Energy Reviews*, 73:291–306.

Papachristos, G. and Adamides, E. (2016). A retroductive systems-based methodology for socio-technical transitions research. *Technological Forecasting and Social Change*, 108:1–14.

Park, H. H. and Rethemeyer, R. K. (2012). The politics of connections: Assessing the determinants of social structure in policy networks. *Journal of Public Administration Research and Theory*, 24(2):349–379.

Pepinsky, T. B. (2005). From agents to outcomes: simulation in international relations. *European Journal of International Relations*, 11(3):367–394.

Petridou, E. (2014). Theories of the policy process: contemporary scholarship and future directions. *Policy studies journal*, 42:S12–S32.

Pierce, J. J. (2011). Coalition stability and belief change: Advocacy coalitions in us foreign policy and the creation of israel, 1922–44. *Policy Studies Journal*, 39(3):411–434.

Pruyt, E. (2015). From building a model to adaptive robust decision making using systems modeling. In *Policy Practice and Digital Science*, chapter 5, pages 75–93. Springer.

Rüb, F. (2016). Agenda-setting and policy-making in time: What the multiple-streams approach can tell us – and what it cannot. In *Decision-Making under Ambiguity and Time Constraints*, chapter 4. ECPR Press.

Sabatier, P. A. (1987). Knowledge, policy-oriented learning, and policy change: An advocacy coalition framework. *Knowledge*, 8(4):649–692.

Sabatier, P. A. (1988). An advocacy coalition framework of policy change and the role of policy-oriented learning therein. *Policy Sciences*.

Sabatier, P. A. (2007). *Theories of the policy process*. Westview Press, second edition.

Sabatier, P. A. and Weible, C. M. (2007). The advocacy coalition framework the advocacy coalition framework - innovations and clarifications. In Sabatier, P. A., editor, *Theories of the policy process*. Westview Press.

Saltelli, A., Bammer, G., Bruno, I., Charters, E., Di Fiore, M., Didier, E., Espeland, W. N., Kay, J., Piano, S. L., Mayo, D., et al. (2020). Five ways to ensure that models serve society: a manifesto. *Nature*, 582.

Schieritz, N. and Grobler, A. (2003). Emergent structures in supply chains-a study integrating agent-based and system dynamics modeling. In *System Sciences, 2003. Proceedings of the 36th Annual Hawaii International Conference on*, pages 9–pp. IEEE.

Schneider, M., Scholz, J., Lubell, M., Mindruta, D., and Edwardsen, M. (2003). Building consensual institutions: networks and the national estuary program. *American journal of political science*, 47(1):143–158.

Scott, T. A., Thomas, C. W., and Magallanes, J. M. (2019). Convening for consensus: Simulating stakeholder agreement in collaborative governance processes under different network conditions. *Journal of Public Administration Research and Theory*, 29(1):32–49.

Shanahan, E. A., McBeth, M. K., and Hathaway, P. L. (2011). Narrative policy framework: The influence of media policy narratives on public opinion. *Politics & Policy*, 39(3):373–400.

Simmons, R. H., Davis, B. W., Chapman, R. J., and Sager, D. D. (1974). Policy flow analysis: A conceptual model for comparative public policy research. *Western Political Quarterly*, 27(3):457–468.

Singh, D., Padgham, L., and Logan, B. (2016). Integrating bdi agents with agent-based simulation platforms. *Autonomous Agents and Multi-Agent Systems*, 30(6):1050–1071.

Smith, A. (2000). Policy networks and advocacy coalitions: explaining policy change and stability in uk industrial pollution policy? *Environment and Planning C: Government and Policy*, 18(1):95–114.

Snijders, T. A. (2001). The statistical evaluation of social network dynamics. *Sociological methodology*, 31(1):361–395.

Snijders, T. A., Van de Bunt, G. G., and Steglich, C. E. (2010). Introduction to stochastic actor-based models for network dynamics. *Social networks*, 32(1):44–60.

Sotirov, M. and Memmler, M. (2012). The advocacy coalition framework in natural resource policy studies - recent experiences and further prospects. *Forest policy and economics*, 16:51–64.

Sprinz, D. F., Bueno de Mesquita, B., Kallbekken, S., Stokman, F., Sælen, H., and Thomson, R. (2016). Predicting paris: Multi-method approaches to forecast the outcomes of global climate negotiations. *Politics and Governance*, 4(3):172–187.

Stave, K. A. (2003). A system dynamics model to facilitate public understanding of water management options in las vegas, nevada. *Journal of Environmental Management*, 67(4):303–313.

Steglich, C., Snijders, T. A., and Pearson, M. (2010). Dynamic networks and behavior: Separating selection from influence. *Sociological methodology*, 40(1):329–393.

Stokes Berry, F. and Berry, W. D. (1999). Innovation and diffusion models in policy research. *Theories of the policy process*, 169.

Stokman, F. N., Van Assen, M., Van der Knoop, J., and Van Oosten, R. C. (2000). Strategic decision making. *Advances in Group processes*, 17(13):1–153.

Stokman, F. N. and Van Oosten, R. (1994). The exchange of voting positions: an object-oriented model of policy networks. *European community decision making. Models, applications, and comparisons*, pages 105–127.

Susi, T., Johannesson, M., and Backlund, P. (2007). Serious games: An overview.

Sutter, A. (2011). Schweizer klimapolitik nach 2012. Master's thesis, University of Bern.

Swanson, D., Barg, S., Tyler, S., Venema, H., Tomar, S., Bhadwal, S., Nair, S., Roy, D., and Drexhage, J. (2010). Seven tools for creating adaptive policies. *Technological Forecasting and Social Change*, 77(6):924–939.

Swinerd, C. and McNaught, K. R. (2012). Design classes for hybrid simulations involving agent-based and system dynamics models. *Simulation Modelling Practice and Theory*, 25:118–133.

Taeihagh, A., Bañares-Alcántara, R., and Millican, C. (2009). Development of a novel framework for the design of transport policies to achieve environmental targets. *Computers & Chemical Engineering*, 33(10):1531–1545.

Taillandier, P., Therond, O., and Gaudou, B. (2012). A new bdi agent architecture based on the belief theory. application to the modelling of cropping plan decision-making. In *International environmental modelling and software society (iEMSs)*.

Teisman, G. R. (2000). Models for research into decision-makingprocesses: on phases, streams and decision-making rounds. *Public administration*, 78(4):937–956.

Thissen, W. A. and Walker, W. E. (2013). *Public Policy Analysis*. Springer.

Thomson, R., Stokman, F., and Torenvlied, R. (2003). Models of collective decision-making: Introduction. *Rationality and Society*, 15(1).

Torenvlied, R. and Thomson, R. (2003). Is implementation distinct from political bargaining? a micro-level test. *Rationality and society*, 15(1):64–84.

van Baal, P. (2016). Business implications of the energy transition in switzerland. Master's thesis, EPFL.

van Baal, P. (2019). The effectiveness of a strategic energy reserve during the energy transition: The case of switzerland. *Competition and Regulation in Network Industries*.

van Baal, P. and Finger, M. (2019). The effect of european integration on swiss energy policy and governance. *Politics and Governance*, 7(1):6–16.

Vincenot, C. E., Giannino, F., Rietkerk, M., Moriya, K., and Mazzoleni, S. (2011). Theoretical considerations on the combined use of system dynamics and individual-based modeling in ecology. *Ecological Modelling*, 222(1):210–218.

Walker, W., Haasnoot, M., and Kwakkel, J. (2013). Adapt or perish: a review of planning approaches for adaptation under deep uncertainty. *Sustainability*, 5(3):955–979.

Walker, W. E. (2000). Policy analysis: a systematic approach to supporting policymaking in the public sector. *Journal of Multicriteria Decision Analysis*, 9(1-3):11.

Walker, W. E., Marchau, V. A., and Swanson, D. (2010). Addressing deep uncertainty using adaptive policies: Introduction to section 2. *Technological Forecasting and Social Change*, 77(6):917–923.

Walker, W. E., Rahman, S. A., and Cave, J. (2001). Adaptive policies, policy analysis, and policy-making. *European journal of operational Research*, 128(2):282–289.

Wang, W., Lin, Y., Zhang, L., Rao, F., and Tan, Y. (2011). Complex patterns in a predator–prey model with self and cross-diffusion. *Communications in Nonlinear Science and Numerical Simulation*, 16(4):2006–2015.

Weible, C. M. (2008). Expert-based information and policy subsystems: a review and synthesis. *Policy Studies Journal*, 36(4):615–635.

Weible, C. M., Heikkila, T., and Pierce, J. (2018). Understanding rationales for collaboration in high-intensity policy conflicts. *Journal of Public Policy*, 38(1):1–25.

Weible, C. M. and Sabatier, P. A. (2005). Comparing policy networks: Marine protected areas in california. *Policy Studies Journal*, 33(2):181–201.

Weible, C. M., Sabatier, P. A., and Lubell, M. (2004). A comparison of a collaborative and top-down approach to the use of science in policy: establishing marine protected areas in california. *Policy Studies Journal*, 32(2):187–207.

Wilder, M. (2016). Whither the funnel of causality? *Canadian Journal of Political Science/Revue canadienne de science politique*, 49(4):721–741.

Wilensky, U. (1997). Netlogo fire model. *ed. Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University*.

Wilensky, U. and Reisman, K. (1998). Connectedscience: Learning biology through constructing and testing computational theories—an embodied modeling approach. *InterJournal of Complex Systems*, 234:1–12.

Wilensky, U. and Reisman, K. (2006). Thinking like a wolf, a sheep, or a firefly: Learning biology through constructing and testing computational theories—an embodied modeling approach. *Cognition and instruction*, 24(2):171–209.

Winkel, G. and Sotirov, M. (2011). An obituary for national forest programmes? analyzing and learning from the strategic use of "new modes of governance" in germany and bulgaria. *Forest Policy and Economics*, 13(2):143–154.

Zafonte, M. and Sabatier, P. (2004). Short-term versus long-term coalitions in the policy process: Automotive pollution control, 1963–1989. *Policy studies journal*, 32(1):75–107.

Zahariadis, N. (2003). *Ambiguity and choice in public policy: Political decision making in modern democracies*. Georgetown university press.

Zahariadis, N. (2007). The multiple streams framework: Structure, limitations, prospects. *Theories of the policy process*, 2:65–92.

Zahariadis, N. (2014). Ambiguity and multiple streams. In Sabatier, P. A. and Weible, C. M., editors, *Theories of the policy process*, number 2, pages 25–59. Westview Press, third edition edition.

Zhang, Y., Wang, Y., Chen, T., and Shi, J. (2020). Agent-based modeling approach for group polarization behavior considering conformity and network relationship strength. *Concurrency and Computation: Practice and Experience*, page e5707.

RAPHAËL KLEIN

Avenue de Tourbillon 36 B
1950 Sion
Switzerland

+41 78 733 69 64
+31 6 13 54 35 78
raphael.klein4573@gmail.com
Linkedin profile

STRENGTHS

- Multi-disciplinary engineer with knowledge both of engineering and political science.
- Ability to think critically and decompose complex problems into simpler entities.
- Enjoy working in multi-disciplinary and multi-national teams and on wide ranging topics.
- Provided and organised executive education in the energy sector, energy governance and related to the energy transition. Enjoy organising and directing workshops or serious game sessions with the goal of disseminating knowledge.
- Experienced in simulation and modelling in a range of different applications and methods.

EDUCATION

2017 – 2020 **Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland**
PhD candidate, Policy and Energy Modelling

2013 – 2017 **Delft University of Technology (TU Delft), Delft, The Netherlands**
- Master of Science M.Sc., in Engineering and Policy Analysis
  Specialization in modelling and simulation
- Master of Science M.Sc., in Aerospace Engineering
  Specialization in aerospace structures and materials

2009 – 2013 **Delft University of Technology (TU Delft), Delft, The Netherlands**
Bachelor of Science B.Sc., in Aerospace Engineering

RESEARCH EXPERIENCE

Doctoral thesis
2017 – 2020
(exp. July 2020)

Exploring policy change through agent-based simulation
MIR, CDM, EPFL, Lausanne, Switzerland
Supervisor: Prof. Matthias Finger
The aim of this work is to demonstrate the usefulness of a modelling and simulation to studying the policy process and advancing the policy process theories. The PhD first refines the work performed in the MSc thesis before developing a common language for modellers to use to simulate the policy process. Furthermore, the PhD thesis demonstrates on how this language is applicable to a number of different cases with varying amount of complexity for the policy process model.

M.Sc. thesis
2016-2017

Policy emergence – An agent-based model approach
TBM, TU Delft, Delft, The Netherlands
Supervisor: Dr. Igor Nikolic
The aim of this MSc thesis was to demonstrate that it is possible to simulate the policy process. The model developed, coded in Python, was demonstrated with a forest fire model and used a mix of a number of different policy process theories including the multiple streams framework, the advocacy coalition framework and the theory of the policy process.

PUBLICATIONS

Manuscripts under review

- Klein R., Ashkenazy A., Nikolic I. and Bots P.W.G. (under review) "Exploring the policy process theories using modelling and simulation", Policy Sciences.
- Klein R., Finger M. (2020) "The long term impact of the electorate on the Swiss electricity market transition". In Kachi, Hettich (ed) *Swiss Energy Governance*.

Manuscripts in preparation

- Klein R. "Modelling the advocacy coalition framework".
- Klein R. Slinger J. "Integrating flood safety simulation and policy beliefs".
- Klein R. "Simulating the policy process".

Conference papers

- Klein, R. (2019) "Understanding policy change using simulation of the advocacy coalition framework", Presented at the 4th International Conference on Public Policy (ICPP4), Montreal, Canada, June 2019.
- Klein, R. (2019) "Simulating policy process theories with agent-based modelling", Presented at the 4th International Conference on Public Policy (ICPP4), Montreal, Canada, June 2019.
- Klein R., Meeuwsen S. and Slinger J. (2016) "Strategies in an uncertain world: A Systems Dynamics analysis of different flood protection strategies", Presented at the 34th International Conference of the System Dynamics Society, Delft, The Netherlands, July 2016.

Conference posters

- Klein, R. "Governing the Swiss electricity sector", Presented at the 2nd International Conference on Energy Research and Social Science, Phoenix, AZ, USA, May 2019.
- Klein, R. "Using the ACF to simulate the policymaking process in socio-technical systems", Presented at the Social Simulation Conference, Stockholm, Sweden, August 2018.
- Klein, R.; Slinger, J. "Integrating flood safety simulation and policy beliefs", Presented at the International System Dynamics Conference, Reykjavík, Iceland, August 2018.

PROFESSIONAL EXPERIENCE

2017 – 2020 — **EPFL – Governing Energy Transition CAS program coordinator**
Lausanne, Switzerland
I devised, organised and executed the Certificate of Advanced Study program called Governing Energy Transition (GET). This program brought in speakers from the energy sector ranging from the government and the private sector to academics. Forty high level professionals of diverse backgrounds in the energy sector have graduated during my chairing and organisation of the program.

2016 — **DNV-GL & TU Delft – Consultant**
Delft , The Netherlands
I was contracted to help for the designing, testing and running of a serious game (Go2Zero) on local initiatives for the energy transition. The game is part of the Citizen Smartcity initiative.

2016 — **TU Delft – Student communications assistant**
Delft, The Netherlands
I was part of the team that organised the Introduction Programme, making sure that 2,000 foreign students are properly welcomed to Delft and introduced to the university.

2014 – 2015 — **NLR – Intern**
Amsterdam, The Netherlands
As an intern, I worked on my thesis at the Dutch Aerospace Institute. My work was focused on the domain of aeroservoelasticity.

2013 – 2014 **The Ocean Cleanup – Boom development coordinator**
Delft, The Netherlands
I was part of the initial team comprised of students and volunteers tasked to determine the feasibility of The Ocean Cleanup concept developed by Boyan Slat.

2013 – 2015 **TU Delft – Student assistant**
Delft, The Netherlands
I was a student assistant for a number of courses in the aerospace engineering bachelor curriculum. This included courses such as rocket propulsion, systems engineering or propulsion and power.

LANGUAGE PROFICIENCY

- English: Fluent, medium of instruction in university
- French: Mother tongue

EXTRA-CURRICULAR ACTIVITIES

Writing and reading in my free time; swimming, skiing; travelling and exploring new destinations.

PERSONAL DETAILS

Age 28, French citizen, holder of a driver's license