Prospect Theory Multi-Agent Based Simulations for Non-Rational Route Choice Decision Making Modelling

Dissertation zur Erlangung des naturwissenschaftlichen Doktorgrades der Julius–Maximilians–Universität Würzburg

vorgelegt von

Gustavo Kuhn Andriotti

aus Passo Fundo - RS - Brasilien

Würzburg, 2009

Eingereicht am: 12. Juni 2009

bei der Fakultät für Mathematik und Informatik

Contents

1	Intr	roduction	13
	1.1	Motivation	13
	1.2	Fundamentals	13
		1.2.1 Rationality	14
		1.2.2 Non-Rationality	14
	1.3	Agent Reasoning	15
	1.4	Contribution	15
	1.5	Experiments And Results	16
	1.6	Main Findings	18
	1.7	State-Of-The-Art Overview	18
	1.8	Text Structure	19
-	Б		0.1
I	Ba	ISIS	21
2	Uti	lity Theory	23
	2.1	Scales Of Measurement	24
	2.2	Prerequisites	25
		2.2.1 A Note On Parameters And Utility Function	26
	2.3	Utility Determinism	26
	2.4	General Properties	26
	2.5	Modelling	27
		2.5.1 Calibrating	28
		2.5.2 Non-Captured Parameters	28
	2.6	Summary	29
3	Rat	ionality And EUT	31
	3.1	Concepts Review	31
	3.2	Saint Petersburg Paradox	31
	3.3	Expected Utility Theory	32
		3.3.1 von Neumann's Formalization	33
	3.4	Discrete Choice Analysis	34
	3.5	Summary	35
4	Bey	ond Perfect Rationality	37
	4.1	Concepts Review	37
	4.2	Criticisms	37
		4.2.1 Bi-Parted Decision-Making System	38

	4.2.1 Bi-Parted Decision-Making System	38
	4.2.2 Reference Dependency And Status Quo	39
4.3	Bounded Rationality	39
4.4	Rationality Deviation	10
	4.4.1 The Allais Paradox	10

	4.4.2	The Ellsberg Paradox	41
4.5	The P	Prospect Theory	42
	4.5.1	Status Quo And Reference Dependence	42
	4.5.2	Probability Distortion Function	43
	4.5.3	Allais Paradox	44
4.6	Editin	ng Phase	45
	4.6.1	Coping With The Editing Phase	45
	4.6.2	Why Not Use Standard Clustering Methods	45
4.7	Advar	nces	46
	4.7.1	Cumulative Prospect Theory	46
	4.7.2	The Prospect Theory For Continuous Prospects	47
4.8	Summ	nary	48

II Contribution

49

· · · ·	$\begin{array}{cccccccccccccccccccccccccccccccccccc$
· · · ·	$\begin{array}{cccccccccccccccccccccccccccccccccccc$
· · · ·	$ \begin{array}{rrrr} & . & . & 53 \\ & . & . & 54 \\ & . & . & 55 \\ & . & . & 56 \\ & . & . & . & . \\ & . & . & . & . \\ & . & . & . & . \\ & . & . & . & . \\ & . & . & . & . \\ & . & . & . & . \\ & . & . & . & . \\ & . & . & . & . \\ & . & . & . & . \\ & . & . & . & . \\ & . & . & . & . \\ & . & . & . & . \\ & . & . & . & . \\ & . & . & . & . \\ & . & . & . & . \\ & . & . & . & . \\ & . & . & . & . \\ & . & . & . & . & . & . \\ & . & . & . & . & . \\ & . & . & . & . & . \\ & . & . & . & . & . \\ & . & . & . & . & . \\ & . & . & . & . & . & . \\ & . & . & . & . & . \\ & . & . & . & . & . & . \\ & . & . & . & . & . & . \\ & . & . & . & . & . & . \\ & . & . & . & . & . & . \\ & . & . & . & . & . & . & . \\ & . & . & . & . & . & . & . \\ & . & . & . & . & . & . & . & .$
· · · ·	54 55 56
· · ·	55 56
· · ·	56
	59
	61
	61
	63
	63
	63
	65
	65
	66
	67
	67
	68
	69
	71
	71
	71
	72
	74
	75

8	Eva	luation	77
	8.1	Concepts Review	77
	8.2	Evaluation Methodology	77
	8.3	Different Algorithm For System 2	78
	8.4	Common Parameters Across The Experiments	79
	8.5	Calibration	79
	8.6	Microeconomics	80
	87	Results Analysis	81
	8.8	Fl Farol	81
	0.0	8.81 Clustering Biss	81
		8.8.2 Scenario Experiments	82
	8.0	Solton	86
	8.10	Burgdorf	80
	0.10		09
	0.11		92
9	Stat	te-Of-The-Art	95
	9.1	Discrete Choice Modelling	95
	0	9.1.1 Utility Based Modelling	95
		9.1.2 Non-Utility Based Modelling	96
	92	Traffic Assignment	97
	0.2	9.2.1 Microeconomics	97
		9.2.2 Artificial Intelligence	98
			00
10	Con	clusion And Future Work 1	101
	10.1	Future Work	102
	10.1		102
		-	102
Bi	bliog	graphy 1	102
Bi	bliog Tecl	graphy 1 hnologies And Algorithms 1	102
Bi A	bliog Tecl	graphy 1 hnologies And Algorithms 1 Notation	L13 L15
Bi A	bliog Tecl A.1	graphy 1 hnologies And Algorithms 1 Notation	L13 L15 115
Bi A	bliog Tecl A.1 A.2	graphy 1 hnologies And Algorithms 1 Notation 1 Super-Network 1 A 2.1 Super network: Formalization	L13 L15 115 116
Bi A	bliog Tecl A.1 A.2	graphy 1 hnologies And Algorithms 1 Notation 1 Super-Network 1 A.2.1 Super-network: Formalisation A.2.2 Super-network	L13 L15 115 116 117
Bi A	bliog Tecl A.1 A.2	graphy 1 hnologies And Algorithms 1 Notation 1 Super-Network 1 A.2.1 Super-network: Formalisation A.2.2 Super-network: Generation Algorithm	L13 L15 115 116 117 119
Bi A	Tecl A.1 A.2	graphy 1 hnologies And Algorithms 1 Notation 1 Super-Network 1 A.2.1 Super-network: Formalisation A.2.2 Super-network: Generation Algorithm A.2.3 Super-network: Example	L13 L15 115 116 117 119 119
Bi A	bliog Tecl A.1 A.2	graphy 1 hnologies And Algorithms 1 Notation 1 Super-Network 1 A.2.1 Super-network: Formalisation A.2.2 Super-network: Generation Algorithm A.2.3 Super-network: Example A.2.4 Navigation	L13 L15 115 116 117 119 119 119
Bi A	A.3	graphy 1 hnologies And Algorithms 1 Notation 1 Super-Network 1 A.2.1 Super-network: Formalisation 1 A.2.2 Super-network: Generation Algorithm 1 A.2.3 Super-network: Example 1 A.2.4 Navigation 1 CFG 1 A.2.1 CFG 1	L13 L15 115 116 117 119 119 119 125
Bi A	bliog Tecl A.1 A.2	graphy 1 hnologies And Algorithms 1 Notation 1 Super-Network 1 A.2.1 Super-network: Formalisation A.2.2 Super-network: Generation Algorithm A.2.3 Super-network: Example A.2.4 Navigation CFG 1 A.3.1 CFG: Parser A.2.0 CFG	L13 L15 115 116 117 119 119 119 125 125
Bi A	bliog Tecl A.1 A.2 A.3	graphy 1 hnologies And Algorithms 1 Notation 1 Super-Network 1 A.2.1 Super-network: Formalisation A.2.2 Super-network: Generation Algorithm A.2.3 Super-network: Example A.2.4 Navigation CFG 1 A.3.1 CFG: Parser A.3.2 CFG: Generator	113 113 115 116 117 119 119 119 125 125 127
Bi A	bliog Tecl A.1 A.2 A.3	graphy 1 hnologies And Algorithms 1 Notation 1 Super-Network 1 A.2.1 Super-network: Formalisation 1 A.2.2 Super-network: Generation Algorithm 1 A.2.3 Super-network: Example 1 A.2.4 Navigation 1 CFG 1 A.3.1 CFG: Parser 1 A.3.2 CFG: Generator 1 A.3.3 Parser States 1	L13 L15 L15 L16 L17 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L19 L119 L119 L119 L119L119 L119 L119 L119 L119 L119 L119 L119 L119 L119 L119 L119 L119 L119 L119 L119 L119 L119 L119 L119 L119 L119 L119 L119 L119 L1119 L119 L1119 L1119 L1119 L1119 L1119 L1119 L1119 L1119 L1119 L1119 L1119 L1119 L11119 L1119 L11111 L1111 L11111 L11111 L11111 L11111 L11111 L11111 L11111 L11111 L11111 L111111 L111111 L111111 L111111 L111111111 L1111111111
Bi	bliog Tecl A.1 A.2 A.3	graphy1hnologies And Algorithms1Notation1Super-Network1A.2.1 Super-network: Formalisation1A.2.2 Super-network: Generation Algorithm1A.2.3 Super-network: Example1A.2.4 Navigation1CFG1A.3.1 CFG: Parser1A.3.2 CFG: Generator1A.3.4 Parser Complexity1	113 115 115 116 117 119 119 119 125 125 127 132 132
Bi	A.3	graphy1hnologies And Algorithms1Notation1Super-Network1A.2.1 Super-network: Formalisation1A.2.2 Super-network: Generation Algorithm1A.2.3 Super-network: Example1A.2.4 Navigation1CFG1A.3.1 CFG: Parser1A.3.2 CFG: Generator1A.3.3 Parser States1A.3.4 Parser Complexity1Dijkstra1	113 115 115 116 117 119 119 119 119 119 125 127 132 132 132 133
Bi	A.3 A.4	graphy 1 hnologies And Algorithms 1 Notation 1 Super-Network 1 A.2.1 Super-network: Formalisation 1 A.2.2 Super-network: Generation Algorithm 1 A.2.3 Super-network: Example 1 A.2.4 Navigation 1 CFG 1 A.3.1 CFG: Parser 1 A.3.2 CFG: Generator 1 A.3.3 Parser States 1 A.3.4 Parser Complexity 1 Dijkstra 1 A.4.1 Modified relaxation function 1	112 113 115 115 116 117 119 119 119 125 127 132 133 133
Bi	A.3 A.4	graphy 1 hnologies And Algorithms 1 Notation 1 Super-Network 1 A.2.1 Super-network: Formalisation 1 A.2.2 Super-network: Generation Algorithm 1 A.2.3 Super-network: Example 1 A.2.4 Navigation 1 CFG 1 A.3.1 CFG: Parser 1 A.3.2 CFG: Generator 1 A.3.3 Parser States 1 A.3.4 Parser Complexity 1 Dijkstra 1 A.4.1 Modified relaxation function 1 A.4.2 Algorithmic Complexity 1	112 113 115 115 116 117 119 119 125 127 132 133 133 133
Bi	A.3 A.4 A.5	graphy 1 hnologies And Algorithms 1 Notation 1 Super-Network 1 A.2.1 Super-network: Formalisation 1 A.2.2 Super-network: Generation Algorithm 1 A.2.3 Super-network: Example 1 A.2.4 Navigation 1 CFG 1 A.3.1 CFG: Parser 1 A.3.2 CFG: Generator 1 A.3.3 Parser States 1 A.3.4 Parser Complexity 1 Dijkstra 1 A.4.1 Modified relaxation function 1 A.4.2 Algorithmic Complexity 1	112 113 115 115 116 117 119 125 127 132 133 133 133 133 133
Bi	A.3 A.4 A.5	graphy 1 hnologies And Algorithms 1 Notation 1 Super-Network 1 A.2.1 Super-network: Formalisation 1 A.2.2 Super-network: Generation Algorithm 1 A.2.3 Super-network: Example 1 A.2.4 Navigation 1 CFG 1 A.3.1 CFG: Parser 1 A.3.2 CFG: Generator 1 A.3.3 Parser States 1 A.3.4 Parser Complexity 1 Dijkstra 1 A.4.1 Modified relaxation function 1 A.4.2 Algorithmic Complexity 1 A.4.2 Algorithmic Complexity 1	112 113 115 115 116 117 119 119 125 125 132 133 133 133 138
Bi	A.3 A.4 A.5 Ext	graphy 1 hnologies And Algorithms 1 Notation 1 Super-Network 1 A.2.1 Super-network: Formalisation 1 A.2.2 Super-network: Generation Algorithm 1 A.2.3 Super-network: Example 1 A.2.4 Navigation 1 CFG 1 A.3.1 CFG: Parser 1 A.3.2 CFG: Generator 1 A.3.3 Parser States 1 A.3.4 Parser Complexity 1 Dijkstra 1 A.4.1 Modified relaxation function 1 A.4.2 Algorithmic Complexity 1 Generators 1 Ta Figures 1	112 113 115 115 116 117 119 125 125 127 132 133 133 138 141
Bi A B	A.3 A.4 A.5 Ext:	graphy1hnologies And Algorithms1Notation1Super-Network1A.2.1 Super-network: Formalisation1A.2.2 Super-network: Generation Algorithm1A.2.3 Super-network: Example1A.2.4 Navigation1CFG1A.3.1 CFG: Parser1A.3.2 CFG: Generator1A.3.3 Parser States1A.3.4 Parser Complexity1Dijkstra1A.4.1 Modified relaxation function1A.4.2 Algorithmic Complexity1Generators1PT $\pi(\bullet)$ Function1	112 113 115 115 116 117 119 1125 125 127 132 133 133 133 138 141
Bi A B	A.3 A.4 A.5 Ext B.1 B.2 D.2	graphy1hnologies And Algorithms1Notation1Super-Network1A.2.1 Super-network: Formalisation1A.2.2 Super-network: Generation Algorithm1A.2.3 Super-network: Example1A.2.4 Navigation1CFG1A.3.1 CFG: Parser1A.3.2 CFG: Generator1A.3.3 Parser States1A.3.4 Parser Complexity1Dijkstra1A.4.1 Modified relaxation function1A.4.2 Algorithmic Complexity1Generators1PT $\pi(\bullet)$ Function1PU bulk in the state1PU bulk in the state1Public to the state1 <t< td=""><td>112 113 115 115 116 117 119 1125 125 127 132 133 133 133 133 134 141 141</td></t<>	112 113 115 115 116 117 119 1125 125 127 132 133 133 133 133 134 141 141

List of Figures

$1.1 \\ 1.2 \\ 1.3$	Prospect Theory functions15El Farol traffic scenario16 $bpr(\bullet)$ characteristic functions17	5 3 7
$4.1 \\ 4.2$	Prospect Theory functions43Weight values for probability and function π 44	} 1
$5.1 \\ 5.2 \\ 5.3 \\ 5.4$	The regular α_n for different α_0 and horizon of 100	5 5 7 8
$\begin{array}{c} 6.1 \\ 6.2 \end{array}$	Travel-time function 68 Fundamental diagram 69	3
$7.1 \\ 7.2 \\ 7.3$	Proposed agent architecture 72 System 1 75 Memory ageing process 76	255
8.1 8.2 8.3 8.4 8.5 8.6	El Farol traffic scenario81Travel-time functions87Travel-time functions by individuals in the Secondary route87Burgdorf scenario90Average occupancy evolution92Mean occupancy and standard deviation for CluPT engine93	2 7 2 3
A.1 A.2 A.3 A.4 A.5 A.6	From a digraph G to a super-network \mathcal{S} .117Labelled graph122Network view123Super-network in "aesthetic" view123One column "real" view from connecting edges124A possible path in the super-network124	723341
B.1 B.2 B.3	Function $\pi(\bullet)$ for different γ values	23

List of Tables

$1.1 \\ 1.2 \\ 1.3$	Occupation results of the Main route for density experiments 1 Table 18 from [Chm05] 1 Simulation results 1	17 18 18			
2.1	Shoe example	25			
$4.1 \\ 4.2 \\ 4.3$	Modified Allais problem 4 Ellsberg problem 4 Modified Allais problem 4	40 41 44			
$5.1 \\ 5.2 \\ 5.3$	Theory classification \dots	53 52 52			
 8.1 8.2 8.3 8.4 8.5 8.6 8.7 8.8 	Simulation parameters 79 Fixed parameters for horizon experiments 83 Simulation parameters for horizon experiments 83 Occupation results of the Main route for horizon experiments with 100 agents 83 Fixed parameters for agent population experiments 84 Simulation parameters for agent population experiments 84 Derived parameters from agent amount experiments with target density 0.3 84				
$\begin{array}{c} 8.9\\ 8.10\\ 8.11\\ 8.12\\ 8.13\\ 8.14\\ 8.15\\ 8.16\end{array}$	density of 0.3	35 35 35 36 38 88 88 89 91			
A.1 A.2	Parsing of $x = n + n$	27 35			

List of Algorithms

5.1	Cluster(•)
A.1	Supernet(•)
A.2	Network(•)
A.3	ConnectNet(•)
4.4	Vertices(•)
A .5	Edges(•)
4.6	$\operatorname{Grammar}(\bullet)$
A .7	LabelComb(•)
A .8	VarsComb(•)
A .9	RecurLbls(•)
A.10	RecurProd(•)
\ .11	Standard Dijkstra algorithm with modified heap
A.12	Relax(•)
A.13	ModRelax(•)
A .14	Implemented Dijkstra algorithm
A.15	ImpRelax(•)
A.16	ModifiedEdgeWeight(•)
A.17	Link-Elimination Shortest-Path

Chapter 1

Introduction

This thesis proposes a novel approach for modelling human decision-makers using Multi-Agent Simulations (MASim) and non-rational behaviour. This non-rational behaviour is here based on the Prospect Theory [KT79] (PT), which is compared to the rational behaviour in the Expected Utility Theory [vNM07] (EUT). This model was used to design a modified Q-Learning [Wat89, WD92] algorithm. The PT based Q-Learning was then integrated into a proposed agent architecture.

Because much attention is given to a limited interpretation of Simon's definition of boundedrationality, this interpretation is broadened here. Both theories, rationality and the non-rationality, are compared and the discordance in their results discussed.

The main contribution of this work is to show that an alternative is available to the EUT that is more suitable for human decision-makers modelling. The evidences show that rationality is not appropriated for modelling persons. Therefore, instead of fine-tuning the existent model the use of another one is proposed and evaluated. To tackle this, the route choice problem was adopted to perform the experiments. To evaluate the proposed model three traffic scenarios are simulated and their results analysed.

1.1 Motivation

The motivation for this work lies on the apparent lack of experimentation with other interpretations of the bounded-rationality pointed out by Simon [Sim55, Sim56]. What seems to be the only interpretation of Simon's critics is: limited/local knowledge and/or limited computational power. This is evidenced by recent surveys, such as [SPG03, SPG04, SPG07, SV00, Sto07] that do not even mention other interpretations. The standard approach is to limit the agent learning to what the agent can grasp from its surroundings and how it can work out this knowledge for its own benefit. But once these limits are eliminated the underlying behaviour is rational.

However, when observing human behaviour, this assumption was demonstrated wrong, i.e. even in problems where neither knowledge nor computational limits are an issue, people do not behave rationally. This points out that the problem is not on how to deal with the limits but rather the model, in this case the rational model. For this reason this thesis investigates an alternative to the rationality: a model where even if no limits are imposed to the agent it may still deviate from the rational behaviour. To demonstrate this, experiments were designed to show when, how, and why (in some extend) the proposed agent behaviour deviates from the rational behaviour.

1.2 Fundamentals

The basis for this work is the Utility Theory [Fis70] (UT). In the utility theory it is assumed that the decision process has two elements: the options and the evaluation function for these options. The options are the elements of the choice set and they represent the solution candidates for the decision problem. The evaluation function, called utility function, is a function that maps each option in the choice set to a numerical value. This means that the function $u : \mathbf{X} \to \mathbb{R}$ is a utility function if \mathbf{X} is the choice set and it maps each element from \mathbf{X} to a value from \mathbb{R}^{1} . The function $u(\bullet)$ is said to model the decision process if it reproduces every preference observed in the individuals being modelled. The preferences can be one of: \sim, \succ , or \prec . This means that given two options; an individual can express no preference (\sim), prefer the first over the second (\succ), or the second over the first (\prec). If the relations established by the $u(\bullet)$ correspond to the preferences, then it is said that $u(\bullet)$ is a valid utility function for the given decision problem.

The Expected Utility Theory [vN28, vNM07] (EUT)² was developed over the UT. This theory is accepted as the standard formalisation for the rational behaviour proposed by Bernoulli [Ber38, Ber54]. Another theory based on the UT is the Prospect Theory [KT79] (PT).³ This theory, however, does not conform to the rational behaviour and therefore is called here non-rational.

1.2.1 Rationality

The rational model analysed is the EUT formalisation, which is the standard for the rational behaviour. According to von Neumann and Morgenstern [vN28, vNM07] the corresponding single valued utility of an option \mathbf{x} in the choice set \mathbf{X} is: $u(\mathbf{x}) = \sum_{\langle x, p \rangle \in \mathbf{x}} = xp$, where $\mathbf{x} \in \mathbf{X}$ is an option with a probability distribution over outcomes [Osb03], $\langle x, p \rangle$ is an item in this probability distribution, x is a possible outcome, and p its associated probability. An option with varying outcomes is called a lottery, which is, as already said, a probability distribution over outcomes.

One of the consequences of this formalisation (see Sec 3.3.1) is that if a utility is, say, 5.0 it does not matter if it is $5.0 = 14.0 \times 0.1 + 4.0 \times 0.9$ or $5.0 = 8.0 \times 0.5 + 2.0 \times 0.5$. This means that probabilities and outcomes are fully compensatory and in this case both options are equivalent: $\mathbf{x} = \{\langle 14.0, 0.1 \rangle, \langle 4.0, 0.9 \rangle\} \equiv \mathbf{y} = \{\langle 8.0, 0.5 \rangle, \langle 2.0, 0.5 \rangle\}$. This is a consequence of the axioms in the rational model.

1.2.2 Non-Rationality

Several experiments show that the rational behaviour is not suitable for modelling the human decision-making. Among them are the Allais [All53, AH79] and the Ellsberg [Ell61] paradoxes. They show that people are non-rational for fairly simple decision problems. In other situations, where the individual must deal with decision chains and other more complex decision problems, the problem is even worse. Some of these issues are addressed by McFadden [McF99], Kahneman [Kah02], and Ariely [Ari08].

To cope with the non-rationality of human decision-making the PT was proposed, which reproduces the Allais paradox (in [KT79]). It is also supported by medical evidences [DMKSD06, KF06] to be a better approach for modelling people when faced with a decision problem.

The main difference between EUT and PT is in how the prospect⁴ is evaluated by the utility function. In the PT instead of the plain summation of the EUT $(eut(\mathbf{x}) = \sum xp)$ it distorts the parameters $(pt(\mathbf{x}) = \sum v(x)\pi(p))$. The functions $v(\bullet)$ and $\pi(\bullet)$ map the human perception of the outcomes and probabilities, respectively. The visual behaviour of both "perception" functions are depicted in Fig. 1.1.

The reason why the PT diverges from the EUT resides mainly in the $\pi(\bullet)$ function.⁵ This function has the following interpretation. The first "bump" (near the origin, in Fig.1.1b) means that people overestimate low probabilities, saying that they are hopeful about seldom outcomes. The other "bump" (near the unity, in Fig.1.1b) represents the scepticism toward high probabilities that are underestimated, i.e. they believe that an "external" factor may interfere and prevent the outcome to happen.

¹ The utility theory is covered by chapter 2.

 $^{^2}$ For an analysis of the EUT please address to chapter 3.

 $^{^{3}}$ The PT and rationality deviations are covered in chapter 4.

 $^{^4}$ The difference between prospects and lotteries is addressed in Sec. 4.6.

 $^{^5}$ This is the relevant point in this thesis.



Figure 1.1: Prospect Theory functions

These are the interpretations of the $\pi(\bullet)$ function behaviour. The first consequence is that probabilities are no longer compensatory as they are for the EUT, mentioned in Sec. 1.2.1, meaning that the options⁶ will have different utilities and are no longer equivalent for the PT. This breaks with some axioms of the EUT.

For the $v(\bullet)$ function the behaviour is the following. It is observed that people are riskier in the losses than in the gains, i.e. if the outcome is perceived as a loss (under the horizontal line in Fig. 1.1a) then people tend to be more sensible to variations. This is evident by the steepness of the function for negative x values. The contrary behaviour is observed for the gains, where the steepness is lower. This means that in gains people are less sensible and tend to behave with more caution.

Editing Phase

The reason why the Prospect Theory has such a name is because the lotteries need to pass through an editing phase to become prospects.⁷ The editing phase lacks a standard and only the concepts were presented in [KT79]. The idea is to group similar outcomes and aggregate their probability. This evokes the need for a clustering method to be run over the lottery formed by the received outcomes. This means that as the outcomes are received they must be aggregated and then are transformed into prospects, which is the product of the editing phase over a lottery. The clustering method adopted is however biased, which is evident in the results.⁸

1.3 Agent Reasoning

To cope with the non-rational behaviour, associated with the PT, it is necessary to have an agent reasoning based on this theory. To address this issue the Q-Learning [Wat89, WD92] algorithm was chosen.⁹ and it provides a fairly simple way to incorporate the PT.

1.4 Contribution

The contribution of this thesis has some major and some minor points. The first major contribution is the Prospect Theory based modified Q-Learning algorithm. This algorithm is similar to the standard Q-Learning [WD92] but instead of being based on the rational behaviour, the Expected Utility Theory, it is based on the non-rational behaviour of the Prospect Theory [KT79]. The

 $^{^6}$ This refers to the prospects ${\bf x}$ and ${\bf y}$ used as examples in Sec. 1.2.1.

⁷ This is the subject of Sec. 4.6.

 $^{^{8}}$ The clustering method is presented in Sec. 5.5 and the bias extensively discussed in Sec. 8.8.1.

⁹ The Q-Learning algorithm is covered by chapter 5.



Figure 1.2: El Farol traffic scenario

second main contribution is the clustering based editing phase that was not formalised in [KT79]. Additionally to the major contributions, an agent architecture is proposed for coping with the practicalities of using agents with the PT based Q-Learning. This architecture also tackles another issue that is the separation between reasoning and intuition.¹⁰ The last contribution is the structure for simulating traffic scenarios including the above mentioned contributions.¹¹

1.5 Experiments And Results

To evaluate the proposal of this thesis, first, a traffic scenario inspired in the El Farol Bar Problem [Art94] is tested. This scenario is a minority game instance, where, given two options one has a lower capacity than the other, i.e. any option is only attractive if the minority of the agents choose that option. The traffic scenario used is depicted in Fig. 1.2. There two routes are offered to go from O (origin) to D (destination). These routes are named *Main* and *Secondary* where the latter has half the capacity¹² of the first.

In this scenario the optimal split is 1/3 of the agents for the *Secondary* and 2/3 for the *Main* route, being this distribution the expected rational behaviour. The objective of the experiments were to verify if and when the PT diverges from the rational choice. If the PT does not diverge from the EUT then it is not relevant for the investigation of traffic but the results show that the PT indeed diverge from the rational behaviour (EUT).

Each experiment was repeated 100 times and the results aggregated into the mean (μ_{\bullet}) and standard deviation (σ_{\bullet}) . In the experiments it shows that the behavioural stability is reached with 1000 or more iterations for this scenario. To the analysis an extra field was added, the "err" field which represents how much the given mean value deviates from the expected rational value (μ_{EUT}) .

Before showing the results it is necessary to explain how the travel-time (the metric for the utility function) is calculated. In each iteration all agents are asked about their decisions and then each route is "burden" with the corresponding amount of agents. After the occupation for each route is determined, the travel-time for each route is calculated. To calculate the travel-time a simple function is used, which is known as the formula from the Bureau of Public Roads [Tra00]. This function determines the density and from this value it derives the vehicle flow, which gives the corresponding travel-time. This relation between density and flow is called the fundamental diagram (depicted in Fig. 1.3).

In the final experiment, the scenario was simulated with different target densities. A target density is the density expected for the rational equilibrium, i.e. when both routes have the same travel-time, which is a consequence of having the same density. The results are shown in Tab. 1.1. The values in Tab. 1.1 are the resulting occupation of the *Main* route over 100 repetitions, using 100 agents. The column *STD* corresponds to the results for agents using the standard Q-Learning, *CluEUT* is the clustered version of the first (to verify if the clustering method is biased), and *CluPT* the PT modified Q-Learning.

 $^{^{10}}$ This is the subject of the Sec. 4.2.1.

 $^{^{11}}$ The traffic modelling is addressed in chapter 6.

 $^{^{12}}$ Capacity refers to the amount of vehicles that fit into a particular road segment, called link. The capacity influences in how sensible a link is, i.e. how the travel-time is affected by the link occupancy (how much vehicles are currently using the link).



Figure 1.3: $bpr(\bullet)$ characteristic functions

Table 1.1: Occupation results of the Main route for density experiments

Density	$\mu_{STD}(\sigma_{STD}): err$	$\mu_{CluEUT}(\sigma_{CluEUT}): err$	$\mu_{CluPT}(\sigma_{CluPT}): err$
0.1	66.45(7.09):-0.21	67.77(7.13):+1.10	69.19(7.24):+2.53
0.2	66.15(7.04):-0.50	67.35(7.13):+0.68	75.57(8.23):+8.90
0.3	66.12(6.99):-0.53	67.25(7.13):+0.59	76.93(8.32):+10.26
0.4	66.43(7.05):-0.23	67.33(7.11):+0.66	74.05(8.28):+7.39
0.5	66.57(7.12):-0.09	67.41(7.15):+0.74	72.99(7.93):+6.32
0.6	66.25(7.11):-0.40	67.89(7.08):+1.22	73.48(7.83):+6.81
0.7	66.18(7.12):-0.47	67.55(7.10):+0.88	73.88(8.13):+7.21
0.8	65.74(7.13):-0.92	67.66(7.12):+0.99	74.24(7.94):+7.58
0.9	65.85(6.96):-0.81	67.77(7.08):+1.10	73.20(7.97):+6.54
μ_{EUT}	$66.\overline{6}$		

The first issue concerns the use of the clustering method. As it can be seen in Tab. 1.1 the values of STD and CluEUT are fairly similar but with a noticeable bias (see the "err" for the CluEUT column in Tab. 1.1). Then it can also be seen that the CluPT column has a consistent divergence from the other two.

The optimal density, where the PT behaviour can be appreciated is in the range [0.2, 0.6]. This density range is a consequence of the fundamental diagram used (in Fig. 1.3) that forces the scenario to deal with constant congestion and travel-time variations.

One point left open is why the *Main* route is the stressed one and not the *Secondary*. The reason is on the shape of the $\pi(\bullet)$ function and because the *Secondary* route is the most "sensible" of them, i.e. a variation of one agent in the occupation has a stronger influence in the travel-time there than for the *Main*. This means that high and seldom travel-times are located at the first "bump" and therefore penalising the *Secondary* route while frequent and not optimal travel-times in the *Main* route are tolerated.

In a second scenario the validity of the PT was tested. In $[SSC^+05, Chm05]$ a scenario similar to the first one was presented to 18 persons that must choose between the two routes and have been monetarily compensated when they chose the fastest route.¹³ The 18 individual competed with each other and performed 100 decision rounds. The experiment was done in four phases, from which only two are of interest here. These two phases were: in a first run the individuals were not aware of the discrepancy in the routes' capacities, then, later, they were informed about it. The data is Tab. 1.2 and the simulated results in Tab. 1.3.

¹³ This data corresponds to the mean occupation for the secondary route in [Chm05], Tab. 18 at page 68 in the rows "Variation I" and "Variation II". The thesis of Chmura is available at http://www.ub.uni-duisburg.de/ETD-db/theses/available/duett-05152005-222337/.

Type	μ_s	σ_s	Deviation
Not-informed	4,50	1.38	-1.50
Informed	4, 44	1.01	-1.56
μ_{EUT}	6.00		

Table 1.2: Table 18 from [Chm05]

		~. · ·	•
Table	1.3:	Simulation	results

Engine		_	Error	
Engine	μ_s	08	Non-Informed (4.50)	Informed (4.44)
STD	5.61	1.47	1.11	1.17
CluEUT	5.16	1.30	0.66	0.72
CluPT	4.10	1.20	-0.40	-0.34
μ_{EUT}	6.00			

It can be observed that the CluPT has the closest results and also reproduce the tendency of preferring the *Main* instead of the *Secondary*.

In a third scenario the city of Burgdorf in Switzerland was simulated. Unfortunately the data available does not allow to evaluate the PT in this scenario.¹⁴ But it shows that the simulation framework is able to simulate real-world scenarios.

1.6 Main Findings

The results show that the PT based Q-Learning does not agree with the rational behaviour (Tab. 1.1). Second, it also shows that the PT based Q-Learning is better at reproducing real data (Tab. 1.3). The last point confirmed by the experiments is that the whole framework does scale for real world simulations.¹⁵ This means that the objectives for this work are fulfilled, i.e. a non-rational learning algorithm based on the Prospect Theory is proposed and evaluated. Not only the Prospect Theory is included in the learning algorithm but also a proposal for the editing phase is made and evaluated. It is also shown that this new algorithm perform better when reproducing real data than its rational counterpart does. Finally, it is shown that this framework does scale for real world scenarios.

1.7 State-Of-The-Art Overview

The decision-making modelling can be divided into some categories, according to the models proposal. The first division is between rational and non-rational based models. For the rational based the standard is the approach made by microeconomics, which is based in the Random Utility Models [GP06] (RUM). The main objective of these models is to extract the correlation among the options in the choice set, the $\mathbf{x} \in \mathbf{X}$. Among the models are: Logit [Ber44, Luc59], Multi Nominal Logit [McF74], Path-Size Logit [BAB99] and Probit [Bli34a, Bli34b].

For the non-rational models are the Prospect Theory [KT79], the Cumulative Prospect Theory [TK92], and their improvements to tackle with continuous prospects. Another proposal is the Theory of Small Feedbacks [BE03], but it has some issues that are discussed in Sec. 9.1.1. But not only utility based models were proposed and among them is the Fast and Frugal Way [GG96] but it is not adequate for general modelling (this is discussed in Sec. 9.1.2). In [BK03, KB04] Multi-agent Simulations are used to tackle with the traffic modelling, which is also the technique used in [BRV+05, BAN06, BMR+08].

¹⁴ This is explained in Sec. 8.10.

¹⁵ This is only mentioned in Sec. 1.5 but fully presented in Sec. 8.10.

This work is more closely related to the work in [SSC⁺05, Chm05], with which the data is shared, and the work in [BK03, KB04, BMR⁺08] for using Multi-agent Simulations. Another closely related work are [Avi06, CS09] where the Cumulative Prospect Theory is used for modelling traffic assignment, even though in both cases microeconomics models are used.

1.8 Text Structure

The remaining is organised into two main parts: theory and application. The theoretical part (from chapters 2 through 4) provides, first, the fundamental theory used along this work: the UT (chapter 2). This chapter is followed by the accepted definition of rationality (chapter 3) and the next chapter presents the core theory defended here: the PT – which is also based on the UT.

In the application part, the contribution of this thesis is presented. The first chapter of this part (chapter 5) presents the Q-Learning algorithm and its modification to incorporate the PT. Then the traffic modelling is presented (chapter 6). In chapter 7, the agent architecture – including the Q-Learning algorithm – is proposed to cope with MASim for discrete choice using the PT. Following this generic and abstract specification comes chapter 8 that presents the evaluation of this thesis. In this chapter the architecture is made concrete and then evaluated with the El Farol Bar Problem modified for the traffic assignment problem. Several experiments were made, presented, and analysed. It also includes the results with a real-world scenario, the city of Burgdorf. The next chapter presents the state-of-the-art and related work (chapter 9). The last chapter discusses the findings and proposes the next steps.

Because some points are left somehow open, such as the specific techniques used for route generation, appendices were necessary. The appendix A contains the different technologies concerning the traffic network representation and its navigation (route generation). The last appendix contains several images that are secondary for the understanding of the text (but if not made available might lead to misunderstandings). It also includes images that otherwise would disrupt the text flow. But when an appendix is necessary in the regular text, it is properly referred. Part I Basis

Chapter 2 Utility Theory

Before starting the formal definitions it is important to emphasise that the Utility Theory [Fis70] (UT) is an attempt to model human choice behaviour. It is also relevant that decision-makers' attitudes toward a set of options is modelled into a function that has this option set as domain and a numerical set, usually \mathbb{R} , as codomain. In this context it is also supposed that the modeller knows the preferences of the decision-makers regarding the choice set, i.e. ranking the options. The modeller must then create a function that assigns a numerical value to each option to reflect the ranking established by the decision-makers for the given choice set. Suppose that an hypothetical choice set $\{A, B, C\}$ is given and a survey that shows people prefer B over C and C over A or formally: $B \succ C \succ A$. Then the modeller must find a function $u : \{A, B, C\} \mapsto \mathbb{R}$ where u(B) > u(C) > u(A). It is irrelevant which particular value $u(\bullet)$ returns for each option, as long as the relations are maintained.

Before introducing the formal definition it is important to keep in mind that the objective is to model human decision-makers and not synthetic beings. This said, the formal structures that support the utility theory must be defined. The first structure is the choice set: the domain of the utility function. This choice set is, in this approach, assumed discrete and countable and each choice has a set of outcomes associated. An example: suppose that the choice set is $\{head, tail\}$ for the flip of a coin and if you choose head and it turns head you win 50\$ then it is said that head has two possible outcomes: 0\$ (if you choose *head* and it turns *tail*) and 50\$ (in case you choose head and it turns head). So, to organise these several outcomes associated with each option, comes a structure called lottery and each single option in the choice set has a lottery that is, according to [Osb03, SLB08], a probability distribution over outcomes. Among the several ways of representing a lottery the structure adopted here is $\{\langle x_0, p_0 \rangle, \langle x_1, p_1 \rangle, \dots, \langle x_n, p_n \rangle\}$, where x_i is the *i*th possible outcome and p_i the corresponding probability. It is also assumed that $\sum p_i = 1$ and $0 < p_i \leq 1$. The lottery is further explored in the next chapter but is introduced here for convenience. Notice though that options and lotteries are conceptually different but they are used interchangeably and therefore bold upper-case letters represent a set of lotteries such as in X, then just X represents the option set (dissociated from the lottery set). The same is valid for an option $x \in X$ and its counterpart element of a lottery $\mathbf{x} \in \mathbf{X}$.

The basic assumption is that – given a choice set and a preference order [Osb03] (which is the choice set ranking) – it is possible to express this ranking using a numerical function (the utility function). The utility function for its turn may need to take into account the options' and individuals' attributes and then uses them to assign a numerical value to each option. One simple example is to survey the people's preferences about different products of the same kind, for instance: shoes x and y from different companies but with the same function (shoes for playing tennis, to give an example). Let $X = \{x, y\}$ be the set of options, then the ranking the possible ranking relations for these options are:

 $x \sim y x$ is no better or worse than y, i.e. the choice between x and y is indifferent (they are equally good).

- $x \prec y$ in this case y is preferred over x, i.e. when choosing between x and y it was observed that the surveyed people preferred y.
- $x \succ y$ this is the contrary of the above relation, i.e. the individuals actually preferred x over y.

According to the UT, the objective is to find the utility function that reproduces these relations. This means that the utility function maps the options to numerical values and these values' numerical relations reflect the preferences. Formally the utility function is expressed in Eq. 2.1.

$$u: X \mapsto \mathbb{R} \tag{2.1}$$

The function $u(\bullet)$ must give a numerical evaluation to each option $x \in X$ and this function is considered a valid model if Eq. 2.2 is true.

$$\forall x, y \in X \mid x \prec y \text{ iff } u(x) < u(y)$$

$$\oplus x \succ y \text{ iff } u(x) > u(y)$$

$$\oplus x \sim y \text{ iff } u(x) = u(y)$$
(2.2)

One consequence of this is that all options must be ranked [Osb03] and the utility function must reflect this ranking through its numerical value [Fis70]. The preferences build, considering the rules above, a partially ordered set, for further analysis of the nature of the utility function it is necessary to introduce the theory of scales of measurements.

2.1 Scales Of Measurement

Because a function returns a numerical value it does not necessarily mean that this value incorporates all properties of its codomain (usually \mathbb{R} or \mathbb{N}). According to the purpose of the function it is possible to classify it using the scales of measurement [Ste46]. This classification system is relevant because it also says which numerical operations are meaningful given the scale. This scale also informs which transformations can be made in the original function without corrupting the functions purpose. The scales according to Stevens [Ste46] are:

- **Nominal** The function can only express equality, i.e. a function of this type can only say if any two elements of the domain (the argument of the function) are equal/equivalent or not. A hash function [CLRS01] is a typical example of a nominal scale, i.e. its value can only express if any two elements are equivalent (assuming a good hash function).
- **Ordinal** A function of this type can express, as in the nominal scale, if any two elements are equivalent or if one of them is higher ranked then the other. The ordering in a deck of cards is an example of such scale (a function gives a natural number to each card). For example, the King is above the Queen and below the Ace, but two Kings of different packs are equivalent.
- **Interval** In this scale, which incorporates the properties of the ordinal scale, it is possible to measure how far the options are from each other. An example of an interval scale is temperature in Celsius or Fahrenheit degrees (not Kelvin), or dates on any calendar, or any other scale with an arbitrary zero. In these examples it is possible to measure the separation between any two options (November 15th is closer to December 15th than to January 17th).
- **Ratio** This is the most complex of the scales. Additionally to the properties of the interval scale, it can also measure how much better an option is compared with any other. The examples are abundant, any physical measurement (with non-arbitrary zero) such as temperature (Kelvin), mass (gram), energy (Joule), and so on. This means that it is possible to measure, for example, if an option is two times better than another.

According to these scales it can be said that the utility function builds an ordinal scale. The reason is because the preferences can only express if any given option is: better, worse, or equivalent to any other. Therefore the utility function (that must be built to reproduce these preferences) can only reflect the same relations observed in the data, i.e. order/rank.

The most relevant conclusion of this is that the utility function, according to the UT, can only be used to rank options. This is not a restriction imposed in this work but imposed by the UT and it also means that any theory derived from the UT has the same limitations, which are imposed by the utility function. A meaningful expansion of the utility function makes necessary to extract/know the desired mathematical properties. To give an example, assuming that an interval utility function is desired, it is necessary to survey how much a person prefers an item over all others. For the shoe problem it corresponds to have both the information if x is better than y, as well as the information of how much it is better than y. In any case, the utility function expressiveness is restricted to the available data (and the relations it establishes).

Because the ordinal scale is simple, compared to the interval and the ratio scales, it also has more freedom in its manipulation. A function in the ordinal scale can be subjected to any affine transformation and this can facilitate the mathematical treatment of such a function. One example of such transformation would be to modify the utility function so that it has the unity interval as its codomain: $u : X \mapsto [0, 1]$. In [vNM07] the utility function is characterised analogous to the heat theory, where a person can tell if it is warmer or colder but not by how much. More precisely, it is addressed, among other places, in Sec. 3.1.2 of [vNM07]

"... In the case of utility the immediate sensation of preference – of one object or aggregate of objects as against another – provides this basis. But this permits us only to say when for one person one utility is greater than another. It is not in itself a basis for numerical comparison of utilities for one person nor of any comparison between different persons. Since there is no intuitively significant way to add two utilities for the same person, the assumption that utilities are of non-numerical character even seems plausible."

2.2 Prerequisites

The prerequisites of the UT are a partially ordered choice set and a function that establishes that ordering. In practical terms it means that the utility function takes into account all relevant attributes and using them reproduces the observed ranking. An attribute is relevant if it is necessary to establish the option ranking. For example in Tab. 2.1, three shoes are compared along with their attributes and the hypothetical observed/surveyed ranking, which is exposed in the last column. If the attributes do not suffice for reproducing the ranking it means that either not all relevant attributes are available, the utility function was not properly chosen, or the problem cannot be modelled using the UT.

Shoe (option)	Price	Colour	Material	Brand	Ranking
a	10\$	Black	Synthetic Leather	X	1st
b	15\$	Blue	Real Leather	Y	2nd
9	13\$	Green	Synthetic Leather	X	3rd

Table 2.1: Shoe example

To express the ranking in the table¹ it is clear that if price (actually $u(x) = 1/x_{price}$) is taken as the only attribute the resulting utility function is invalid. The preferences expressed by this simple utility function are u(a) > u(c) > u(b) and the preferences are: $a \succ b \succ c$ (the utility function violates the $b \succ c$ preference). The colour is the only attribute that could be used to

¹ Recall that the objective of a utility function is to reproduce the ranking using the attributes, assuming that they suffice.

alone establish the preferences. One of the solutions would be to have a utility function that attributes 3 to Black, 2 to Blue, and 1 to Green. But ignoring this particular case, that is likely to be wrong for any extrapolation of this case, some multi-attribute alternatives for the utility function can be realised. One of them would be: $u(x) = 1/x_{price} + 0.01 * f(x_{material})$, where $f(x_{material}) = 1$ if $x_{material} =$ Real Leather and 0 other else. This particular example reproduces the preferences $-u(a) = 0.1, u(b) \simeq 0.07667$, and $u(c) \simeq 0.07692$; which yields u(a) > u(c) > u(b). The $f(\bullet)$ could also have another rule: $f(x_{brand}) = 1$ if $x_{brand} = Y$ and 0 other else.

It can also be seen from this simple example that the function $u(\bullet)$ can have numerous variations and can go under an even higher amount of transformations being still valid. Just to give an example of a family of possible linear transformations: u'(x) = au(x) + b for any $a \in \mathbb{R}^+$ and $b \in \mathbb{R}$. As long as $u'(\bullet)$ is made strictly increasing (considering $u(\bullet)$ the domain) it can be used as a valid transformation. The advantage of such flexibility is to manipulate the utility function in a way that best suits the modeller.

2.2.1 A Note On Parameters And Utility Function

It may cause some misunderstanding that in the example a utility function uses/combines another utility function. This is not a utility combination, since this is not possible, but an entirely new utility function that by chance has a common element with another utility function, which is not facilitated by the fact that some times a single attribute is used as the utility. Attributes, numerical or not, are not utilities, only the utility function has a utility. Therefore the numerical value used to represent an attribute, which contributes to the utility, is called worthiness. This nomenclature is used along the text and is an abstract distinction, to avoid giving the impression that sometimes utilities can be mathematically combined with other utilities, which is the utmost "sin" when manipulating utility functions.

2.3 Utility Determinism

When adopting a utility function to represent the decision process, the decisions are always deterministic, i.e. the option with the highest utility is the option chosen. According to the UT, if the utility function is correctly selected it expresses a sharp logic: "the winner takes all". In the shoe example the preferences say that if a store has the three types of shoes the costumer will always choose the shoe a. If the store does not have a then the costumer will go for the second best (shoe b) and then, as a last option, the shoe c.

A consequence of this is that the utility function is personal, i.e. each person has its own evaluation process – for this reason it is also called Subjective Utility Theory [Fis81] (SUT). Moreover, according to the UT if the attributes are kept constant the decision does not change. The only possibility for change is if the attributes or parameters of the utility function change, but the function itself is fixed and constant. Among the parameters of the utility function one can find socio-economic attributes, such as monthly income, or preferences such as: "if the individual already has the pair of shoes in question, prefer a pair of another type". Nevertheless, the changes occur exclusively in the parameters and never in the function. The parameters taken into account are fixed too.

Of course this discussion refers to the theoretical base of the UT, stating the axioms and assumptions over which it is built.

2.4 General Properties

The UT has some general properties, some of them were already mentioned and some are derived from its formalisation.

Partial Ordering The preferences build a partially ordered set of options and so does the utility function.

- **Ordinal Scale** Because the preferences do not express more than the above mentioned partial ordering, the utility function has only the \langle , \rangle , = as valid comparison operators over the utility value, corresponding to the ordinal scale explained in Sec. 2.1.
- **Compensatory Effect** If the model does not explicitly incorporate restrictions then it is compensatory by nature, i.e. the worthiness of one attribute can compensate the lack of another. In the shoe example, the lack of real leather in shoe c can be compensated by a lower price, i.e. the price worthiness can compensate the lack of worthiness of the material.
- **Immutability** This is more an axiom than a property. It says that the utility function is fixed for each individual. Therefore, if the model does not reproduce the choices observed it means that the model is wrong, not that the person violates its internal utility function.
- **Determinism** The choice is always made picking the best option (according to the utility function). If the model reproduces some of the choices but not all of them it means that the model is wrong and probably does not take into account all relevant parameters (recall the previous property).
- Laplacean Demons A Laplacean Demon, a prerequisite to conform the utility theory [Wim76], is a term used to call an entity that has unlimited resources and computational power. Then because for the UT the complete ranking is necessary, it results in modelling persons as Laplacean Demons because an individual is assumed to always calculate the utility of all options and to make a perfect choice.

2.5 Modelling Using The Utility Theory

From the issue just mentioned in the previous section, follows that a utility function is inevitable to have errors in the model when modelling. The reason lies in the capability of capturing all relevant parameters as well as the individual preferences (the individual/personal utility function). Because it is almost impossible to collect all necessary data, just a partial survey is made. This is the first drawback: the model does not have all necessary variables. Another problem is to model all individual preferences. Usually a (generalised) utility function is sought that can represent all or a class of individuals, thus the second issue: the model does not model an individual but a group of them.

The common praxis in modelling decision-making processes using a utility function is twofold. First, the necessary option/environmental attributes are collected/surveyed (building a data structure similar to the one in Tab. 2.1). Second, the socio-economic attributes of the decision-makers being modelled are surveyed and if significant differences among them are observed then probably a segmentation of the individuals (in groups) is worth investigating (meaning multiple utility functions). Having all the data is the first step in modelling; then the contribution of each different attribute for the utility must be investigated (usually guessed). A common arrangement is the linear combination of all attributes (as in the utility function for the shoe example).

Once the relationship among the attributes is determined, they are expressed in a crude form as in Eq. 2.3, where the functions $f(\bullet)$, $g(\bullet)$, and $h(\bullet)$ are fixed/predetermined by the modeller and returns the worthiness of each attribute. In the shoe example they where: $f(x_{price}) = 1/x_{price}$, $g(x_{brand}) = 0$, and $h(x_{material}) \equiv 1$ iff $x_{material} =$ Real Leather. This leaves some room for calibration/estimation: the values of β_{price} , β_{brand} , and $\beta_{material}$. Because these values depend on the individuals they are left to be calibrated/estimated according to the data.

$$u(x) = \beta_{price} f(x_{price}) + \beta_{brand} g(x_{brand}) + \beta_{material} h(x_{material})$$
(2.3)

Therefore, a calibration/estimation process is necessary to "discover" the "betas" in order to minimise the error between the data and the results of the model. This means that for the above example, a vector $\boldsymbol{\beta} = [\beta_{price}, \beta_{brand}, \beta_{material}]$ is searched that fits the data best. Again using the shoe example, the solution is not unique and among them is: $\boldsymbol{\beta}^* = [1.0, 0.0, 0.01]$ – which is absolutely accurate for the data provided and thus an optimal $\boldsymbol{\beta}$.

2.5.1 Calibration Methods

It is necessary to state that this is not an evaluation method, even though the evaluation function is usually the fitness function for the calibration procedure. Among the several calibration methods one element is constant: they all need a fitness function, i.e. a quality measurement (a function that can determine if β' is better than β''). The two most used fitness functions are the Mean Squared Error [Bie05] (MSE) and the Likelihood [Fis22] function (or the Maximum Likelihood Estimator, MLE for short). Both have advantages and disadvantages and both are biased.

It is argued [Ber93] that the MSE is biased because, among other things, it weights large errors more than small errors. This means that it is not suitable for a model that must be calibrated to reproduce a behaviour instead of an overall fit. An example would be to reproduce the behaviour expressed in a data set. In this case it is better to "hit" most of the points, even if a small amount of them have huge errors, than "almost hit" on all of them.

The main criticism of to the MLE is that it does depend on large amount of data to reduce its bias. One simple example is that if from n numbers (homogeneously distributed from 0 to n-1) just one is drawn then the MLE will not considered the existence of other possible n-1 elements and estimate the mean as the value of this one value. It is said [Edw92] that: "the likelihood function depends only on what actually happened, and not on what could have happened."

2.5.2 Non-Captured Parameters

Because of the restrictions in collecting the data and modelling each individual afterwards the model usually does not feature all relevant attributes/parameters. This is considered an error in the model. Nevertheless, there is a technique that can be used to reduce this error. It relies in a generalisation about the attributes not captured. The technique consists in separating the utility function into two parts, as in Eq. 2.4 where v(x) is equivalent to the u(x) in Eq. 2.3. It contains the part of the utility function that can be specified by the data, called the observable part of the utility. This leaves the e(x) part to be modelled that is called the unobservable part of the utility. It means that all parameters and attributes incorporated into e(x) are not present in v(x) and are still relevant for the decision model, i.e. can influence the option ranking. But this part, as the name says, is not observable (either it cannot be measured or it is left out for convenience, to reduce the model's complexity).

$$u(x) = v(x) + e(x)$$
(2.4)

Even though it is not observable, some assumptions might be made about it. If e(x) cannot be specified/observed but its behaviour does and this behaviour is independent from x then this behaviour can be incorporated into the model. This term accounts for the violations in v(x), i.e. the cases where although v(x) establishes a ranking this very ranking is not followed by some individuals. For example, if v(x) says that a is preferred over b but it is observed that in some cases b is preferred over a, then this is called a violation of the utility function. The violations are then attributed to e(x) and because e(x) was observed to not depend on x then its behaviour, called ε , is added to the model, as in Eq. 2.5

$$u(x) = \beta_v v(x) + \beta_\varepsilon \varepsilon \tag{2.5}$$

Note that Eq. 2.5 has also two additional calibration parameters β_v and β_{ε} . These two parameters are there because it is easier to give either v(x) and ε^2 a fixed form and then manipulate the contribution of each. This gives the advantage of comparing the ratio of information and error,³ v(x) and ε respectively. Recalling once again the shoe example, its crude utility function (with error) will then be in the form of Eq. 2.6.

$$u(x) = \beta_{price} f(x_{price}) + \beta_{brand} g(x_{brand}) + \beta_{material} h(x_{material}) + \beta_{\varepsilon} \varepsilon$$
(2.6)

² The ε is a Probability Distribution Function (PDF) such as the N(0, 1).

 $^{^{3}}$ For this it is assumed that both codomains are in comparable ranges.

The problem in incorporating the error term ε is that the calibration/estimation process must be aware of it and deal with it too. The main issue is that ε has a stochastic behaviour and the assumptions about this behaviour influence the type of calibration used. The reason lies on the calibration methods that can handle just some types of Probability Distribution Functions (PDFs). This topic will again be addressed when discussing econometrics in the next chapter.

2.6 Summary

In this chapter the following main ideas were presented. The utility theory needs a utility function and the numerical values of this function must build the same ranking as the one established by observed preferences. This means that Eq. 2.2 must be true for a function $u(\bullet)$ to be a valid model for the decision process being modelled. The values of the utility function may not be used to perform comparisons or operations that do not correspond to the ones allowed by the ordinal scale (Sec. 2.1). The utility theory assumes some properties to be hold (Sec. 2.4) which are: partial ordering, $u(\bullet)$ builds an ordinal scale, probabilities and outcomes are compensatory, the $u(\bullet)$ is immutable, the choices are deterministic (the winner takes all), and the beings being modelled are assumed Laplacean Demons (they have perfect knowledge and unrestricted resources).

Moreover, some aspects of the decision process may remain non-captured by the model and for them a stochastic behaviour is assumed (Sec. 2.5). This aspects are called non-captured parameters and they must be estimated for the model to be valid, which requires a calibration process (Sec. 2.5.1).

Chapter 3

Rationality And Expected Utility Theory

In the previous chapter the general lines of the utility theory were discussed with a simple example. In this chapter this theme is deepened to present the fundamentals of rationality and its developments.

The basic concept behind rationality is utility maximisation. According to the rational behaviour and utilitarianism human beings, when confronted with a decision problem, evaluate the utility of each option and choose the one with the highest utility. This decision is always, as explained in the previous chapter, deterministic and as long as the environment does not change (the attributes remain the same) the choice is then the same.

3.1 Concepts Review

From the previous chapter it is important to bare in mind the following key issues. First that the utility function $u: X \mapsto \mathbb{R}$ builds an ordinal scale¹, i.e. the values yield by $u(\bullet)$ can only be used to rank the options $x \in X$, nothing more. Second, the ranking built by the $u(\bullet)$ is deterministic, i.e. if u(a) > u(b) the choice will always be $a \in X$. Third, the preferences must be known for all options, i.e. the choices build an partially ordered set. The last point is that the nature of the function $u(\bullet)$ does not matter as long as the Eq. 2.2 remains true, meaning that it reproduces the observed preferences for all options.

3.2 Saint Petersburg Paradox

The problem with the utility maximisation theory is that it is limited by one choice, i.e. it does not account for a sequence of choices. This was noticed by Bernoulli [Ber38, Ber54] when studying a hypothetical problem, called the Saint Petersburg paradox [Mar04]. The paradox is in establishing a price for a hypothetical game; whose expected accumulated reward (mathematical expectancy) is infinity but in reality could only request a small entry fee, when compared with the final expected prize. The hypothetical game is as follows: the game participant must pay an initial fee x and a fair coin is tossed. According to the rules: if the coin turns tails then the participant win the pot, otherwise the coin is tossed again (without an additional fee) and the pot is doubled. The initial pot v is lower than the initial fee x, which means that if the participant wins early in the game he/she actually looses money. The game implies that for k tosses the pay-off will be $v2^{k-1}$. The expected final pay-off extrapolation is in Eq. 3.1.

 $^{^{1}}$ The scales of measurement are the subject of Sec. 2.1.

$$E = \frac{1}{2}v2^{0} + \frac{1}{4}v2^{1} + \frac{1}{8}v2^{2} + \dots$$
$$= \sum_{k=1}^{\infty} \frac{1}{2^{k}}v2^{k-1}$$
$$\lim_{n \to +\infty} E \Rightarrow \lim_{k \to +\infty} v \sum_{k} \frac{1}{2} = +\infty$$

In the paradox, with this simple extrapolation, it can be seen that the game has unlimited worthiness, i.e. the fee should be set to the expected value of the pay-off, i.e. the maximum amount of money available (due to lack of a better practical definition of infinity monetary value). But it turns out that people were not willing to pay the expected game reward as fee but instead a small multiple of the initial pot. The paradox is in the fact that people apparently could not realise the worthiness of the game, i.e. they were "miscalculating" the game's utility and therefore the paradox. What Bernoulli [Ber38, Ber54] realised is that people do not maximise the utility, but the utility's accumulated diminishing worthiness. This means that money, in this example, loses its worthiness with time and effort. For Bernoulli it meant that people's time and effort have costs that diminish the game utility.²

After Bernoulli, rational behaviour is characterised by the maximisation of the diminishing utility.

3.3 Expected Utility Theory

(:

k

When dealing with decision choices approached by the rational behaviour the basic structure is the lottery, as mentioned in the previous chapter. This means that a decision problem is reduced to choose among several lotteries. The lottery structure was already introduced in the previous chapter but in Eq. 3.1 it is presented again. It is important to notice that hereafter lotteries and options are no longer disentangled but used as synonyms³ because an option without the corresponding lottery cannot be evaluated using the Utility Theory (UT). Another point that may seem confusing is the use of the word outcome/pay-off for values that are actually possible utility values. They are indeed not conceptually equivalent because a utility is by definition a single value that represents the worthiness of an option regardless its possible outcomes. In essence the utility represents the way in which one or more outcomes are combined into a single numerical value. Therefore, even though some outcomes have the same numerical value of the corresponding utility they are not called utilities.

$$\mathbf{x} = \{ \langle x, p \rangle \mid x \in \mathbb{R} \land p \in (0, 1] \}$$

$$\sum_{x, p \rangle \in \mathbf{x}} p = 1$$
(3.1)

Then a lottery is a set of pairs $\langle x, p \rangle$ where x represents a possible outcome for the lottery/option **x** and p a probability (**x** is a probability distribution function over outcomes [Osb03, SLB08]). The values of x and p are later explained, first it is necessary to differentiate between two types of lotteries:

Simple Are lotteries with only one pair, i.e. $\mathbf{x} = \{\langle x, 1 \rangle\}$, where the outcome x becomes the utility of the lottery/option \mathbf{x} .

 $^{^{2}}$ He rationalised that unlimited money does not worth infinity if it costs infinity time and effort to reach.

 $^{^{3}}$ Here the terms lottery and option are used interchangeably but they are actually different. An option is something more abstract and represents an element without any numerical value associated with it. A lottery on the other hand is the numerical transformation of an option, i.e. a concrete option with its possible outcomes. The reason why both are used as synonyms relies on the fact that only the numerical characteristic of an option is of interest in this text.

Mixed Different from the simple lotteries, the mixed lotteries have more than one pair, i.e. $|\mathbf{x}| > 1$ and the utility must be calculated aggregating the several possible outcomes.

A simple lottery has the same aspect of the alternatives in the shoe problem from the previous chapter (presented in Tab. 2.1). There, the x value corresponds to the pay-off value⁴ $x = p(o) = 1/o_{price} + 0.01 * f(o_{material})$ where o is the option itself and x the corresponding outcome. To make it clearer, the shoe choice problem would be to choose among the lotteries in the set **X** specified in Eq. 3.2.

$$\mathbf{X} = \{\mathbf{x}_{a}, \mathbf{x}_{b}, \mathbf{x}_{c}\}$$
(3.2)
$$\mathbf{x}_{a} = \{\langle p(a), 1 \rangle\} \Rightarrow \mathbf{x}_{a} = \{\langle 0.1, 1 \rangle\}$$
$$\mathbf{x}_{b} = \{\langle p(b), 1 \rangle\} \Rightarrow \mathbf{x}_{b} = \{\langle 0.07667, 1 \rangle\}$$
$$\mathbf{x}_{c} = \{\langle p(c), 1 \rangle\} \Rightarrow \mathbf{x}_{c} = \{\langle 0.07692, 1 \rangle\}$$
$$p(o) = 1/o_{price} + 0.01 * f(o_{material})$$

3.3.1 von Neumann's Formalization

Before von Neumann⁵ [vN28, vNM07] (among several other improvements brought by his work) the Expected Utility Theory (EUT) could only model problems with simple lotteries. This means that if any of the variables in the model had a stochastic behaviour it could not be properly modelled. To give an example, if in the shoe model the price varies according to the store where it is bought, then the shoe *a* instead of a fixed price 10\$ has a varying price with the following distribution: 10\$ for 90% of the stores and 12\$ for the remaining 10%. This makes $\mathbf{x}_a = \{\langle 0.1, 0.9 \rangle, \langle 0.08333, 0.1 \rangle\}$.

In the UT the ranking is made based on a utility function that returns single values (one for each option/lottery) and these values are then compared with each other, where the highest value results in the choice. The utility of a lottery, for the simple lotteries, is the outcome/pay-off of the options, i.e. $u(\mathbf{x}) = p(o)$. This is not valid for the mixed lotteries because the utility theory demands a function that returns a single value, where the highest value yields the choice. For this problem von Neumann proposed to aggregate the different utilities into the mathematical expectancy over the possible outcomes (the multiple x values in the lottery). Then the utility of a mixed lottery is given by Eq. 3.3.

$$u(\mathbf{x}) = \sum_{\langle x, p \rangle \in \mathbf{x}} xp \tag{3.3}$$

Then for the shoe example, the new utility of the stochastic a is $u(\mathbf{x}_a) = 0.09833$ instead of the previous (for the non-stochastic version) $u(\mathbf{x}_a) = 0.1$. For this assumption (the utility of a mixed lottery to be equivalent to the mathematical expectancy) some restrictions are imposed to the model, so that it remains a valid utility model. Because in [vNM07] the axioms are restricted to a more basic mathematical structure⁶ and in [Osb03] (also used as reference) the concepts are only informally presented, the axioms given below were extracted from [SLB08].⁷

(3.1.1) Completeness The preferences among all options must be known, i.e. given any two lotteries $\mathbf{x}, \mathbf{y} \in \mathbf{X}$: $\mathbf{x} \succ \mathbf{y}, \mathbf{x} \prec \mathbf{y}$, or $\mathbf{x} \sim \mathbf{y}$.

(3.1.2) Transitivity If $\mathbf{x} \succ \mathbf{y}$ and $\mathbf{y} \succ \mathbf{z}$ then $\mathbf{x} \succ \mathbf{z}$.

⁴ Notice that here it is called pay-off and not utility, since utility is the value resulting of the evaluation of the lottery and not the outcome in it. It is important to say that in the shoe example of the previous chapter the function was indeed the utility function but here it is used to yield the outcome and therefore it was renamed to $p(\bullet)$. ⁵ The original idea and formalism [wN28] was from you Neumann alone (mentioned in the introduction

 $^{^{5}}$ The original idea and formalism [vN28] was from von Neumann alone (mentioned in the introduction of [vNM07]) and Morgenstern helped to expand the formalisation of von Neumann into a more didactic form of a book [vNM07], this is why in this text the credit is sometimes given to von Neumann alone and not both.

⁶ See Sec. 3.6 from [vNM07] for further reading.

 $^{^7}$ The numeration give here is the same as it appears in [SLB08].

- (3.1.3) Substitutability If any two outcomes are equivalent, i.e. $\langle x_i, p_i \rangle, \langle x_j, p_j \rangle \mid x_i = x_j$ then if $p_i = p_j$ the pairs may be used interchangeably: $\langle x_i, p_i \rangle \equiv \langle x_j, p_j \rangle$.
- (3.1.4) **Decomposability** It says that for two given lotteries where all pairs are equivalent then the lotteries are equivalent as well: $\forall \langle x_i, p_i \rangle \in \mathbf{x}, \forall \langle x_j, p_j \rangle \in \mathbf{y} \mid x_i = x_j \land p_i = p_j \Rightarrow \mathbf{x} \sim \mathbf{y}.$
- (3.1.5) Monotonicity Given any two outcomes where $x_i \succ x_j$ and two mixed lotteries with only these outcomes, then the agent will choose the lottery with the highest probability assigned to the preferred outcome. Let $\mathbf{x} = \{\langle x_i, p \rangle, \langle x_j, 1-p \rangle\}, \mathbf{y} = \{\langle x_i, q \rangle, \langle x_j, 1-q \rangle\}$, and p > q, then $\mathbf{x} \succ \mathbf{y}$.
- (3.1.6) Continuity It states that if three simple lotteries exists such that $\mathbf{x} \succ \mathbf{y} \succ \mathbf{z}$ then $\exists p \in [0,1] \mid \mathbf{y} \sim \{ \langle x_{\mathbf{x}}, p \rangle, \langle x_{\mathbf{z}}, 1-p \rangle \}.^{8}$

Special attention must be given to a derived axiom from axioms 3.1.3 and 3.1.4 that is here called independence. If the axioms 3.1.3 and 3.1.4 are accepted then the following must be too. Given two equivalent options/lotteries $\mathbf{x} \sim \mathbf{y}$ then their utilities must be equivalent as well, i.e. $u(\mathbf{x}) = u(\mathbf{y})$ (form the previous chapter). Since axiom 3.1.3 says that if two outcomes are equivalent then they stay equivalent when they have the same probability of happening. It derives that $pu(\mathbf{x}) = pu(\mathbf{y})$, being $p \in [0, 1]$. Moreover, if the probability of all outcomes in two given lotteries are equal and their probability as well (axiom 3.1.4) then the following equivalence is true: $\mathbf{v} \sim \mathbf{w} \mid \mathbf{v} = \{\langle u(\mathbf{x}), p \rangle, \langle u(\mathbf{y}), 1 - p \rangle\} \land \mathbf{w} = \{\langle u(\mathbf{x}), p \rangle, \langle u(\mathbf{y}), 1 - p \rangle\}$ for any $p \in [0, 1]$. Then from 3.1.3 comes that $\mathbf{w} \equiv \{\langle u(\mathbf{y}), p \rangle, \langle u(\mathbf{x}), 1 - p \rangle\}$, since the outcomes are equivalent. This means that for any $p \in [0, 1]$ it is true that $u(\mathbf{v}) = u(\mathbf{w})$ and since the utility of a mixed lottery is $\sum xp$ (Eq. 3.3) it follows that $pu(\mathbf{x}) + (1 - p)u(\mathbf{y}) = pu(\mathbf{y}) + (1 - p)u(\mathbf{x})$. This implies that equivalent options/lotteries can be combined for any probability $p \in [0, 1]$ regardless the region (in the interval), i.e. using a p = 0.1 is as good as p = 0.9.

For completeness, it is also assumed that when referring to the utility of a lottery the necessary attributes were taken into account for calculating the outcomes x in the pairs $\langle x, p \rangle$. The utility from Eq. 3.3 will be renamed as $eut(\mathbf{x})$ to refer to the specific utility calculated according to the EUT.⁹

3.4 Discrete Choice Analysis

The Discrete Choice Analysis (DCA) is the statistical sub-field that provides the means for fitting/calibrating a choice model to a data set. Because, as argued in Sec. 2.5.2, usually not all parameters are used in modelling the utility function, it means that the modelling technique sees the utility function in two parts. First, the nature of the known utility – v(x) from Eq. 2.5 – and second the error term ε . A common problem is that v(x) is correlated among the options, i.e. some attributes, or attribute values, are shared by more than one option. Moreover, the options must be assumed independent from each other, i.e. a random variable where the stochastic element is given by the ε . This also means that the option random behaviour is given by the probability distribution function associated with ε .

Because the options must be random variables the chosen DCA model must correctly extract the correlation among the options to conform the independence of each option. Then it means that a DCA model extracts the correlation existing in v(x) and makes the options random variables, which have the stochastic behaviour expressed by ε . Examples of such models are listed below.¹⁰

Logit Developed by Berkson [Ber44], extended by Luce [Luc59], and improved by McFadden [McF74] into the Multi Nominal Logit (MNL) to cope with several options (the original Logit can

⁸ Since simple lotteries have only one outcome (with probability 1) the notation x_x refers to this single outcome unambiguously.

 $^{^{9}}$ A function to be called utility function must not conform the von Neumann-Morgenstern [vNM07] axioms but give a ranking to the several options in the choice set. Therefore the shift to a new nomenclature when referring to specific ways for calculating the utility of a lottery.

¹⁰ For a review and comparison of the different models for route choice can be found in [Ram02].

model the choice between only two options). The Logit model assumes that the ε in the data is a logistic distribution.

- **Probit** A more complex model than the Logit, developed by Bliss [Bli34a, Bli34b], fits the data to a standard Normal curve.
- **Nested-Logit** First derived by Ben-Akiva [BA73], it overcomes a limitation of the MNL models and permits the model to organise the options' correlation in "nests".
- **Cross-Nested Logit** It improves the original Nested-Logit by allowing an option to belong to multiple nests. It was developed by Vovsha [Vov97].
- **Path Size Logit** Developed by Ben-Akliva and Bierlaire [BAB99], it captures the correlation among the options assuming that they share some parts (of the path/route).¹¹ Thus the correlation depends on the size (of the path/route) shared and therefore the name path size.
- Mixed Logit This is a further improvement over the MNL model developed by McFadden and Train [MT00] that uses a restricted form of correlation extraction available in the Probit model.¹²

The differences in the above mentioned models are twofold: how the correlation among the options is captured and which probability distribution function is used/assumed for ε . The latter is usually to choose to the most convenient for the calibration method. For the correlation, the objective is to reduce the options to be independent and identically distributed (i.i.d or iid) – also referred as independent from irrelevant alternatives (i.i.a or iia). This property demands the choice's stochastic behaviour to be reduced to a random variable, so it conforms the utility maximisation principle [Mar60].

These models must then be calibrated to fit the data and the mathematical framework for calibrating DCA models is called econometrics [Woo03]. These methods are used to fit a DCA model to a data set. This combination of econometrics for DCA models calibration is called the Random Utility Theory [Man77] (RUT, reviewed in [Fis81]), whose models are called Random Utility Models (RUM).

3.5 Summary

This chapter extends the ideas presented in the previous chapter and formalise the concepts there presented. The concept of lottery (presented in the previous chapter) is further explained and presented under the formalism from von Neumann and Morgenstern [vNM07]. The Expected Utility Theory (EUT) is presented (Sec. 3.3) as well as its formalism (Sec. 3.3.1). The main consequence of this formalism is summarised by the final paragraphs in Sec. 3.3.1, which is called the independence. This is the main concept exploited in the next chapter when presenting the Prospect Theory (PT).

The last section of this chapter shows an overview of the main theories derived from the EUT. There the main issue is how to organise the options so that they remain independent from each other, i.e. they remain i.i.d. (a fundamental requirement for the UT). This means that these theories main objective is to capture the correlation among the options so that they are, after removing the correlation, independent from each other.

¹¹ This model was developed for route choice modelling.

¹² The Probit is considered the most complete and complex way of extracting the correlation among the options, but the computational requirements for calibrating such a model are considered prohibitive and therefore a restricted form of Probit is used in the Mixed Logit model.
Chapter 4 Beyond Perfect Rationality

The term non-rationality is here used to define any theory that does not follow the rational behaviour defined in the previous chapter. This means that a theory that does not follow the axioms in Sec. 3.3.1 (to be specific) will be here called non-rational.

The term bounded-rationality [Sim55, Sim56] is avoided because it is usually associated with a rational behaviour over imperfect knowledge or limited computational resources. It is however understood that bounded-rationality is the superset of what will be here defined as non-rational. Under the umbrella of bounded-rationality are not only imperfect knowledge and limited computational resource but also behaviours that do not converge to the rational behaviour. Therefore the term non-rationality seemed more appropriated to refer to the models that even under perfect knowledge and unlimited resources do not converge to the rational behaviour with probability 1 – to be mathematically precise. The theory here adopted to express non-rationality is the Prospect Theory [KT79] (PT). Another reason is in the evidences shown in [DMKSD06, KF06] that support the PT as a valid hypothesis for the human decision behaviour.

4.1 Concepts Review

In this chapter the critics to the last chapter are presented. Therefore it is important to keep in mind the general properties of the utility function (Sec.2.4), specially the assumption that people are instances of Laplacean Demons. Then when discussing the Prospect Theory (PT) in Sec. 4.5 the independence property derived from the Expected Utility Theory (Sec. 3.3.1) is the main issue. This property says that when aggregating probabilities and outcomes the region where the probabilities are does not matter, i.e. 0.1 = 0.4 - 0.3 has the same effect as 0.1 = 0.9 - 0.8. In other words, a 10% probability near zero is the same as near 100%. This is not true for the PT.

In this chapter the lottery concept (presented in chapter 2 and expanded in chapter 3) is discussed. This discussion leads to the presentation of what a prospect is (Sec. 4.6) and how it differs from a lottery.

4.2 Criticisms To The Rational Model

Among the critics of rationality is Simon [Sim55, Sim56] who claims that the rational model is not appropriate to describe the human behaviour. He says that the pure rational model – that requires a Laplacean Demon (Sec. 2.4) – is not feasible¹ and proposes the idea of boundedrationality [Sim82]. The term coined by Simon [Sim82] refers to the limits for the human beings to apply the rational thinking. He argues that a person has neither unlimited resources nor unlimited computational power and therefore is restricted to its own limitations in calculating the option's

¹ An instance of the Laplacean Demon requires unlimited time and unlimited resources, which is definitely not the case for human beings in the general case and rarely in any particular problem instances.

utility. This limitations can be of mathematical complexity, i.e. the person does not possess the mathematical skill to evaluate the options; or it can be a psychological interference that prevents the individual from calculating the utility correctly. What is observed in the praxis, specially in the AI field, is that individuals are modelled having partial/local knowledge and utility functions with mathematical complexity restrictions. But if the problem is simple enough (does not exceed the individual's limitation) then the model is strictly rational, i.e. it complies the rational behaviour. In the following sections a short list is given for some of the issues that can induce a deviation on the rational behaviour. For an extended list of them please refer to [Kah02]; for an informal presentation to [Ari08]; and for criticisms in the context of the economic modelling to [McF99]. To finalize a quotation from Simon [Sim86] (page 223) about the rational behaviour assumption:

"First, I would recommend that we stop debating whether a theory of substantive rationality² and the assumptions of utility maximization provide a sufficient base for explaining and predicting economic behavior. The evidence is overwhelming that they do not."

4.2.1 Bi-Parted Decision-Making System

According to several researchers – among them Epstein [Eps94], Sloman [Slo96, Slo02], and Tversky and Kahneman [TK83, KF02, KF05] – empirical evidences suggest that the human decisionmaking process is bi-parted, i.e. two different reasoning processes are used when a decision must be made. The two systems received several different names but the characteristics of them are in agreement across the researchers. For convenience the Kahneman and Tversky's nomenclature is used: System 1 and System 2. The first part (System 1) is the intuition level and the second the reasoning level (System 2).

The intuition level is responsible for a quick response, sometimes referred as a "quick-anddirty" response. This response is based on information that is easily available and it depends on the familiarity of the decision-makers with the choice problem. The decision at this level is described as effortless. When the intuition solves a problem the process is sometimes described by the respondents as: "the answer just popped-up into my mind".

The reasoning level, on the other hand, demands effort and time. It is a cumbersome process. At this level the choices are individually evaluated, compared, and several parameters and variables are taken into account. The individuals tend to evaluate all options and consciously to choose the best of them.³ It is said "tend" because the individual may not be aware of all options (local knowledge).

What was noticed is that the individuals have the tendency to first try to solve problems using intuition, trading a possibility of better gains for time and effort (that would be demanded to solve the problem using reasoning). The two systems are not independent, they work in a hierarchy, where the reasoning level (System 2) monitors the intuition. The System 2 is then responsible for correcting the System 1, if the estimated outcomes do not satisfy the individual's constrains. In this case the problem is reviewed and the choice is made by the System 2. But as experienced with the problem the person becomes as reliable the System 1 becomes.

Rubinstein [Rub07] observed that the response time and decisions' quality suggest that the quick decisions, done by intuition, are of lower quality than the ones with larger response times, done by reasoning. About the implicit learning, in [Reb89] some explanation for this type of learning is given, suggesting that implicit learning (by intuition) is observed to be more robust than explicit learning.

One consequence of this is that sometimes the individuals fails to notice changes on the options (change on the attributes) or options set (new options are added or old options are removed). In this regard this behaviour reproduces the bounds specified by Simon, where the person is biased by its own knowledge, whose capacity is limited.

 $^{^{2}}$ Referring to the SUT.

 $^{^3}$ The best for the individual is not necessarily the rational choice.

4.2.2 Reference Dependency And Status Quo

The reference dependency, also known as *status quo* [KT79, GSP⁺89], is the point that determines if an outcome is a gain or a loss. The principle is that when faced with a decision problem the person analyses the option's profit or loss based on her/his current state, the *status quo*. This reference works as a behaviour trigger. The usual observed attitudes towards the options is how risky one behaves depending on the *status quo* that is, according to [KT79], riskier in the losses than in the gains.

To give an example; a risky game is proposed: 25% of winning 100; 50% of winning 10; and 25% of losing 120\$. In this game it is asked if the participant is willing to participate or not. It is also asked if the choice is affected by informing that the current wealth is -50\$ or 50\$. It is expected that when the current wealth is -50\$ people are more inclined to accept the game and when it is 50\$ not. The game has however an expected value of zero ($100 \cdot 0.25 + 50 \cdot 0.10 - 120 \cdot 0.25$). It seems that the individuals see a chance of at least diminishing their losses when the wealth is said -50\$; and see the chance of being in the loss, after the game, too threatening when the wealth is said 50\$.

Investigating this influence in traffic, Horowitz [Hor84] observed the user-equilibrium of a twolink network. There, deterministic and stochastic models are used for link cost evaluation that had had different initial states (link loads). It is observed that the equilibrium of the stochastic models can depend on the initial states (the reference) and achieve different equilibrium states or even oscillating states (oscillating link loads). This can be interpreted as an effect of the initial state acting as a reference and then creating a bias for the subsequent decisions.

4.3 Bounded Rationality

Bounded rationality has several definitions, none concrete, and this usually leads to using a narrow interpretation instead of the broader. To state clear what is meant by bounded rationality the abstract definition from Simon [Sim92] (pages 353–354) is quoted below:

"...One procedure already mentioned is to look for satisfactory choices instead of optimal ones. Another is to replace abstract, global goals with tangible subgoals, whose achievement can be observed and measured. A third is to divide up the decisionmaking task among many specialists, coordinating their work by means of a structure of communications and authority relations. All of these, and others, fit the general rubric of "bounded rationality", and it is now clear that the elaborate organizations that human beings have constructed in the modern world to carry out the work of production and government can only be understood as machinery for coping with the limits of mans abilities to comprehend and compute in the face of complexity and uncertainty."

From this definition it can be said that humans do not optimize but satisfy their needs through the realisation of the steps that lead to this satisfaction. This can be also read as saying that people has limitations on their knowledge (limited by their own observations) and computational power (compute only until the satisfaction level is achieved). Gigerenzer [GG96] argues, on the other hand, that this is a rather too narrow interpretation of what Simon had meant with bounded rationality. He also says (page 651): "For most part, however, theories of human inference have focused exclusively on the cognitive side, equating the notion of bounded rationality with the statement that humans are limited information processors, period." What is meant by Gigerenzer is that modelling human decision-maker only by limiting the knowledge and computational power is not the answer.

This is why here the term bounded rationality is avoided and the term non-rationality used instead. Just limiting the information processing mechanism is not the solemn condition for bounded rationality but just one of them.

Problem 1:					
Opt.	Outcome	Prob.	eut(ullet)	Pref. (in $\%$)	
	2500	0.33		18	
A	2400	0.66	2409	10	
	0	0.01			
В	2400	1.0	2400	82	
Problem 2:					
Opt.	Outcome	Prob.	eut(ullet)	Pref. (in %)	
С	2500	0.33	825	83	
	0	0.67	020		
ת	2400	0.34	816	17	
D	0	0.66	010	11	

Table 4.1: Modified Allais problem

In [Sim79] Simon says that the classical models (rational models) must be replaced by models that take into account: our imperfect knowledge about the alternatives and their consequences, our limited computational power, our lack of a consistent utility function for heterogeneous alternatives, and our uncertainty about exogenous events. This definition goes clearly beyond just limiting the knowledge about the world and limited complexity of the utility function.

4.4 Systematic Rationality Deviation

In Sec. 4.2 some aspects of human psychology were presented that interfere in the options' evaluation. These issues are more limitations, as Simon specified, than systematic deviations; they vary with the problem formulation and current status (wealth), which are represented by the individual's parameters. But two paradoxes show that rationality is systematically violated even when the problem is simple enough to not exceed the mathematical skills of the individuals. The paradoxes are formulated to avoid the *status quo* influence as well, by dealing only with gains.

These two formulations are the Allais [All53, AH79] and the Ellsberg [Ell61] paradoxes. The Allais paradox shows that people do not use the $eut(\bullet)$ function (Eq. 3.3) to rank the options and so they violate consistently and systematically the rational choice theory. The Ellsberg paradox is more complex in its formulation but shows that the individuals change their estimates during the choice problem, i.e. they changed their utility functions. These paradoxes show that the rational hypothesis is not an adequate model for describing human behaviour.

4.4.1 The Allais Paradox

The Allais paradox demonstrates that the rational behaviour is violated in a simple monetary choice problem. To illustrate it, in Tab. 4.1 a modified version of the Allais paradox (from [KT79]) is shown. The modified version is chosen because for this version the corresponding data is available.⁴ The problem (Tab. 4.1) was proposed as follows: to each person two choice problems were given: first to choose between lotteries/options A and B (Problem 1:) and then between C and D (Problem 2:). There, the column "Outcome" corresponds to the monetary gain and the column "Prob." the corresponding probability for that outcome. To the participants only the first three columns were presented and their answers collected (aggregated in the column "Pref. (in%)"). For reference, in the column "eut(•)" is the rational utility – calculated using the Eq. 3.3.

 $^{^4}$ The original Allais problem was informally proposed in a conference and the data was not published, only the consequences.

Problem 1:					
Choice	Ball	Outcome	Odds		
A	red	100\$	30:90		
B	black	100\$	X:90		
Problem 2:					
Choice	Ball	Outcome	Odds		
C	red or yellow	100\$	(30 + (60 - X)) : 90		
D	black or yellow	100\$	(X + (60 - X)) : 90		

Table 4.2: Ellsberg problem

What can be seen is that the utility does not reproduce the preferences observed in the data. These preferences are: $A \prec B$ and $C \succ D$; but the rational choice yields eut(A) > eut(B) and eut(C) > eut(D). Because both problems use the same parameters (monetary values) and the person's *status quo* is unlikely to change between one problem and another; then the attribute money can only have either a positive utility or a negative utility (hypothetical assumption). This means that a person can only choose between seeing money as desirable or undesirable. If the monetary utility is desirable the ranking is the same as the one yield by the $eut(\bullet)$ function. If, on the other hand, it is undesirable⁵ it yields eut(A) < eut(B) and eut(C) < eut(D), which violates the $C \succ D$ preference. Either way this means that only two explanations are possible. The first is that the problem cannot be modelled using the utility theory since no utility function was found that maps the preferences. The second is that $eut(\bullet)$ is not an acceptable model for the individual's utility function. The latter is shown by Kahneman and Tversky [KT79] to be the explanation – see Sec. 4.5. Thus, the Allais paradox can be used to show that limiting the knowledge and computational power do not suffice to account for deviations in rationality.

4.4.2 The Ellsberg Paradox

The Ellsberg paradox [Ell61] has a similar problem setting to the Allais paradox. For the participants it is given an urn containing 90 balls and it is said that from these 90 balls 30 are red. The remaining 60 balls are an unknown mixture of black and yellow balls. The task is to guess the colour of a ball (before drawing it) and bet on this ball's colour, then the ball is drawn and its colour inspected. The possible bets are depicted in Tab. 4.2.

In the first problem the participant may choose between picking a red or a black ball. This means that if he/she thinks that the chances of picking a red ball is higher compared to picking a black ball he/she guesses red. If, on the other hand, he/she reasons that the chances of picking a black ball is indeed higher than picking a red ball then black is guessed. Then, if the participant guessed right he/she becomes a monetary outcome of 100\$, other else nothing (wrong guess). For the second problem the setting is changed a little, as shown by Tab. 4.2, but in essence the choice problem is similar. The only difference is that if a yellow ball is draw the participant gains regardless of his/her choice (in the first case he/she looses regardless). It is important to say that the same person participate in both choice problems, one after another.

Repeating: after presenting the problems' conditions to the participants, they must choose a lottery/option and draw a ball from the urn (returning it to the urn after inspecting its colour). First they must choose between A and B (*Problem 1*:) and then between C and D (*Problem 2*:). To the participants the column "Odds" is not shown.

Analysing the choice problems, they are reduced (in both cases) to estimating the amount of black balls, which in Tab. 4.2 is represented by X in "Odds", and then betting against it. Because the outcomes (in column "Outcome") are the same; the participant must only estimate the amount

⁵ For example using $x_X = -x_{Outcome}$ or $x_X = e^{-x_{Outcome}}$ as the monetary attribute utility

of black balls. In the first problem, if the participant estimates an amount of black balls higher than 30 (X > 30) he/she will choose B. For the second choice, the problem remains the same, i.e. to estimate the amount of black balls and to bet against it (as evidenced by column "Odds" in Tab. 4.2). This means that if a person chooses A for the first problem he/she should also choose C and, conversely, if she/he chooses B then also D. In the praxis the most observed combination is A with D, which is inconsistent.⁶

The conclusion for this paradox is the same as for the Allais paradox: either the utility theory does not apply for the problem or the rational model does not correctly address the option's utility.⁷ This paradox explores the "sure-thing" principle that says that a person prefers to bet against known chances (for A they are 30 in 90 and for D 60 in 90).

4.5 The Prospect Theory

The Prospect Theory [KT79] (PT) was developed to tackle the Allais paradox. The mathematical model is presented in Eq. 4.1 and it is also a utility based theory, but not rational. The main difference between PT and the Expected Utility Theory (EUT) is in how the outcomes are combined with their respectively probabilities. In PT two new functions arrive: $v(\bullet)$ (Eq. 4.2) and $\pi(\bullet)$ (Eq. 4.3). These functions are non-linear transformations for the outcomes and probabilities, respectively, and they represent the participant's perception of both values. In Eq. 3.3 the utility is calculated as $eut(\bullet) = \sum xp$ and in Eq. 4.1 it is calculated as $pt(\bullet) = \sum v(x)\pi(p)$.

$$pt(\mathbf{x}) = \sum_{\langle x, p \rangle \in \mathbf{x}} v(x)\pi(p)$$
(4.1)

$$v(x) = \begin{cases} x^{\alpha} & x \ge 0\\ -\lambda(-x)^{\beta} & x < 0 \end{cases}$$

$$(4.2)$$

$$\pi(p) = \frac{p^{\gamma}}{(p^{\gamma} + (1-p)^{\gamma})^{1/\gamma}} \begin{cases} \gamma = \gamma^+ & x \ge 0\\ \gamma = \gamma^- & x < 0 \end{cases}$$
(4.3)

To have a visual comparison between EUT and PT, regarding the manipulation of the outcomes and probabilities refer to Fig. 4.1. There the outcome transformation, corresponding to the EUT, is depicted by the curve v(x) = x in Fig. 4.1a. In Fig. 4.1a is also the $v(\bullet)$ function used by the PT: curve $v(x)^{\alpha}$ for positive outcomes and $v(x) = -\lambda(-x)^{\beta}$ for negative outcomes. For the $\pi(\bullet)$ function, the curve identified by $\gamma = 1.00$ in Fig. 4.1b represents the EUT "distortion" of a probability p and the curve $\gamma = 0.61$ represents the distortion for positive and $\gamma = 0.69$ for negative outcomes. Thus, it can be said that Eq. 4.1 introduces two elements: the status quo (reference dependence) and probability distortion. An additional comment must be made: if all parameters (α , β , λ , and γ) are reduced to 1 then the model reduces to the EUT utility function.

4.5.1 Status Quo And Reference Dependence

In PT the utility evaluation (Eq. 4.1) depends on the *status quo*, i.e. the participant's current state. Therefore in $pt(\bullet)$ (Eq. 4.1) it is also assumed that the outcomes (the x values) are already manipulated to assume the zero as the *status quo*.⁸

As discussed in Sec. 4.2.2, the *status quo* represents the current individual state and it determines how the individual perceives the outcomes. This perception can be seen in Fig. 4.1a, where the steepness in the losses (region identified by $v(x) = -\lambda(-x)^{\beta}$) is higher when compared to the gains (region identified by $v(x) = x^{\alpha}$). According to [KT79] this is necessary to express the

⁶ Unfortunately, the original article [Ell61] does not give any numerical evaluation of the problem.

 $^{^{7}}$ In [SLB08] it is called the Monotonicity axiom, which informally says that an agent always want more of something good. For the Ellsberg paradox it translates into betting in the higher guessed probability since the outcomes are the same.

⁸ If not then v(x) = v(x - sq), sq is the status quo value, and γ is adjusted according to the x - sq signal.



Figure 4.1: Prospect Theory functions

observed behaviour towards monetary lotteries: riskier in the losses than in the gains. This means that the individuals are less willing to risk when in a profit situation. They perceive the outcome differences $-\Delta x$ – less significant than they really are in gains and are more sensible to differences in losses, i.e. the differences (Δx) appear greater than they really are.

4.5.2 Probability Distortion Function

When analysing the influence of the *status quo* in the function $\pi(\bullet)$ (depicted in Fig. 4.1b) the difference is not as salient as for $v(\bullet)$ (Fig. 4.1a). The differences, on the other hand, between $\pi(\bullet)$ and just p ($\gamma = 1.00$ in Fig. 4.1b) are visible: the "bumps" of $\pi(\bullet)$.

The objective of $\pi(\bullet)$ is to model the "hope" and "sceptic" observed. When individuals evaluate outcomes with low probabilities, they believe that their chances are actually higher than the real probabilities, but for the high probabilities the effect is inverted, people are pessimistic about them. An explanation given in [KT79] is that when a person is confronted with an outcome having a high probability she/he acts as some external factor could intervene and lower her/his chances. For these reasons the function $\pi(\bullet)$ has the depicted inverted "S" shape (Fig. 4.1b), with a "upper bump" near the origin and a "depression" near the certainty. This means that the function $\pi(\bullet)$ (Eq. 4.3) over compensates low probabilities⁹ and under compensates high probabilities.¹⁰ These distortions expressed by the functions are necessary to reproduce the behaviour observed in the experiments.

The function $\pi(\bullet)$ is the main responsible for the ability of the PT to reproduce the Allais paradox.¹¹ To better understand the influence of the $\pi(\bullet)$ function, in Fig. 4.2 the lottery A from Tab. 4.3 is depicted with only the values of $\pi(\bullet)$ corresponding to each outcome. There it can be seen that the probability of the outcome 2400 is severely decreased¹² while the corresponding to the outcome 2500 is elevated.¹³

In the graphic (Fig.4.2) both features/distortions shown by the function $\pi(\bullet)$ can be observed. The probability of outcome 0 is overcompensated;¹⁴ but the effect is clearer by the outcome 2400. At this particular outcome the $\pi(\bullet)$ function has a strong influence by greatly diminishing the real probability and this is responsible for making the utility of A lower than the utility of B (the utility referred here is the one calculated by the $pt(\bullet)$ function).

The reasons for the deviations are because the function $\pi(\bullet)$ violates the independence axiom (Sec. 3.3). This axiom says that if two lotteries are equivalent, i.e. their utilities are equal, then

⁹ In the region near the origin, when comparing with the EUT evaluation, shown by the dotted line in Fig. 4.1b. ¹⁰ In Fig. 4.1b comparing the EUT (dotted line) with the $\pi(\bullet)$ for the region near the certainty.

¹¹ The outcome distortion function is also necessary but it clearly does not have great influence when compared with the $\pi(\bullet)$ function, because the outcomes are restricted to only one region of the $v(\bullet)$ (the positive).

 $^{^{12}}$ It has a high probability and therefore it is under compensated.

¹³ The probability is located at the first "bump" in Fig. 4.1b.

 $^{^{14}}$ It has no influence because de outcome value annuls it.



Comparison between eut() and pt() regarding probabilities (for lottery A)



Table 4.3: Modified Allais problem

Problem 1:					
Opt.	Outcome	Prob.	eut(ullet)	pt(ullet)	Pref. (in $\%$)
A	$\begin{array}{c} 2500\\ 2400\\ 0\end{array}$	$0.33 \\ 0.66 \\ 0.01$	2409	806.5	18
B	2400	1.0	2400	943.2	82
Problem 2:					
Opt.	Outcome	Prob.	eut(ullet)	$pt(\bullet)$	Pref. (in $\%$)
С	$\begin{array}{c} 2500 \\ 0 \end{array}$	$0.33 \\ 0.67$	825	326.7	83
D	$\begin{array}{c} 2400 \\ 0 \end{array}$	$\begin{array}{r} 0.34 \\ 0.66 \end{array}$	816	320.1	17

they may be combined using any probability and the mixture remains has the same utility. This is not the case for the PT, where the region does matter (see Fig. 4.1b)

4.5.3 Reproducing The Allais Paradox Preferences

When the parameters from Eq. 4.1 are set to $\alpha = \beta = 0.88$ and $\lambda = 2.25$ (Eq. 4.2) as well as $\gamma^+ = 0.61$ and $\gamma^- = 0.69$ (Eq. 4.3) the function $pt(\bullet)$ yields utilities in the "eut(•)" column from Tab. 4.3. Revisiting the Tab. 4.1 in Tab. 4.3, where the column " $pt(\bullet)$ " is added, it is possible to see that this utility function reproduces the preferences.¹⁵ This shows that the utility theory applies to the Allais paradox (a model exists) but not the rational model (EUT).

¹⁵ Recollecting, the preferences are: $A \prec B$ and $C \succ D$. Then looking at the values given by $pt(\bullet)$, it yields: pt(A) < pt(B) and pt(C) > pt(D).

4.6 Prospects And Editing Phase

The main reason why the Prospect Theory is so called is because of the prospect concept. A prospect is similar but not equivalent to a lottery (Sec. 3.3), which is a simple pair set. A prospect is not. A prospect is a lottery after the editing phase but its mathematical representation is equivalent: $\{\langle x_0, p_0 \rangle, \ldots, \langle x_n, p_n \rangle\}$.

The editing phase, just informally described in [KT79], corresponds to aggregate similar outcomes assigning a new corresponding aggregated probability. Some proposals were presented and among them is the work from Narens [Nar04] which establishes a clear theory for not only editing prospects but also to reproduce the Ellsberg paradox. The theory is a development from the Support Theory [TK94, RT97] and is based on an algebra defined over cognitive symbols and semantic meanings. The problem with this theory is that it requires to know how each of the used cognitive symbols (such as the elements involved in a lottery) are mapped into semantic elements (associated with a meaning) for each person. It ultimately means that propositions such as the semantic of a colour must be scrutinised for each person and it is most likely to be infeasible.

To use the PT, a method based on the approach of [DS05] is used for transforming a lottery into a prospect. The basic idea is to edit a lottery by aggregating elements from it assuming a threshold ϵ , then a clustering method is applied to it. Because this method is part of the contribution of this work and has some issues to be discussed it is presented in the next chapter.

4.6.1 Coping With The Editing Phase

In the literature two distinctive approaches are found to cope with the editing phase. The first is to avoid it completely. This is what was used by Kahneman and Tversky [KT79] when presenting the PT. They have not allowed the outcomes to be too similar, i.e. the possible outcomes were so chosen that they do not require an editing phase. This is the same approach used in [Avi06] and also in [SK04]. In the first the Cumulative Prospect Theory (presented in Sec. 4.7) is used to model route choice and it keeps the possible outcomes separated from each other, to avoid editing the lotteries. In the second, which uses PT for departure time decision modelling, the options are also kept discrete and separated.

Another approach is to use continuous prospects, which also avoids the editing altogether. This is what is done in [CS09] and in [RW08]. The continuous case of the Prospect Theory is discussed in Sec. 4.7.2.

4.6.2 Why Not Use Standard Clustering Methods

The striking question for coping with the editing phase is why not use a standard clustering method because the editing phase consists in aggregating similar outcomes and assigning an appropriate probability to the aggregated outcomes. But this implies that a structure exists that stores the outcomes received and an algorithm exists that manipulates this structure to transform the stored data into a lottery and later into a prospect. Lotteries can be easily built over the histograms of the outcomes received by the agent, where the probability is simply the frequency over the total of outcomes received.

This suggestion has however a serious drawback: it forces the agent to remember all outcomes and their frequency and this is potentially inconvenient to the space complexity since the outcome stream¹⁶ must not have a foreseeable limit. For the general purpose it is supposed that neither the amount of iterations nor the amount of possible outcomes are limited.¹⁷ To alleviate this condition it must be said that a good sample from the histogram is as good as the histogram itself.¹⁸

 $^{^{16}}$ It is called stream because the sequence of outcomes received by the agent has a time component (one outcome per step) and therefore it is called stream here.

¹⁷ This is of course a theoretical speculation since it is always limited by the amount of simulated steps. Nevertheless in some cases, as in the discussed here, some unknown limited is as good as unlimited because it can make the implementation infeasible.

 $^{^{18}}$ For the EUT this is almost obvious, since it calculates the mathematical expectancy. For the PT it is important that relevant features are not excluded from the sample – by relevant features it is meant the outcomes located at

Along with these properties it is necessary that the algorithmic complexity of the learning method suffers no significant increase. This means that the algorithm that will store the lot-tery/prospect must not demand an increasing memory space nor computational time as the agent gets experienced with the world (as it accumulates experience).

To cope with these restrictions a good clustering method solves the problem. Not any clustering method, however, because it is run over dynamic data (the outcomes) and to store the histogram is prohibitive. In data mining this is called "Clustering Evolving Data Streams".¹⁹ Among the methods is the lossy counting algorithm [MM02]. But it has two major drawbacks: the space complexity is unbounded and new data is as good as old data. Another alternative is to use the sliding window clustering [HK05] but this also presents a drawback: at each step the window must be shifted (so that the oldest data leaves the window) and it implies calculating the histograms/clusters each time. This can be time consuming when the amount of agents doing this operation is large. It also may be restrictive if the window size required as well as the amount of agents is large.

4.7 Advances In The Prospect Theory

After the original model was proposed [KT79], some improvements were developed. Among them are the Cumulative Prospect Theory [TK92, WT93] (CPT), the Prospect Theory for continuous distributions [RW08] (ContPT), and the Continuous Cumulative Prospect Theory [DS05] (ContCPT).

4.7.1 Cumulative Prospect Theory

The CPT is an improvement over the PT and includes another psychological feature. This feature is the evaluation of a prospect looking at the outcomes through the cumulative effect of the probabilities. Some evidences in [FW97] shown that the CPT is more descriptive than the PT, even though other evidences [CH94, WG96] suggest the contrary. Therefore which theory is the best is still to be investigated.

To let it be clear, the model is presented in Eq. 4.4. A prerequisite for the CPT is that the prospect must be ordered by the prospect outcomes, i.e. the lowest outcome is the first to appear and the highest the last.²⁰ Therefore the pairs have indices (indicated by a subscription) and the index *o* corresponds to the *status quo* index. The terms \mathbf{x}^+ and \mathbf{x}^- refer to the positive part of the prospect \mathbf{x} (indices higher than *o*) and to the negative part of \mathbf{x} (indices lower than *o*).

$$cpt(\mathbf{x}) = V(\mathbf{x}^+) + V(\mathbf{x}^-) \tag{4.4}$$

$$V(\mathbf{x}^{+}) = \sum_{i=0}^{n} \pi_{i}^{+} v(x_{i} - x_{o})$$
(4.5)

$$V(\mathbf{x}^{-}) = \sum_{i=-m}^{o} \pi_{i}^{-} v(x_{i} - x_{o})$$
(4.6)

$$v(x_i) = \begin{cases} x_i^{\alpha} & x_i \ge 0\\ -\lambda(-x_i)^{\beta} & x_i < 0 \end{cases}$$

$$(4.7)$$

The $cpt(\bullet)$ function (Eq. 4.4) aggregates the cumulative utility of the positive (function $V(\mathbf{x}^+)$ in Eq. 4.5) and negative (function $V(\mathbf{x}^-)$ in Eq. 4.6) parts of the prospect. These two functions are similar to the $pt(\bullet)$ function (Eq. 4.1) but not entirely; they differ in the $\pi(\bullet)$ function (the function $v(\bullet)$ in Eq. 4.7 is the same as in Eq. 4.2). In the regular $\pi(\bullet)$ the evaluation is the same as in the functions $w^+(\bullet)$ (Eq. 4.12) and $w^-(\bullet)$ (Eq. 4.13), also depending on the signal of the

the "bumps" of the $\pi(\bullet)$ function.

 $^{^{19}}$ According to [HK05] in Sec. 8.1.5.

²⁰ This means that $x_i < x_{i+1}$ for all pairs $\langle x_i, p_i \rangle$ in the prospect.

 π_i

outcome. For the CPT the probability distortion function is a little more elaborated and depends on both the current pair index and the signal of the pair's outcome being evaluated.

$$\pi_n^+ = w^+(p_n) \tag{4.8}$$

$$\pi_{--}^- = w^-(p_{--m}) \tag{4.9}$$

$$\pi_{-m}^{+} = w^{+} \left(\sum_{j=i}^{n} p_{j} \right) - w^{+} \left(\sum_{j=i+1}^{n} p_{j} \right)$$

$$(4.9)$$

$$o \le i \le n-1 \tag{4.10}$$

$$= w^{-} \left(\sum_{i=1}^{i} p_{i} \right) - w^{-} \left(\sum_{i=1}^{i-1} p_{i} \right)$$

$$\begin{pmatrix} \sum \\ j=-m \end{pmatrix} \begin{pmatrix} \sum \\ j=-m \end{pmatrix} \begin{pmatrix} \sum \\ j=-m \end{pmatrix}$$

$$1-m \le i \le o$$
(4.11)

$$w^{+}(p) = \frac{p^{\gamma^{+}}}{(p^{\gamma^{+}} + (1-p)^{\gamma^{+}})^{1/\gamma^{+}}}$$
(4.12)

$$w^{-}(p) = \frac{p^{\gamma}}{(p^{\gamma^{-}} + (1-p)^{\gamma^{-}})^{1/\gamma^{-}}}$$
(4.13)

The original $\pi(\bullet)$ (Eq. 4.3) function is in $w^+(\bullet)$ and $w^-(\bullet)$. The essential difference is in taking the difference between the distortion of the cumulative probabilities for the *i*th pair instead of only manipulating the probability of the pair itself. This is calculated by π_i^- (Eq. 4.11) and π_i^+ (Eq. 4.10), the negative and positive outcomes respectively. The exceptions are the in the extremes, i.e. the first²¹ and last pairs which are calculated by π_m^- (Eq. 4.9) and π_n^+ (Eq. 4.8) and have no cumulative aspect. The accumulation can be described as:

- 1. Consider the probabilities as discrete points in a probability distribution function²²;
- 2. For positive outcomes set the accumulation limit to the last pair and for negative to the first;
- 3. Set the first probability region to be from the *i*th pair to the limit and calculated the probability distortion;
- 4. Set the second region to be from the next pair closer to the limit and also calculate the distortion;²³
- 5. The difference between the distortion of the first and second regions yields the outcome probability perception.

Taking the example of prospect A in Tab. 4.3, when analysing the perception of the probability for the outcome 2400, the informal interpretation is as follows. At 2400 the remaining chunk is 0.99 (0.66 + 0.33) and the next remaining chunk is 0.33, so the perceived probability is the perception of 0.99 discounted the perception of 0.33. Nevertheless the advantages of the CPT over the PT are not clear, as mentioned before.

4.7.2 The Prospect Theory For Continuous Prospects

The continuous version of the PT [RW08] (ContPT) works by assuming that an outcome threshold ϵ exists to transform the continuous prospect into a discrete prospect. The strategy used was to

²¹ Recall that for the CPT the prospect must be ordered in increasing manner according to the outcome value, i.e. the first pair in the prospect has the lowest outcome value and the last the highest.

 $^{^{22}}$ This works because the pairs are ordered by the outcome in the lottery.

²³ For positive outcomes it means the next pair (i + 1) and for negative outcomes the previous pair (i - 1).

derive the resulting PT model when ϵ tends to 0, in this case transforming the PT model into the ContPT. For the continuous CPT [DS05] (ContCPT) the problem is more complex. Because the cumulative aspect, which is based on the concept of capacity [Cho54], the accumulations (from the $\pi(\bullet)$ function) must be taken into account when extending the CPT for the continuous case. This process strongly depends on the Probability Distribution Function (PDF) adopted and therefore a model must be derived for each PDF.

The continuous models are not shown because they are not investigated here, which uses only discrete prospects and lotteries. However if it were the case²⁴ then it would mean that a PDF must be chosen and the appropriated model derived.

4.8 Summary

This chapter started with the criticisms to the rational behaviour for modelling human decisionmaking (Sec. 4.2). It also points out two specific issues: the bi-parted decision making process and the *status quo* dependency. The first refers to the separation between intuition and reasoning when someone makes a decision. By intuition it is understood the "quick-and-dirty" response to a decision problem, i.e. a decision that is immediately available. This is here called the System 1 decision-making process. The reasoning system, called System 2, tackles a more cumbersome decision process that does a more profound analysis of the alternatives and tries to make the best decision.

The latter issue, the reference or *status quo*, refers to the point where gains are separated from losses, i.e. which outcomes are classified as gains and which as losses. This classification influences in how a person behaves. In the losses persons are said to behave more risky and with more precaution in the gains (Sec. 4.2.2).

The next relevant issue in this chapter is the systematic violations observed in the rational behaviour, which are explicitly shown in the Allais and Ellsberg paradoxes (Sec. 4.4). Then the Prospect Theory (PT) is presented (Sec. 4.5) that copes with the Allais paradox (Sec. 4.5.3). The PT tackles with it by distorting the outcomes and probabilities in a prospect.²⁵ These distortions are necessary to reflect the human perception of outcomes and probabilities. The first distortion function $v(\bullet)$ (Eq. 4.2) makes the explicit separation between gains and losses (Sec. 4.5.1) and the function $\pi(\bullet)$ (Eq. 4.3) distorts the probabilities. The probability perception function $\pi(\bullet)$ reflects the optimism regarding low probabilities and the pessimism for the high probabilities (Sec. 4.5.2). The optimism refers to the behaviour observed in people when confronted with low probabilities outcomes. They tend to perceive the probability higher than it really is and the inverse is observed for the high probabilities, which are perceived lower than they really are.

The editing phase that transforms a lottery into a prospect is discussed in Sec. 4.6. This procedure corresponds to the transformation of a lottery, by aggregating similar outcomes, into a prospect. It is also discussed how this is tackled in the literature (Sec. 4.6.1) and why standard clustering methods are not adequate for this task (Sec. 4.6.2).

The last part of this chapter presents the advances in the PT (Sec. 4.7). The Cumulative Prospect Theory (CPT) is presented as well as the continuous versions of the PT and CPT.

 $^{^{24}}$ An example would be to assume that the travel-time is distributed according to a particular PDF.

 $^{^{25}}$ Prospects are the subject of Sec. 4.6.

Part II Contribution

Chapter 5

Q-Learning Based On The Prospect Theory

This chapter presents the Q-Learning [Wat89, WD92] algorithm to tackle with decision under uncertainty.¹ Its main point is to present the standard Q-Learning and the modifications to transform it from a rational decision-making algorithm to a non-rational one based on the Prospect Theory (PT). It also presents presents a proposal of algorithm for tackling with the editing phase, which as explained before lacks a formalism (Sec. 4.6).

5.1 Concepts Review

The necessary concepts from the previous sections are first that a utility is an aggregation over outcomes and their probabilities.² Then the important problem of transforming a lottery into a prospect, the editing phase (Sec. 4.6). This process has not a defined formalisation (Sec. 4.6) and is usually avoided by forcing the options or outcomes to be separated enough to make the editing phase unnecessary (Sec. 4.6.1). It is also important to know that the standard clustering methods are not the best choice for coping with it (Sec. 4.6.2) and therefore a new algorithm is proposed (Sec. 5.5).

5.2 Why Learning

In the previous chapters the two main utility theories for decision-making modelling were presented and discussed – the Expected Utility Theory (EUT) in chapter 3 and the Prospect Theory (PT) in chapter 4. These theories are only capable of solving part of the problem: they provide methods for option ranking based on the expected/estimated outcomes, but they do assume that this information is given beforehand. This is not always the case, depending on the problem formulation. Therefore a learning algorithm is necessary. In any case, knowing or not the internals of the problem, the objective is always to reach the optimal policy. A policy is a function, or look up table, that associates a state with an action. An optimal policy is a policy whose state associated action is, on the long run, the one that maximises the accumulated utility.

To restrict the problems tackled by this work, it is assumed that the problems to be solved have the Markovian property [Mar71], i.e. the transitions among the different world states do not depend on the state's history. This means that the only elements that influences the world's state transition are the current state and agent's actions. Another approach is to suppose a world with

¹ Knight [Kni21] defined that decisions under risk are those made with known stochastic outcomes, as in the Allais paradox from Sec. 4.4.1, and under uncertainty are those where the outcomes are not informed, i.e. the individual must infer/estimate them.

 $^{^{2}}$ This is true for the Expected Utility Theory (Sec. 3.3) and the Prospect Theory (Sec. 4.5).

Bayesian properties [Bay63], where the previous state of affairs do matter for the transitions. The latter is not discussed in this thesis.

Another point is that learning is not necessary if the state transitions are known, i.e. the decision-makers are aware of the state transition and reward functions for the given problem. In this case the use of most of the learning techniques are counter-productive because they are usually less efficient than analytical solvers.³ However, if the transition and reward functions are not known then learning algorithms seem to be the best strategy [SLB08] for repetitious decisions with a temporal⁴ component. The basic scenario under which the agents make their decisions is both iterative and under uncertainty, because the agents are not aware of the outcomes of their actions before executing them. In this scenario the task, or tasks, is repeated indefinitely, which characterise the conditions for an implicit decision problem.⁵ Another issue is the uncertainty about the outcomes; they are not beforehand known to the agent, i.e. the agent is not aware of the options' costs before making a decision.

As said before, it is assumed that the world/problem has the Markovian property [Mar71] – the historical state track does not interfere with future state transitions. Formally it is assumed that the world is a Markov Decision Process [Mar71, Bel57b, Bel57a] (MDP), but other formalisms have been also considered and this discussion is addressed in Sec. 5.3.

Agent learning is needed because in the used MDP the transitions and outcomes are not known to the agents. By this it is meant that to "solve" the choice problem imposed by the MDP the agent must adopt an optimal policy, i.e. to have an action associated with each possible state, where the expected accumulated outcome is optimal. To achieve this, without knowing the transition and reward functions, the agent must first learn how states are reached and rewards are assigned, meaning to learn the Probability Distribution Functions (PDFs) associated with each function. The idea of informing the state transition and reward functions may seem appealing but its applicability rapidly diminishes as the amount of actions, states, and rewards increases. The stochastic behaviour of the state transition and reward functions may turn the enumeration, prerequisite for an analytical solution, into an infeasible criterion for real world scenarios. No to mention the problem when several agents are interacting with each other at each iteration, as exposed by [SLB08]. Because of this, the enumeration of all possibilities may exceed the computational capabilities available, making this strategy less desirable. Learning, on the other hand, can be very effective in regard to this space problem, since a good learning algorithm learns tendencies (where the mathematical expectancy of each function is located) and not subtleties.

Therefore the learning strategy is adopted and the specific Q-Learning algorithm [Wat89, WD92] was chosen. It is chosen because it is a reinforcement learning algorithm and second because it can be easily adapted for the use of EUT and PT – the two utility based theories being compared. This is not the case of $TD(\lambda)$ [Sut84, Sut88] because it takes advantage of the rational utility for mixed lotteries – the utility of mixed utilities in EUT is the mathematical expectancy (Sec. 3.3). For the PT, the outcome PDF must be known/inferred, but in $TD(\lambda)$ the strategy is to only accumulate the outcomes and correct the estimated expectancy value, it does not keep the PDF behind the outcome function.

A last explanation worth noting is that here only the learning aspects of the decision-making process are approached. This means that the algorithmic part that switches between exploration and exploitation is left out of its specifications. It is so because this part is approached in the next chapter when presenting the agent architecture. But for the sake of the argument in Eq. 5.1 a simple version of it is given. There: $D_n(s)$ is the effective decision made at step n for state s; Rand(0,1) is a function that draws an number from an uniform distribution in the interval [0,1]; $r_{explore} \in (0,1)$ is the exploration rate;⁶ \mathcal{A}_s is the set of all possible actions for the state s;

³ In theory, most of the learning algorithms require unlimited time to guarantee the convergence to the optimal policy, even though in practical terms this is not the case.

 $^{^4}$ By temporal it is meant the fact that the notion of order in the sequence of decisions exists.

 $^{^{5}}$ It refers to the bi-parted decision-making system where familiar decisions tend to be (implicit) made using intuition, the System 1 (Sec. 4.2.1).

⁶ The extremes are excluded because 0 means that the agent never explores and 1 that it never exploits. Usually a conservative low rate is chosen, such as $r_{explore} = 0.1$, but with one exception: $r_{explore} = 1$ iff n = 0, i.e. the

 $\operatorname{Rand}(\mathcal{A}_s)$ randomly chooses one of these actions; and $V_n(s)$ is the decision made based on what has been learnt about the decision process, i.e. the exploitation part.

$$D_n(s) = \begin{cases} \text{iff } \text{Rand}(0,1) \le r_{explore} & \text{Rand}(\mathcal{A}_s) \\ & \text{else} & V_n(s) \end{cases}$$
(5.1)

5.3 Why MDP

Among the several formalisms characterised by the Markovian property are: Markov Decision Process [Mar71, Bel57b, Bel57a] (MDP), Stochastic Games [Sha53] (SG), Partially Observable Markov Decision Process [Lov91] (POMDP), and Partially Observable Stochastic Games [HBZ04] (POSG). These four models are classified in Tab. 5.1, where "States" means that the agent(s) is/are aware of the current state and "Observations" means that the agent(s) has/have an indirect⁷ perception of the state. Moreover, "Single Agent" means that the agent(s) is/are not aware of any other possible agent in the world; and "Multi-Agent" means that the agent(s) is/are aware of the existence of other agents.

Table 5.1: Theory classification

	Single Agent	Multi-Agent
States	MDP	SG
Observations	POMDP	POSG

Formally an MDP world is defined by the tuple $\langle S, A, T, R \rangle$ where S is the state set, A the action set, $T : S \times A \mapsto \Pi(S)$ the transition function where $\Pi(S)$ is the set of all PDFs over the state set, and $R : A \times S \mapsto \mathbb{R}$ the reward function. The transition function $T(\bullet)$ is a function that returns a PDF over the state set S for each state-action pair. The function $T(\bullet)$ can be changed, for convenience, to T(s, a, s') = p where: s is the current state; a the action chosen at state s; s' a possible next state; and $p \in [0, 1]$ the probability for reaching the state s' from state s using the action a. The reward function $R(a, s) \in \mathbb{R}$ returns a numerical value (the outcome) for reaching state s through action a, regardless of the previous states.

The POMDP world has additional structures and its tuple is $\langle S, A, T, R, O, \Omega \rangle$ where Ω is the observation set and $O : S \times A \mapsto \Pi(\Omega)$ the sensor function, i.e. what the agent observe from the world given the current state and the action chosen. The sensor function can be transformed into O(s, a, o) = p where: s is the current state s; a the action chosen; o a possible observation; and $p \in [0, 1]$ the probability of receiving the observation o given the state-action pair (s, a).

The use of MDP or POMDP depends on the problem being modelled. In traffic both formalisms are possible. If the state is the network status (such as link load for each link) and the observations represent the experienced link loads of the chosen route, then POMDP is the correct formalism. Consequently, the agent does not know the current status (the load of each single link) but only its own experience. But if the network state is not relevant for the agent⁸ and the route experience is sufficient then the route experience is the state and the MDP is the correct formalism. This is the approach used, because it is supposed to be seldom the case where a driver is interested in knowing the complete network state and even rarer the case in which the driver keeps track of such states. Another approach to the problem is to consider the observation, or the feedback about the world state, as the reward. In this case the reward function is defined as $R : \mathcal{A} \times \mathcal{S} \mapsto \Pi(\mathbb{R})$ where $\Pi(\mathbb{R})$ is the set of all possible PDFs over the real numbers. This means that the function $R(a, s) = \pi(\mathbb{R})$ returns a particular PDF over \mathbb{R} for each (a, s) pair. Again in the route choice example: a state will be characterised by a transportation need⁹ (a route) for a given situation, the route set will

first decision is always an exploration.

 $^{^{7}}$ Just some features of the state, the observations, are given to the agent(s) as feedback about its/their current state.

 $^{^{8}}$ The case where the agent is not concerned about the state of the each link in the whole network.

 $^{^9}$ A transportation need is task where an individual must go from A to B using a transportation mean.

be the action set, and the reward function builds a cost PDF for each route. To give a concrete example in commuting from home to work and vice-versa (the commuting scenario): two MPD states are possible, where the first state is the route from home to work and the second from work to home. If, in the meanwhile, another transportation need approaches, it will be treated as another state, and so on. Therefore the partial observability is not an issue¹⁰ and avoiding it is also desired for complexity reasons. Because a POMDP needs to guess the current state according to the observation, to solve a POMDP requires also being able to correctly guess the current state and therefore able to make a good action choice. Therefore partial observability is used not in this work.

For the multi-agent property the possibilities are SG and POSG. In both cases the "awareness" of other agents is added to the tuples. For the SG the tuple is $(\mathcal{S}, \mathcal{A}, T, R, \mathcal{I})$ where \mathcal{I} is the agent set. Moreover, the functions $T(\bullet)$ and $R(\bullet)$ must be changed to take all agents into account. Therefore the transition function is defined as $T: \mathcal{S} \times \mathcal{A}^n \mapsto \Pi(\mathcal{S})$, or seen as $T(s, \mathbf{a}, s') = p$ where $\mathbf{a} = \langle a_0, a_1, \dots, a_n \rangle$ is the joint action vector (each agent $i \in \mathcal{I}$ makes a decision about the action a_i to use). The reward function also needs the joint action vector and then it is defined as $R: \mathcal{S} \times \mathcal{A}^n \mapsto \mathbb{R}^n$ or $R(s, \mathbf{a}) = \mathbf{r}$, where $\mathbf{r} = \langle r_0, r_1, \ldots, r_n \rangle$ gives the reward to each action for the corresponding agents. An analogous process also modifies the POMDP into POSG; but as the POMDP is not considered, the formalism for POSG is not presented.

The relevance of SG is that it takes into account all agents, which is the case for multi-agents. However, here the perspective is not the one from the world but from the agent. This means that, the agent perception of what the world is, is important and not what the world actually is. In [ER07] it is shown that the contribution of each agent declines as the amount of agents increases and strategies that ignore the other agents are also effective (in environments where the amount of agents is high). Another point is that humans are unlikely to keep track of each other competitor for environments where the social interaction is minimal – as observed in [AOR05] and in [KR00]. Therefore the explicit representation of each other agent by the agent itself is not made, i.e. the agent does not represent the other agents in its internal structures – thus the MDP representation is chosen. This was also the opinion of von Neumann and Morgenstern [vNM07], explicitly made in section 2.4.1 and again in 2.5:

[From 2.4.1]: "... When the number of participants becomes really great, some hope emerges that the influence of every particular participant will become negligible, and that the above difficulties¹¹ may recede and a more conventional theory become possible. These are, of course, the classical conditions of "free competition." Indeed, this was the starting point of much of what is best in economic theory."

[From 2.5]: "... Sometimes free competition is assumed, after the introduction of which the participants face fixed conditions and act like a number of Robinson Cru $soes^{12}$ – solely bent on maximizing their individual satisfactions, which under these conditions are again independent."

5.4Standard Q-Learning

The original Q-Learning [Wat89, WD92] is a reinforcement learning algorithm for solving MDP problems¹³ without knowing the environment beforehand. This means that the learning procedure needs to know neither the outcomes of its actions nor the transition function, prior to its decision, to achieve the "maximum accumulated discounted wealth/reward". Because the algorithm does not need the options' outcomes it is suitable for learning in environments with decisions under uncertainty, i.e. where the outcomes are stochastic and the reward's PDF is unknown to the

¹⁰ In this case instead of using observations the agent assumes the feedback as the state itself.

¹¹ It refers to the problem of multiple players, where a 2 players game is different from a 3 players and this is different from a 4 players game and so on.

¹² In [vNM07] Robinson Crusoe is how a single player is called, i.e. a player that acts as if he/she is alone in the game. ¹³ As said before, MDP problems assume a single agent trying to find the best policy for the given problem.

agent. The general idea of the algorithm is to explore the environment and to learn how actions influence the accumulated reward and to increasingly make better decisions (as the experience accumulates) that maximise its discounted accumulated wealth. It is shown in [Wat89, WD92] that given enough iterations the algorithm converges to the rational behaviour. To achieve this it starts by heavily exploring the environment (trying all the possible actions) and as the experience is accumulated it decreases the environment exploration, thus increasing the exploitation of own knowledge. In layman's terms it is: the agent starts trying everything to understand the nature of the environment and then, as the experience is accumulated, relies more and more on own estimates, ending up using only the accumulated knowledge.

The algorithm does that by keeping an estimate of each action's reward and updating it as the action is reused. This process slowly converges to the mathematical expectancy of the actions' reward (the EUT utility in Eq. 3.3), without requiring x_i or p_i to be known (the lottery pairs describing the outcomes' PDF). Eq. 5.3 shows the original specification of Q-Learning, where $V_n(s)$ returns an action a for iteration n at state s. As n grows $V_n(s)$ tends to be the best action, based on the accumulated knowledge about the received rewards that are stored in the table/function $Q(\bullet)$ (Eq. 5.2). In Eq. 5.2 r_n is the immediate reward (received at iteration n), γ a fixed discount factor (to recursively decrease the influence of the possible subsequent actions), and α_n the decreasing $(\lim_{n\to\infty} \alpha_n = 0)$ learning factor to help $Q(\bullet)$ to converge to the $eut(\bullet)$ value.

$$Q_{n}(s,a) = (1 - \alpha_{n})Q_{n-1}(s,a)$$

$$+ \alpha_{n} \left[r_{n} + \gamma \sum_{s' \in S} V_{n-1}(s') \right]$$

$$V_{n}(s) \equiv \operatorname{argmax}_{a \in \mathcal{A}} [Q_{n}(s,a)]$$
(5.3)

5.4.1 Learning Factor α And Exploration Rate

The learning controlling factor α (Eq. 8.4) is defined in [Lit94] and explicitly given in Eq. 5.4.¹⁴ The problem with such definition is that it requires a considerable large horizon to reach learning "accommodation" – to reach $\alpha_n \leq 0.1$, it requires n > 500560. In other words, the original definition requires more than 500560 steps to reach a learning rate of 10%. This is way too high a value for practical experiments, i.e. a too long simulation time. The other simple solution is to decrease α_0 to fit a desirable accommodation speed, i.e. to establish the desired learning rate by a wanted horizon n and then to calculate the α_0 . For instance, to achieve the $\alpha_{100} \leq 0.1$ it would require that $\alpha_0 \leq 0.97723$, but it also makes $\alpha_{50} \leq 0.31623$ and $\alpha_{25} \leq 0.562341$. Which is a too quick accommodation compared with the original: $\alpha_{50} \leq 0.999770026$. This could compromise the learning performance. To give an idea of the α_n with different α_0 see Fig. 5.1,¹⁵ where the curve with $\alpha_0 = 0.9999954$ (the original value) has no noticeable decrease (it is almost a flat line at the value 1.0).

$$\alpha_0 = 0.9999954$$
$$\alpha_n = \alpha_0^n \tag{5.4}$$

A second approach is to change the function to a more suitable for short horizons. To tackle this, the Richards' function [Ric59] was chosen. This function is an adjustable sigmoid function (Eq. 5.5) where α_{init} is the initial learning factor (usually $\alpha_{init} = 1$), α_{final} the desirable value for

¹⁴ The values were taken from [Lit94] because in the original article [WD92] no equation for α is given, only the guidelines that $(\lim_{n\to\infty} \alpha_n = 0)$. In the thesis from Watkins [Wat89] there is also little clarification. In some cases α is kept constant and in other it is dependent from state variables. Therefore, the equations given in [Lit94] are used.

 $^{^{15}}$ For a higher horizon, please see Fig. B.2.



Figure 5.1: The regular α_n for different α_0 and horizon of 100



Figure 5.2: Comparison between the two possibilities for the α_n function

 α when reaching the wanted horizon n. The other factors $(Q, B, M, \text{and } \nu)$ are flexible, but the chosen values give a fairly symmetric curve along the simulation horizon. A comparison between this function and the original exponential curve, with $\alpha_0 = 0.97723$, is depicted in Fig. 5.2.¹⁶ There the parameters are: $\alpha_{init} = 1$, $\alpha_{final} = 0.1$, and n = 100.

$$A = \alpha_{init}$$

$$K = \alpha_{final} - \alpha_{init}$$

$$Q = 0.5$$

$$B = 10/n$$

$$M = n/2$$

$$\nu = 0.5$$

$$\alpha_n = A + \frac{K}{\left(1 + Qe^{-B(i-M)}\right)^{1/\nu}}$$
(5.5)

The exploration rate, i.e. how much does the agent experiment with the environment is also supposedly decreasing with n. To cope with this the same strategy as for α is used and the Eq. 5.5 is used, except that for exploration the $\alpha_{init} = 0.2$ instead of 1.

5.4.2 How To Transform The Standard Q-Learning To Be PT Based

As the original Q-Learning algorithm converges to the rational behaviour, changing it to become non-rational (based on the PT) means that $Q(\bullet)$ must return a $pt(\bullet)$ "converging" value instead of

 $^{^{16}}$ For a visualisation of the impact of the different parameters please address to Fig. B.3.



Figure 5.3: Single-state Q-Learning $Q_n(a)$ evolution

an $eut(\bullet)$ "converging" value. To make clear how the standard Q-Learning eventually converges¹⁷ to the $eut(\bullet)$ value, the algorithm is simplified to the "single-state" case shown in Eq. 5.7. The difference between Eq. 5.3 and 5.7 is in the $Q(\bullet)$ function working (Eq. 5.2 and 5.6). From the original function (Eq. 5.2) the MDP function $T(\bullet)$ influence is removed (the $\gamma \sum_{s' \in S} V_{n-1}(s')$ term), i.e. it does not need to take into account the influence of the action a transitioning to another state as well as to calculate this next state worthiness. In a simplified explanation, the reduced Q-Learning in Eq. 5.7 is "stateless", i.e. it is assumed that the agent remains in the same state.

$$Q_n(a) = (1 - \alpha_n)Q_{n-1}(a) + \alpha_n r_n$$
(5.6)

$$V_n \equiv \operatorname*{argmax}_{a \in A} \left[Q_n(a) \right] \tag{5.7}$$

To conduct some experiments, a hypothetical MDP (dissociated from any particular scenario) is formulated in a way that the agent acts over a single state with a single action and an hypothetical reward function $R(\bullet)$. With this oversimplified MDP it is possible to better observe how the $Q(\bullet)$ function/table evolves as iterations are accumulated.¹⁸ This experiment is designed to only evaluate the convergence of values in $Q(\bullet)$, in particular to which direction.

For the deterministic case¹⁹ the "accommodation" of the $Q(\bullet)$ table/function is rather quick, as shown in the Fig. 5.3a where the iteration step (or the $Q(\bullet)$ index) is in the x axis. For the stochastic case, in Fig. 5.3b, the "accommodation" is slower but also reaches the $eut(\bullet)$ value (Eq. 3.3). In the experiments shown in Fig. 5.3a and 5.3b, the simplified algorithm from Eq. 5.7 was used and the value of $Q(\bullet)$ (Eq. 5.6) recorded as the experiment evolved. In the stochastic version the lottery A from the Allais *Problem 1*: (Tab. 4.3) was used.²⁰ The proof that the Q-Learning always converge to the mathematical expectancy is given in [Wat89, WD92].

This means that to modify the Q-Learning algorithm for returning the PT utility (Eq. 4.1) it is necessary to modify the $Q(\bullet)$ function to converge to a $pt(\bullet)$ equivalent value, instead of the $eut(\bullet)$ from the original formulation. The simple examples from Fig. 5.3a and 5.3b are easy to transform for using the $pt(\bullet)$ function. The results are presented in Fig. 5.4a and 5.4b. This modification required an extra table to record all different rewards, but when a value is repeated

 $^{^{17}}$ The algorithm is proven to converge only at infinity, which is for all practical applications an unrealistic criterion. Therefore it is expected that the algorithm reaches the optimal policy in a discrete amount of steps, which also means that it only "eventually" reaches an optimum point in a limited amount of steps.

¹⁸ The inclusion of more actions and states makes the analysis a little more complicated because not only the specific accommodation for the current state and action is learnt but their effect on the possible forthcoming states and other possible actions.

¹⁹ Deterministic means that $R(\bullet)$ has a PDF with only one possible outcome, which has probability 1. The stochastic counterpart has PDF with several possible outcomes, each with its own probability.

 $^{^{20}}$ For each iteration step a weighted roulette wheel was used to select the outcome in the function $R(\bullet)$.

the corresponding counter was incremented.²¹ This way the extra table can be used to construct the prospect. Then instead of using $\alpha_n r_n$, $\alpha_n pt(P_n)^{22}$ was used, where P_n corresponds to the prospect built based on the current reward. This is shown in eq. 5.13, where P_n is the prospect generate from an auxiliary "table" \mathbf{C} . The table \mathbf{C} simply stores each received reward in pairs $\langle c, f \rangle$ where c is the reward value and f is the frequency with which the reward has appeared. This means that each time a reward is received the corresponding counter f is incremented or a new pair $\langle c, f \rangle$ is included with f = 1 – assuming that no other pair has a $c = r_n$, being r_n the reward received at step n.

$$\mathbf{C} = \{ \mathbf{c} \in \mathbf{C} \mid \mathbf{c} = \langle c, f \rangle \}$$
(5.8)

$$\mathbf{C} = \{ \forall \mathbf{c}, \mathbf{d} \in \mathbf{C} \mid c_{\mathbf{c}} \equiv c_{\mathbf{d}} \Rightarrow \mathbf{c} = \langle c_{\mathbf{c}}, f_{\mathbf{c}} + f_{\mathbf{d}} \rangle \land \mathbf{C} = \mathbf{C} - \{\mathbf{d}\} \}$$
(5.9)

$$\mathbf{C}_n = \mathbf{C}_{n-1} + \{\langle r_n, 1 \rangle\} \tag{5.10}$$

$$P_n = \left\{ \forall \langle c, f \rangle \in \mathbf{C} \exists \langle x, p \rangle \mid x = c \land p = \frac{f}{\sum_{\langle c, f \rangle \in \mathbf{C}_n} f} \right\}$$
(5.11)

$$Q_n(a) = (1 - \alpha_n)Q_{n-1}(a) + \alpha_n pt(P_n)$$
(5.12)

$$V_n \equiv \operatorname{argmax}_{a \in \mathcal{A}} [Q_n(a)] \tag{5.13}$$

It is possible to observe that the $Q(\bullet)$ function converges to the $pt(\bullet)$ values, depicted in Fig. 5.4a and 5.4b. The reason why the PT based Q-Learning is faster (in converging) than the regular version is because of the score/reward table. This table gives more information than just an aggregation over the past rewards, which is the case of the value iterated strategy of the original Q-Learning. This means that as soon as the table has a significant sample²³ over the reward function the $Q(\bullet)$ function returns a value close enough to the expected $pt(\bullet)$. The objective of the figures was not to compare the performance of both version but to show that both converge to the theoretical expected value. The formal definition of convergence is in Eq. 5.14 (for the EUT/normal version) and in Eq. 5.15 (for the PT version).

$$\lim Q_n(a) = eut(r(a)) \tag{5.14}$$

$$\lim_{n \to \infty} Q_n(a) = pt(r(a)) \tag{5.15}$$

The proof of Eq. 5.15 is rather simple: $\lim P_n = R(\bullet)$, i.e. the experienced the agent gets to the MDP (the higher n becomes) the more its P_n evolves closer to the PDF in $R(\bullet)$ and since this is the only prerequisite to return a correct $pt(\bullet)$ equivalent value then it converges to it in $n \to \infty$.²⁴

The information issue, i.e. the extra storage needed by the PT version is approached in the next section, where the issue not covered by the previous chapter is recapitulated: the editing phase.

 $^{^{21}}$ The method used should not be confused with the adopted for the agent Q-Learning because it has several drawbacks, explained in the very next section. It is suitable, nevertheless, for this particular experiment. The reason for using a different algorithm is again to give a clear view of a single aspect (in this case the convergence of $Q(\bullet)$ instead of trying to explain every aspect together. ²² The $pt(\bullet)$ function is the same as in Eq. 4.1.

²³ The frequencies stored in the extra table are close enough to the real PDF in $R(\bullet)$.

 $^{^{24}}$ The numerical and practical issues are ignored here, i.e. when n is actually ∞ all values in P_n are also numerically ∞ and therefore the computation of the probabilities is not possible. But as discussed before, $n = \infty$ is not an acceptable condition for any practical algorithm and therefore it can be said that when P_n reaches a close enough sample of $R(\bullet)$ its values can correctly construct $R(\bullet)$ and therefore the correct $pt(\bullet)$ value.



Figure 5.4: Simple Q-Learning $Q_n(a)$ evolution, using PT

5.5 Editing Phase

In Sec. 4.6.2 it is discussed why the standard clustering methods are not appropriate for coping with the editing phase. Therefore here a simple algorithm is proposed that combines part of the simplicity of the lossy counting and part of the strategy in [ALSS95]. This editing method is also inspired by [DS05], despite of the drawbacks discussed in Sec.4.6.

The algorithm is a stream clustering algorithm based on centroids (geometrical centres), i.e. the outcomes are seen as clusters and the derived prospect corresponds to the centroids of those clusters, having the sum of the outcomes' probabilities as the cluster probability and the mean value (of the outcomes) as the representative outcome. Because $R(\bullet)$ is unidimensional the clustering is made over a single dimension. This method also requires a separation threshold, called ϵ .

As a consequence of the clustering method, the reward function codomain is limited and partially known, because ϵ must be given. It means that the derived Q-Learning algorithm needs to accumulate more knowledge about the world than the original version, because it must have an idea of the reward function codomain. This means that ϵ should not be calibrated but informed (by the modeller) to help the algorithm. This implies too, that the agent is not completely "blind" when looking at the reward function $R(\bullet)$ (it knows how to aggregate the rewards through ϵ).²⁵ Another consequence is that the $R(\bullet)$ function must map a closed interval in \mathbb{R} .²⁶

This cluster structure is defined by Eq. 5.16, where **C** is the cluster structure, which is a set of pairs $\langle c, f \rangle$. These pairs represent the centroid value in c and the amount of points in that centroid in f. In this structure the centroids are unique, i.e. no two pairs share the same centroid value c.

$$\mathbf{C} = \{ \mathbf{c} \in \mathbf{C} \mid \mathbf{c} = \langle c, f \rangle \land c \in \mathbb{R} \land a \in \mathbb{N}^* \}$$

$$\mathbf{C} = \{ \forall \ \mathbf{c} \in \mathbf{C}, \not \exists \ \mathbf{d} \in \mathbf{C} \mid c_{\mathbf{c}} = c_{\mathbf{d}} \}$$
(5.16)

In Eq. 5.16 only the structure is formalised but not how to transform it into a prospect \mathbf{x} . This process is rather simple and shown in Eq. 5.17. There $\mathbf{x}_{\mathbf{C}}$ is the prospect generated from the cluster structure \mathbf{C} ; $\langle c, f \rangle$ is a centroid/accumulator pair from \mathbf{C} ; $\langle x, p \rangle$ is the resulting outcome/probability pair (for the corresponding $\langle c, f \rangle$); and x receives the c value and p receives the frequency f converted to the corresponding probability.

 $^{^{25}}$ The agent builds the prospects using the editing algorithms but the value ϵ is given by the modeller and not "discovered" by the agent.

²⁶ If the interval is not closed the amount of items in the extra table can grow unlimited. A simple example is to try to segment in a limited number of intervals (higher than 1) the entire \mathbb{R} : the number of intervals possible are 1 and ∞ .

$$prospect(\mathbf{C}) = \mathbf{x}_{\mathbf{C}} \tag{5.17}$$

$$\mathbf{x}_{\mathbf{C}} = \left\{ \forall \langle c, f \rangle \in \mathbf{C} \ \exists \langle x, p \rangle \in \mathbf{x}_{\mathbf{C}} \ \bigg| \ x = c \land p = \frac{f}{\sum_{\langle c, f \rangle \in \mathbf{C}} f} \right\}$$

The last part of the algorithm builds the structure C from the rewards received by the agent, as it experiences it. A simple version of the algorithm is shown in Algo. 5.1. The function $Cluster(\bullet)$ receives as arguments: the current cluster structure, the clustering threshold, and the new reward to be included in the cluster. This algorithm avoids the linear growth in both memory and computational complexity. This means that once the reward function codomain is partitioned the amount of pairs in C stays constant. Another advantage of this algorithm is that it builds the structure as the rewards are received, which relaxes the codomain restrictions (the codomain must be limited but the limits must not be known beforehand).

Algorithm 5.1: Cluster(•)
Data : $\langle \text{centroid}, \text{accumulator} \rangle \text{ set } \mathbf{C}$
Data: the threshold ϵ
Data: the reward r to be included
Result: Updated set C
1 $\mathbf{c}_{min} \leftarrow NIL$; /* \mathbf{c}_{min} is the closest centroid to r */
2 $\Delta d \leftarrow \infty$; /* Δd is the distance between r and c $_{min}$ */
/* finds the closest centroid to the reward r */
3 forall pair $c \in C$ do
4 if Distance(c, r) < Δd then
5 $\Delta d \leftarrow \text{Distance}(c, r);$
$6 \mathbf{c}_{min} \leftarrow \mathbf{c};$
7 end
8 end
/* checks if the centroid is suitable for aggregation */
9 if $\Delta d \leq \epsilon$ then
$10 c_{\mathbf{c}_{min}} \leftarrow \frac{r + a_{\mathbf{c}_{min}} c_{\mathbf{c}_{min}}}{1 + a_{\mathbf{c}_{min}}} ; \text{/* new centroid value */}$
11 $a_{\mathbf{c}_{min}} \leftarrow a_{\mathbf{c}_{min}} + 1$; /* increments counter/accumulator */
12 else/* it is a new centroid */
13 $\mid \mathbf{C} \leftarrow \langle r, 1 angle$; /* add the new centroid */
14 end

The editing method presented has also drawbacks: it has a fixed threshold that may not be the most appropriate clustering method.²⁷ Another issue is its strong dependency on the initial values. If the algorithm is used as in Algo. 5.1 then the first rewards must be a characteristic draw from $R(\bullet)$, i.e. a well distributed sequential draw from $R(\bullet)$. For example, let $\epsilon = 5$ and a reward sequence be $\langle 5, 15, 7, 13, 9, 11, 9, 11, 9, 11 \rangle$. In this case $\mathbf{C} = \{\langle 7.8, 5 \rangle, \langle 12.2, 5 \rangle\}$, which is clearly not an acceptable clustering. A better clustering would be $\mathbf{C} = \{\langle 6, 2 \rangle, \langle 10, 6 \rangle, \langle 14, 2 \rangle\}$

To overcome this problem some modifications are possible. The most simple is to fix where the centroids can be; in the example, let only multiples of the ϵ to be the centroids (such as: -5, 0, 5, 10, and so on).²⁸ Another approach would be to have, besides the threshold, the sample size that is considered statistical significant. Then before this sample is available the algorithm keeps all outcomes. Thus the behaviour would be: if the sample is not large enough then create the

²⁷ An ϵ depending on the region of $R(\bullet)$ may be needed. A simple example would be the Normal curve, where near 0 it is interesting to have a smaller ϵ than near the extremes.

 $^{^{28}}$ This approach simplifies the computational complexity of the algorithm but does not reflect the abstract concept of a centroid, which is the "centre" of the cluster.

cluster structure every step. Then, when it has reached a sample large enough, it switches to the Algo. 5.1 using the last created cluster as the starting structure. This last suggestion implies that the reward function acts as a random variable (if the PDF is not known) and that the sample size must be estimated/known.

5.5.1 Bias

It is also noteworthy saying that this clustering method is biased. The bias exists because it gives the same weight to new and old experiences. This means that a past too "bad" or too "good" result will be remembered as "fresh" and any other experience, even though it may not reflect the current world state, i.e. the rewards received recently. As it will be seen in chapter 8, this has a marginal impact on the results but may not be ignored. The improvement of the clustering method is left as future work.

5.6 Modified *Q*-Learning

The modifications necessary are a combination of what was presented as the single-state Q-Learning in Sec 5.4.2 and the clustering method from the previous section. As said before in Sec. 4.6, the editing phase has not been formalised and the standard clustering methods are not suitable for the task (Sec. 4.6.2). Therefore a new clustering method is proposed (Sec.5.5) to cope with the editing phase. This algorithm must be included into the modified Q-Learning. The complete modified Q-Learning is presented in Eq. 5.21. It is assumed that $\mathbf{C}_n \leftarrow \mathsf{Cluster}(\mathbf{C}_{n-1}, r_n)$ represents the centroid set at step n and $\mathsf{Cluster}(\bullet)$ the function in Algo. 5.1, then the modified Q-Learning algorithm is given by Eq. 5.18, 5.19, 5.20, and 5.21; where $v(\bullet)$ and $\pi(\bullet)$ are the same distortion functions from the PT specification (Eq. 4.2 and 4.3). It is worth noticing that the Q-Learning has no complexity increase because the only modifications are the inclusion of the set \mathbf{C} and the $pt(\bullet)$ function. The $pt(\bullet)$ function depends on the size of \mathbf{C} and on the function $\mathsf{Cluster}(\bullet)$, which for its turn only depends on the size of \mathbf{C} . Yet \mathbf{C} has a limited size (significantly smaller than the amount of iterations) meaning that no increase of computational complexity occurs.

$$pt(\mathbf{C}) = \sum_{\mathbf{c}\in\mathbf{C}} v(c_{\mathbf{c}})\pi\left(\frac{a_{\mathbf{c}}}{\sum_{\mathbf{c}\in\mathbf{C}} a_{\mathbf{c}}}\right)$$
(5.18)

$$\mathbf{C}_n = \mathsf{Cluster}(\mathbf{C}_{n-1}, r_n) \tag{5.19}$$

$$Q_n(s,a) = (1 - \alpha_n)Q_{n-1}(s,a)$$
(5.20)

$$+\alpha_{n} \left[pt(\mathbf{C}_{n}) + \gamma \sum_{s' \in \mathcal{S}} V_{n-1}(s') \right]$$
$$V_{n}(s) \equiv \operatorname*{argmax}_{a \in \mathcal{A}} \left[Q_{n}(s, a) \right]$$
(5.21)

To make it clear how the algorithm works a simple example is given step-by-step. The clustering parameter is $\epsilon = 5$ and in Tab. 5.2 and 5.3 is the evolution of each of the structures according to the received reward, in the second column of Tab. 5.2. For the *Q*-Learning parameters it is assumed: $\alpha_n = \alpha_0^n$, $\alpha_0 = 0.9999954$, $\gamma = 0.9$, and $Q_0(\bullet) = 0$ – which are the parameters suggested in [Lit94].

In the first table (Tab. 5.2) it is shown what happens with the centroid set \mathbf{C}_n as it incorporates each new reward (column " r_n " in Tab. 5.2). Then in Tab. 5.3 one can see how the \mathbf{C}_n set is converted into a prospect, where the accumulators are converted into probabilities, and what is the $pt(\bullet)$ value for this prospect (columns "prospect(\mathbf{C}_n)" and " $pt(\mathbf{C})$ " in Tab. 5.3, respectively). The last column in Tab. 5.3 shows the Q values at each step.

n	r_n	\mathbf{C}_n	Operation in \mathbf{C}
1	100	$\{\langle 100,1\rangle\}$	includes a new pair $\langle 100, 1 \rangle$
2	102	$\{\langle 101,2\rangle\}$	aggregates with pair $\langle 100, 1 \rangle$
3	101	$\{\langle 101,3\rangle\}$	aggregates with pair $\langle 101, 2 \rangle$
4	101	$\{\langle 101,4\rangle\}$	aggregates with pair $\langle 101, 3 \rangle$
5	110	$\{\langle 101,4\rangle,\langle 110,1\rangle\}$	includes new pair $\langle 110, 1 \rangle$
6	105	$\{\langle 101.8,5\rangle,\langle 110,1\rangle\}$	aggregates with pair $\langle 101, 4 \rangle$
7	113	$\{\langle 101.8, 5 \rangle, \langle 111.5, 2 \rangle\}$	aggregates with pair $\langle 110, 1 \rangle$
8	120	$\{\langle 101.8, 5 \rangle, \langle 111.5, 2 \rangle, \langle 120, 1 \rangle\}$	includes new pair $\langle 120, 1 \rangle$

Table 5.2: Evolution of \mathbf{C}_n (part 1)

Table 5.3: Evolution of \mathbf{C}_n (part 2)

\overline{n}	$prospect(\mathbf{C}_n)$	$pt(\mathbf{C})$	$Q_n(s,a)$
1	$\{\langle 100,1\rangle\}$	57.54	57.54
2	$\{\langle 101,1\rangle\}$	58.05	109.84
3	$\{\langle 101,1\rangle\}$	58.05	156.90
4	$\{\langle 101,1\rangle\}$	58.05	199.26
5	$\{\langle 101, 0.8 \rangle, \langle 110, 0.2 \rangle\}$	51.58	230.91
6	$\{\langle 101.8, 0.8 \rangle, \langle 110, 0.2 \rangle\}$	52.19	260.01
7	$\{\langle 101.8, 0.71 \rangle, \langle 111.5, 0.28 \rangle\}$	51.43	285.44
8	$\{\langle 101.8, 0.625 \rangle, \langle 111.5, 0.25 \rangle, \langle 120, 0.125 \rangle\}$	60.97	317.86

5.7 Summary

In this chapter the Q-Learning algorithm is approached and its modifications to conform the PT. It is argued that learning is necessary because the Markovian state transition function is not known beforehand and therefore the agents must learn it (Sec. 5.2). Then it is said that here the MDP approach is used (Sec. 5.3) because the amount of agents is high and it is not suppose that persons keep track of each single other competitor when its amount is high, which is also the opinion from von Neumann and Morgenstern (see quotation at the end of Sec. 5.3).

Following this discussion the standard Q-Learning is presented (Sec. 5.4) as well as the first step towards its modification to be PT based (Sec. 5.4.2). Then the final issue is approached: the editing phase (Sec. 5.5). It is also important to mention that the clustering method adopted is biased, as discussed (Sec. 5.5.1). With all necessary algorithm it is then possible to present the modified Q-Learning (Sec. 5.6), which includes a simple example to illustrate how the algorithm works. Therefore the complete Q-Learning proposal to tackle the PT based learning algorithm is presented in this chapter.

Chapter 6 Traffic And Route Choice

In this chapter the general aspects of traffic modelling and then the specific aspects of route choice modelling is presented. The modelling framework used here is the four-step model, which is addressed first in this chapter. When the "big picture" is shown then the specifics of the route choice problem are discussed and how it is tackled in this work. It also includes the mapping from traffic to Markov Decision Process (MDP), which is necessary since the MDP is the adopted modelling approach for the agent (Sec. 5.3).

6.1 Concepts Review

The only necessary concept necessary here is the MDP modelling adopted by the learning algorithm (Sec. 5.3). This means that it is necessary to map each element of the MDP to the corresponding traffic element. The basic property of an MDP is that it does no depend on historical evolution of the states, i.e. the next state depends only on the current state and the action of the agent. Specially important is to define which are the actions, the states, and the rewards.

6.2 Traffic Modelling

The goal of any given model is to better understand the modelled phenomenon. This means that a model is only useful if it casts some light over the phenomenon and helps to explain its observed behaviour and also to predict future behaviour under other conditions. In a nutshell: understanding and prediction are the key features of a good model. According to [HB07], for traffic this translates into understanding its dynamic and also its behaviour in a future condition.

The practical applications of a traffic model are diverse and depend on who is looking at the model. For traffic authorities the goal is to understand the network use (link¹ load) and how conditions will change, e.g., when a new highway is built or the impact of a new industry facility for the vicinities of it. For logistic companies it is to find optimal ways to navigate through the network with minimal cost (avoid overloaded links for example) and help to plan future paths according to the future conditions. For public transportation companies it is interesting for planning schedules, and, of course, profitability (how may persons are in each line at each time), to plan future extensions (such as new lines or schedules) to improve satisfaction and profit. For transportation researchers a model helps to understand how people make decisions to fulfil their transportation needs.

In the previous paragraph, except for the last example, all applications are basically interested into link load and its behaviour. This is also the focus in this text, since this seems to be the only available standard against which a model can be thoroughly evaluated. The reason relies on the

 $^{^{1}}$ Link is the term to specify a segment in a street that connects two crossings and link load is the measurement of the amount of vehicles in this segment at given time.

available data, which is mainly link load. Nevertheless some hope can be expected from projects such as the Berkeley's Mobile Millennium Project² that collects a comprehensive route decision dataset, although their findings are still to be published. For the time being, the only practical benchmark for traffic models is link load, i.e. a model is as good as it is capable of reproducing the network load.

6.3 The Four-Step Model

The most well-known traffic modelling technique is the four-step model [Bat07, McN07] (FSM). This model requires, as the name says, four steps to be fulfilled before modelling traffic. These steps are:

- **Trip Generation** Based on various data (such as socio-demographic and land use surveys) determine which are the potential origins and destinations.
- **Trip Distribution** When all origins and destinations are realised, it is necessary to establish how many individuals go from each origin to each destination. This means that the origin destination matrix (OD matrix) is specified.
- Mode Split Once all transportation needs are determined it is necessary to decide, for each trip, which transportation mode will be used. For instance how many individuals (for each OD pair) are using a private vehicle, a public transportation (and its kind), or another transportation mode such as car-pool or bicycle.
- Assignment When all previous steps are fulfilled, then the modeller must find the routes for each trip (with its corresponding mode) that better corresponds to the observed link loads (occupation).

Even though assignment is seen as a single step by the FSM, it can be sub-divided into other four steps [Bat07]:

- Route Choice For each OD pair a route will be assigned.
- Load Aggregation The link load will be aggregated based on the chosen routes, i.e. how many routes share each one of the links.
- **Capacity Evaluation** Since each link has a specific capacity, the corresponding side-effects must be taken into account, such as travel-time penalties.
- **Cost Evaluation** For each OD pair the travel cost, based on the routes, will be evaluated and given as feedback.

The whole process is usually iterative running until an equilibrium is reached, i.e. the individuals can no longer ameliorate their costs. The Wardrop [War52] equilibrium definition is normally used, which is also known as user equilibrium (UE). When the equilibrium is reached, the model can be evaluated, verifying how good it actually reflects the observations. Since link load/occupation is the most largely available data from a network this is used along with a fitness function (Sec. 2.5.1) to inform how the model fits the data. In the praxis this means that a model must ultimately reproduce the observed link loads through route choice modelling, although the ultimate goal is to use the model to forecast future conditions of the network or how will the load be distributed in some future time. The logic behind the model is this: if all steps have been correct and the model reproduces the observed network load (with an acceptable accuracy) then the decision model must be correct. This means that the route choice model also correctly captures how people decide to go from A to B (the transportation need). Thus, because the way how people make decision remains the same,³ when the scenario changes the results (link load) given by the model will be correct.

² http://traffic.berkeley.edu/, launched on November 10th of 2008.

 $^{^{3}}$ This is based on the utility theory, which says that the utility function is unlikely to change (chapter 2).

6.3.1 Route Choice Problem

From the four steps in the FSM the last one (assignment) is the most critical because it is the only that just observing the behaviour does not bring much information to the modelling process. In comparison, to find out which are the origins and destinations (the trip generation step) is rather simple and can be estimated by locating the several zones in the modelled region (living, shopping, and working zones) and any future situation can also be estimated by the land use planning. This information is all what is needed to build the necessary model, because the observations are the parameters. For the next step (trip distribution) the same is true, even though harder to acquire, knowing the amount of individuals in each zone is sufficient. These parameters can be collected from demographic and socio-economic surveys. The mode split can also be estimated by the number of registered vehicles, public transportation use, and so on. It is fair to say that this is not as easy as the previous two but the observation gives sufficient details about its nature.

The last step is the most difficult one because only observing it (assuming that all routes are known) does not give much insight on how the routes were chosen. Since the goal is to capture how people choose their routes. It is assumed that it is unlikely to change over time. Thus, having all routes provides the gold standard against which the model must be evaluated, not the model itself. This means that a choice model is still necessary.

A route choice is a discrete choice model, where a set of finite and countable choices (choice set) is given and the individual chooses the element that best suits him/her. If this choice problem can be modelled using a utility function it means that a mathematical function can be built that mimics the individuals' choices. Since a utility function is used it is meaningful to discuss about its modelling and inevitably about its states of equilibrium.

6.3.2 Equilibrium And Utility Functions

The concept of user-equilibrium is only meaningful with an associated fitness function (in this case the utility function) that says when a change in the choice does not improve the expected reward (the equilibrium definition). It is then meaningful to discuss about rationality and non-rationality in this context. According to [CS09], the rational equilibrium is said to be the objective user optimal decision and for the non-rational equilibrium it is said to be the subjective optimal decision. But regardless of the theory the optimal choice is always the one that yields the maximum utility, according to the used utility function.

In traffic as already discussed, the standard approach is to assume rational travellers, on the other hand it seems that the Prospect Theory (PT) is gaining more and more attention from the traffic researchers. But in a recent article [CS09] a continuous derivation of the Cumulative Prospect Theory (CPT) is investigated for modelling a five link network. The main argument for the use of the CPT is that it is – arguably – a better approach to model how people evaluate choices under uncertainty, i.e. when the costs are not known beforehand. On the other hand, the very existence of an equilibrium in traffic is questioned in [Goo98], where it is argued that an equilibrium cannot be supposed valid for traffic when the data is not collected in such state and no evidence can be shown to support that the system moves towards the equilibrium. That said, it is supposed here that an equilibrium exists and, as in [CS09], this equilibrium is the one established by the perceived cost, i.e. the utility function does not reflect necessarily the real cost of a route, but what the traveller/agent perceives as its cost, its subjective cost.

Regarding the utility function, the only variable is the travel-time. But before presenting the travel-time calculation method (in the next section) it is necessary to discuss how relevant the travel-time function is for the decision process. Recalling the discussion about the utility function in chapter 2 it is said that the utility function forms an ordinal scale and that any monotone affine transformation is possible. That said, the only important feature of a travel-time function for the route choice model is that it must be proportional to the link load, i.e. the more vehicles in the link the higher is the travel-time in that particular link. The argument is the following: suppose that the real travel-time function⁴ (if such modelling is possible) is an exponential curve,

⁴ The very nature of the travel-time as a function can be questioned because. Since it is independent from the

such as the one supposed by the Bureau of Public Roads [Tra00] (BPR) and that the drivers take only the travel-time into consideration. For the utility function side, only the ordering is relevant, i.e. how the links/routes are scored (which is the best, the second best, and so on). With this in consideration a utility function that uses the exponential values of the travel-time function is equivalent to another function that uses the logarithm of the travel-time, because the logarithm does not change the ranking of the links/routes (even though the actual utility values are different).

Now suppose that the travel-time function is hypothetically linear. Then again both the examples of utility function explained before are still equivalent and valid, because the travel-time is still proportional to the amount of vehicles and the utility functions still rank more occupied links below less occupied. This means that taking the point-of-view of the agents and the decision-making, the specific travel-time function is almost irrelevant. It must only reflect the same behaviour observed in reality, which is higher occupation implies higher travel-time.

This remark has a caveat nevertheless. If a travel-time function is used that is only loosely related to the real travel-time function then this model cannot be used to model problems where the departure and arrival times are relevant. In this case none of the so called day-activity planning models can use this function. Another issue is the lack of a specific travel-time function for urban traffic. The BPR function is explicitly said to only model highway traffic [Tra00]. A more specific modelling is attempted in [DG08] using fundamental diagrams [RPM04]⁵ for urban links in San Francisco (USA) and Yokohama (Japan). The approach was to extract four linear regions from the estimated fundamental diagrams and then to calculate the travel-time according to the linear region corresponding to the occupation. For compatibility reasons the BPR method is used, which is presented in Sec. 6.4.1.

6.4 Modelling Traffic Assignment

In this work it is assumed that the first three steps are already fulfilled. This means that the model needs as input the OD matrix (steps 1,2, and 3) and the objective is to model the traffic assignment. To accomplish this task some prerequisites are necessary. First, the choice set must be provided, i.e. the route set for each OD pair.⁶ The method used is the so called Link Elimination Shortest-Path [ACaERSMM93] (LESP) with one modification: only one link is eliminated per algorithmic run – this is formally presented in Sec. A.5. Second, a utility function structure is provided and here only the travel-time is taken into account. This means that it is supposed, in this first model, that only travel-time is relevant for the individual route evaluation. Since an iterative process is used to achieve the network's final state, an equilibrium definition is necessary. To cope with that both concepts are used: rational user-equilibrium (EU) and the subjective userequilibrium, as [CS09] defined it. The first definition is obviously used for the rational models and the second for the non-rational PT based model proposed.

As mentioned in the previous paragraph the only variable in the utility function is the traveltime, but travel-time is a cost. Then the simplest way to extract the worthiness of travel-time is in Eq. 6.2, where: r is the route being evaluated, $l \in r$ is a link used in the route r, and t_l is the time consumed to transverse⁷ the link l. In a simpler explanation: the utility of a route is equal the negative value of the route's travel-time that is the sum of the consumed time in each of its links. This leaves the t_l to be calculated that is done by the function $bpr(\bullet)$ (Eq. 6.3) explained

specific drivers in the link. Therefore the travel-time, per definition, is not a function (the counter-domain is not unique for the same domain value). This can be better understood if different quotes of aggressive (move faster) and non-aggressive (slower) drivers are mixed and, with a fixed amount of drivers each time, made to travel the same link. It is expected that this will produce different behaviours and travel-times as well.

⁵ Fundamental diagrams is a graphic plot that has the vehicular density as the x-axis and flow in the y-axis. Density, for its turn is simply the amount of vehicles currently on the link divided by the link capacity (the maximum amount of vehicles that fit on the link). The flow, on the other hand, is the amount of vehicles per time unit that travels the link.

⁶ Because in this first approach no modus is modelled, it is assumed that each OD pair has an associated modus. 7 Transverse is the technical term to designate the time consumed to travel a link.

 v_t

in the very next section.

$$u(r) = v_{travel-time}(r) \tag{6.1}$$

$$_{ravel-time}(r) = -\sum_{l \in r} t_l \tag{6.2}$$

The network is not, however, used as it comes to the simulation. First the input graph is transformed into a super-network (Sec. A.2) that stratifies the input graph according to its multiple modi. This means that each modus has its own network, with the relevant links for this particular modus, and then the different networks are interconnected, so that modus change is allowed. This structure provides a better way to identify and control the modus change when generating a path in the network.⁸ Therefore a modified shortest-path algorithm is used (Sec. A.4) for generating routes.

6.4.1 Travel-Time Calculation

For the travel-time calculation the formula provided by the Bureau of Public Roads [Tra00] (BPR) is used. The reason for its adoption is compatibility, i.e. because most of the commercial traffic simulation programs, such as VISUM,⁹ also use it. This function is in Eq. 6.3 below, where, according to [Avi06], $\theta = \rho = 2.0$, l_{load} is the occupation of link l, C_l is the capacity of l, l_{length} is the length of l, v_{max} is the maximum speed allowed, and l_{lanes} the amount of lanes of l.

$$bpr(l) = t_{freeflow}(l) \times \left(1 + \theta * \left(\frac{l_{load}}{C_l}\right)^{\rho}\right)$$
(6.3)

$$t_{freeflow}(l) = l_{length}/v_{max} \tag{6.4}$$

$$C_l = 0.3 \left(l_{lenght} l_{lanes} \right) \tag{6.5}$$

It is, however, worth noticing that this is not appropriated for urban traffic scenarios because it was explicitly developed for highway traffic. The travel-time curve for a single lane link with 100m is depicted in Fig. 6.1 and the fundamental diagram in Fig. 6.2.

6.4.2 Translating Route Choice To Link Load

As explained in Sec. 6.3 the method for translating route choice to link load is to collect all decision (routes) and to count how many drivers took the same link and then calculate the corresponding travel-time informing that back to the drivers. The formalization of this concept is made below:

- 1. Ask all agents about their route choices: $\mathbf{a} = \langle a_0, a_1, \dots, a_n \rangle$, where \mathbf{a} is the joint action vector and a_i is the route choice of agent *i*.
- 2. Reset all link loads: $\forall e \in E_S \Rightarrow e_{load} = 0$, where S is the graph representing the network, E_S the link set of the graph, and e_{load} is the link occupation (amount of agents occupying this particular link).
- 3. Set the new link loads: $\forall e \in E_{\mathcal{S}} \Rightarrow e_{load} = \sum_{a \in \mathbf{a}} \delta_e$, where *a* is a route (in this case a list of links) and $\delta_e = \begin{cases} 1 & iff \ e \in a \\ 0 & \text{other else} \end{cases}$.
- 4. Inform back to the agents the travel experience: $\mathbf{R} = \langle \mathbf{r}_0, \mathbf{r}_1, \dots, \mathbf{r}_n \rangle$, where $\mathbf{r}_i = \langle r_0, r_1, \dots, r_m \rangle$ and $r_j = bpr(e_j \in a_i)$ (Eq. 6.3). This means that the agent receives back a tuple (\mathbf{r}_i) where each element corresponds to the travel-time, which is calculated by function $bpr(\bullet)$, for all links (e_j) of the chosen action (a_i) .

 $^{^{8}}$ A Context-Free Grammar is used for this purposes (Sec. A.3).

⁹ http://www.ptvag.com/



Figure 6.1: Travel-time function

6.5 Traffic As An MDP

As said in Sec. 6.1 and explained in Sec. 5.3 the world is, for the agents, an MDP instance. This means that the traffic assignment elements must be mapped into the MDP elements. First, each step is defined as transportation event, i.e. a step encloses an entire trip/travel (from an origin to a destination). Second, the traffic network must provide the states and rewards to the agents.

To cope with the last requirement, the traffic network is transformed into a super-network [CFCLB03, FCHLvN03, FCvNB04], which is formalised in Sec. A.2. A super-network can be said to be a stratified labelled digraph, i.e. each modus is represented by one stratus where all links have the particular modus of the stratus. The links are also labelled according to which modus they have. This structure is useful because it not only provides means to assign special costs (such as the cost of buying a bus ticket) but also to provide a multi-modal navigation through the network. This means that using the super-network representation this work can be easily extended to provide not only the traffic assignment modelling but also the modus split modelling (left as future work).

The states, as already mentioned, are the OD pairs, i.e. a state is defined by an origin and a destination node/vertex. Therefore the action set is the route set for the given state (OD pair).

The last element of the MDP specification is the reward function. This is done by informing the travel-time experience of each link in the agent's chosen route. Formally, the traffic instance of the MDP tuple $\langle S, A, T, R \rangle$ is in Eq. 6.6. In this definition some other structures are defined, which are: S for the input super-network; V_S for the vertex set of S; AllPaths(•) an algorithm that generates all paths for a given OD pair (an example is in [MMS90]); r is a route; $l \in r$ is a link in route r; bpr(•) a function that returns the travel-time for the given edge/link.



Figure 6.2: Fundamental diagram

$$\begin{split} MDP &= \langle S, A, T, R \rangle \tag{6.6} \\ S &= \{ \forall \langle o, d \rangle \in V_S \times V_S | o, d \in V_S \land o \neq d \} \\ A &= \cup_{\langle o, d \rangle \in S} A_{\langle o, d \rangle} \\ A_{\langle o, d \rangle} &\subseteq \texttt{AllPaths}(o, d) \\ T(\langle o, d \rangle, r, \langle o^{'}, d^{'}) \rangle &= \begin{cases} 1.0 & iff \langle o^{'}, d^{'} \rangle \equiv \langle o, d \rangle \forall r \in A \\ 0.0 & \text{other else} \end{cases} \\ R(s, r) &= -\sum_{l \in r} bpr(l) \end{split}$$

In Eq. 6.6 the action set $A_{\langle o,d \rangle}$ is only a subset of AllPaths(o,d) because it depends on the path generation algorithm and, as argued before, the path generator used is the LESP (Sec. 6.4).

6.6 Summary

In this chapter the traffic modelling is approached. The modelling paradigm used is the four-step model (FSM), which is the standard approach (Sec. 6.3). The focus is on the route choice problem (Sec. 6.3.1), which is the last step in the FSM. It is also stressed that modelling traffic assignment usually means to reproduce the link loads observed in the data because that is normally what is available. In Sec. 6.4 it is discussed how route choices are translated into link occupation and then into travel-time experiences. In the last section, the mapping from traffic to an MPD instance is made. This step is necessary to make it possible to use the Q-Learning formalism (chapter 5) for modelling the route choice problem, which is the scenario used to evaluate this work.

Chapter 7

Agent Architecture

In this chapter the elements are put together into an agent architecture. These elements are the learning engine (Sec. 5.6), the split reasoning (Sec. 4.2.1), sensory (in this case travel-time feedback, Sec. 6.4.1), and memory management (explained in details in this chapter). In the proposed architecture the concepts of split reasoning (discussed in chapter 4) and non-rationality, as well as rationality, are incorporated. A quick note about the typefaces here employed: regular typeface is used as in "System 1" it refers to the idea evoked by the concept and the typewriter typeface appears as in "System 1" it refers to the architectural mechanism that, in this case, cope with the "System 1" concept.

7.1 Concepts Review

Here the only concept that is needed from the previous chapters is the bi-parted reasoning system (Sec. 4.2.1). The bi-parted system means that the decision-making is performed at two different levels. The first, called System 1, is the intuitive level and corresponds to the "quick-and-dirty" decision-making. This part accounts for the decision with which the agent is accustomed with, i.e. the decisions made by System 1 refer to problems that are more-or-less known by the agent and whose solutions are rapidly available. The second level is the System 2 that accounts for the decisions that need a more elaborated decision process. At this level the options are collected and analysed before any choice is made. This is a more cumbersome process and therefore the agent, when making the decision using the System 2, tries to truly optimise its choice.

7.2 Why An Agent Architecture

From what has presented it may seem unnecessary to have an agent architecture since the elements to be put together present low or no complexity. The main reason is to organise the concepts and make it simpler to understand how data, reasoning, and sensory fit together. Another one, not less important, is to make it simpler to extend this architecture into one that includes new features (some examples are given in chapter 10). The last major contribution of an architecture is to lay the foundations for implementation by explicitly pointing out the basic building blocks.

It also makes it simpler to show how the specific features used here are, such as the split reasoning (chapter 4) and the flexibility to switch between the Expected Utility Theory (EUT) and the Prospect Theory (PT). Some secondary features are: the switch between exploration (experimenting with the world) and exploitation (using the own world model to make decisions); as well as to have an intuitive decision-making level (required by the split reasoning model).



Figure 7.1: Proposed agent architecture

7.3 Proposed Agent Architecture

The proposed architecture is depicted in Fig. 7.1. In this architecture two reasoning levels are present and they are named System 1 and System 2 after [TK83]. As discussed in Sec. 4.2.1; the bi-parted decision-making engine is supported by several researchers and medical evidences, therefore here only two decision-making modules are used. The first is the intuitive level (System 1) that account for known situations and for those whose options are already available. The decision-making at this level is sometimes described as a decision that "pops-up" to mind and in general it is said to be "quick-and-dirty". This means that the decision is not necessarily the best (whenever the criteria are) but the best among the ones that are promptly available to the agent.

In System 2 on the other hand the highest reasoning is situated, where more options are analysed (or even generated) and their estimated outcomes taken into account. This means that the System 2 is responsible for, at a first glance, supporting the System 1 in the decisions that it (the System 1) cannot make. In this case System 2 acts as a solution provider when the System 1 falls short. Another function of the System 2 is to supervise the activities of System 1, i.e. if the outcome (or expected outcome) of System 1 is below the agent's acceptance level (Sec. 4.2.1) it takes over the decision processes and corrects the System 1's decision.

Another point is that here a split memory mechanism is proposed where each module has its own memory. The Short Term Memory is the only memory available to System 1 while System 2 has access to both, the short and long term memories. The reactive behaviour is in System 1 and, as explained in Sec. 4.2.1, the System 2 is only triggered when System 1 performed poorly or does not know how to decide, which is not the case in reactive responses. In this model the supervision (dash-dotted line identified with Supervision in Fig. 7.1) is explicit depicted as well as the option supplier (the extra command line Options? from System 1 to System 2 in Fig. 7.1).

7.3.1 Memories

As discussed in Sec. 4.2.1, the bi-parted reasoning system asks for a decision from the intuition level (System 1) and this first decision can be corrected by the reasoning level (System 2) and that if the System 1 cannot solve the problem it is escalated to the System 2. It is also known [AF00] that as easy an information is available as probable it is that it will be used – this is called accessibility bias. In practical terms it means that if a decision must be made and the short-term memory already contains any options it is probable that the choice this will be made among them. Then, if no option is promptly available, the System 2 will be triggered to supply an option set from which a decision will be made. Therefore to support this decision-making mechanism a suitable memory organisation must be provided. In this case the memory is split into two: short and long
term memories.

Short Term Memory

To support the intuitive decision-making module (System 1), the Short Term Memory is provided and it has some defined features. First, it is restricted in size, i.e. how many options it can hold. Therefore, according to [Mil56], the Short Term Memory is limited to 7 items per context, and the context here is represented by a state, i.e. each different state is interpreted as a different context. Specializing this concept for traffic, it means that each origin destination pair (OD pair) corresponds to a state and, the routes being the available actions, each state holds a maximum of 7 routes. So the state "home-to-office" holds a maximum of 7 routes and "office-to-groceries-store" holds a different set of routes, which again can hold a maximum of 7 routes.

Along with the short term memory items a score (a single numerical value) is also kept with each one of them. This score is the "intuitive" utility of each item. This means that when the System 1 makes a decision it ranks each of the items in the Short Term Memory and chooses the very best according to the scores/utilities associated with each item. For the current application the scores kept with each item are the latest evaluation made by the Q-Learning, i.e. the $V_n(s)$ values (Eq. 5.3 and 5.21). This way the supervision of System 2 is "pro-active" and keeps the System 1 always up-to-date with the actual utilities.

It must be noticed that this is not entirely correct under the psychological point-of-view, because a separation between them exists and the scores in the short term memory are not promptly updated. This means that the System 2 must have an internal threshold that establishes when the utility, or expected reward, does not meet the "acceptable" standards. However, no model for this behaviour is available and establishing an "ad-hoc" interpretation of this will introduce another variable in the model and one that will probably be wrong. Thus this lack of decoupling between System 1 and System 2.

Aside from this, the short term memory suppose to be up-to-date, i.e. it must contains relatively "fresh" items. It means that items that are old are not supposed to be in it. To cope with that, in Short Term Memory the items are also tagged with their age, i.e. how long the item in the Short Term Memory was not being used. In practice, this means that each time the agent is in a given state and it makes a decision the non-chosen options have their age increased and the chosen has its age set to zero. As soon as an item reaches the maximum age threshold it is excluded from the Short Term Memory. This threshold, in the application, is set to be a week, supposing that each state is reached once a day. This means that, if a route hasn't been used a week long it no longer belongs to the short term memory, it is not "fresh" any more.

Another feature of the Short Term Memory is that if the System 2 decides that new items are necessary and the memory is full, then items must be excluded. This means that the items in the Short Term Memory can be recycled and items with poor performance (with the lowest scores) are removed to make room for new items coming from System 2.

Long Term Memory

For the long term memory no limit is built in and it can hold as many items as available. This memory is the information storage, i.e. all data available and necessary for the high reasoning decision-making (System 2) is registered there. This is more a general storage space without a preestablished form. Translating to the used Long Term Memory for the traffic scenarios, it contains the $Q(\bullet)$ function/table (Eq. 5.2 and 5.20) from the Q-Learning algorithm as well as the vectors C (Eq. 5.19 and Algo. 5.1). It also contains a table that associates links with travel-times, i.e. for each used link its travel-time is recorded and stored for use of the route generation module (explained later). As it can be noticed, the Long Term Memory is heterogeneous and can be seen as collection of different types of information sets.

7.3.2 Action Choice Generation

The choice set, i.e. the set of the possible actions for any given state can be beforehand given or on-demand generated. Either way it must be provided to the System 2, which is not responsible for its generation. The strategy adopted depends on the which algorithm is used to be the System 2. Because here the *Q*-Learning is used and the test scenario is traffic assignment some advantage can be taken from this set-up. In essence the choice generation is, for traffic assignment, a route generator and for this task the algorithm used is a variation of the Link-Elimination Shortest Path (LESP), explained in Sec. 6.4. Then a simplification is made: not all possible routes are generated for each possible OD pair, but enough routes to fill the Short Term Memory.

Then, as the routes are being forgotten or recycled (as explained before), new routes are generated on-demand, at the request of System 2. This event, of demanding new routes, are only triggered when the exploring behaviour is activated (introduction of chapter 5). Where new options are generated until the Short Term Memory capacity is exhausted and than an action is chosen randomly from the options available in the Short Term Memory. This is explained later when presenting the System 1.

7.3.3 Environment

It was already explained, in Sec. 5.3, that the environment is assumed an MDP instance. This means that the **Perception** receives, as feedback, the current state and reward received for the previous chosen action. It also means that the **System 1** must inform an action for the current state for each simulation step. In traffic, the states are the transportation needs, i.e. the OD pairs, for which the agent must choose a route. As a feedback it also receives the reward for the last route (a function of the travel-time as explained in Sec. 6.4). A slight modification is made here, where the agent receives the travel-time reward corresponding to each of the links that compound the chosen route and the final reward is the sum of links' travel-times converted to utilities (Eq. 6.2).

7.3.4 System 1

Some of the aspects of the System 1 were already mentioned but its internal mechanisms were not fully explained, which is the subject of this section. The System 1 is meant to behave as a "quick-and-dirty" decision-maker (an effortless decision) and it is helped by the Short Term Memory. Thus, the basic algorithm implemented for System 1 is in Fig. 7.2.

The algorithm is rather simple; it just selects the best option currently available in the Short Term Memory. In the case none is available it requests a new option from System 2. When the decision is made it is then informed to the Environment. The decision's score is, in the next step, then up-dated by the System 2 that becomes the feedback from Perception.

7.3.5 Situation Recognition

Before explaining the System 2 it is necessary to discuss the Situation Recognition. The main task is to inform the decision-making modules, System 1 and System 2, about the current state and that a new decision is necessary. Then it looks if the System 1 can make the decision, i.e. if the corresponding Short Term Memory has any item in it. In case it does, the System 1 is responsible for making the decision. But if the corresponding Short Term Memory is empty then the System 2 is informed to make the decision.

Another responsibility of the Situation Recognition is to forward actions' feedbacks, i.e. rewards/costs from the last action, to the System 2. It is also in this module that the "time" track is kept, i.e. it is responsible for keeping track of where in time the agent is. In the MDP world it means to know what is the current step.

This time tracking function also incorporates the ageing mechanism of the Short Term Memory. Therefore, the Situation Recognition recognises that time is passing and that the corresponding entries in the Short Term Memory must have their age increased. Thus, if any entry reaches the age



Figure 7.2: System 1

threshold the System 2 is informed about this and acts accordingly. The basic ageing algorithm is presented in Fig. 7.3.

7.3.6 System 2

The System 2 (Sec. 4.2.1) is the highest reasoning processes. This means that at this level the cumbersome reasoning is made, where, when necessary, information is analysed and a more elaborated decision is made. At the System 2 the decisions that cannot be handled by the System 1 are made. It also corrects/supervises the activities of the System 1.

By cumbersome reasoning it is meant to analyse a multi-dimensional information space and to extract from it the relevant choices that lead to a decision. In the case of Q-Learning, it means to analyse the multi-dimensional $Q(\bullet)$ table/function, evaluating the best choice and, in the case of the PT modified Q-Learning, also to analyse the various vectors **C**. For the traffic specific case, it also means to provide the route generator (Sec. 6.4) with updated link weights, which are the travel-times collected as feedback by the **Perception** and stored in the Long Term Memory.

The supervision is made in two opportunities. The first is the already mentioned "pro-active" supervision that updates the Short Term Memory scores when it receives the feedback from the last action. The second supervision or correction is when the System 2 decides to explore instead of exploit. Then System 2 takes over the decision process and randomly chooses an option to be the next action. It can also ask the action generator for more options, which in case of traffic means new routes, when they are necessary.

7.4 Summary

Here the agent architecture is presented. Its main objective is to lay the foundations for the practical application of the different concepts presented here. How to organise the memory, the decision flows, and how the different modules interact with each other (Sec. 7.2). This architecture is designed to accommodated the envisioned agent aspects (Sec. 7.3). This means that an MDP



Figure 7.3: Memory ageing process

world is assumed and special attention is given to the interaction between the two reasoning systems: System 1 and 2 (Sec. 7.3.4, 7.3.5, and 7.3.6).

Chapter 8

Evaluation

In this chapter the results and evaluation of the agent architecture and learning algorithm are made. The objective is to evaluate in which conditions the Prospect Theory based Q-Learning diverge from the rational behaviour. Then to verify if it is indeed better at reproducing real data and has a last point to see if the whole framework can scale to a real world scenario.

8.1 Concepts Review

The first point is the Q-Learning algorithm (chapter 5) an its modification to be based on the Prospect Theory (Sec. 5.6). An important issue is the clustering algorithm (Sec. 5.5) used to cope with the editing phase necessary but not formalised for the Prospect Theory (Sec. 4.6). One may not forget that the clustering algorithm is biased and this bias was discussed in Sec. 5.5 and again when presenting the results (Sec. 8.8.1). Another concept is the bi-parted reasoning system (Sec. 4.2.1) which is also approached by the architecture (Sec. 7.3.6). The bi-parted system refers to the split of the decision-making process into two levels. The first, called System 1, is the "quick-and-dirty" decision-making that is triggered when the problem is well known by the agent and whose options are promptly available to the agent (in the short-term memory, Sec. 7.3.1 and 7.3.4). The second is called System 2 and here accounts for the learning algorithm, i.e. the Q-Learning (chapter 5).

8.2 Evaluation Methodology

The objective of the experiments done and described here are to first show how and in which conditions the Prospect Theory (PT) based Q-Learning (chapter 5) diverge from the rational behaviour. This means that the main goal is to provide the conditions under which the rationality is violated and why. For the experiments the traffic scenarios were adopted. Here three algorithms were used for the System 2 reasoning engine: the standard Q-Learning (*STD*), the so-called clustered Q-Learning (*CluEUT*), and the clustered PT based Q-Learning (*CluPT*). The standard Q-Learning (Sec. 5.4) is the gold standard for the rational behaviour, i.e. this is the algorithm that behaves according to the von Neumann definition of rationality [vNM07] and this was proven correct in [WD92]. This means that if the presented architecture works then the standard Q-Learning must converge to the expected rational behaviour. Second, because the clustering method, presented in Sec. 5.5, was not proven to be correct the claims must be supported by evidences. Therefore, the standard Q-Learning was modified to use the clustering method (Sec. 5.5) proposed for the editing phase (Sec. 4.6) and it must, as the standard Q-Learning, converge to the rational behaviour. Then, assuming that the previous two assumptions are correct, the PT modified Q-Learning (Sec. 5.6) can be evaluated.

What is meant by the previous paragraph is the following. Given a simple enough scenario, where the expected rational behaviour is known, the validity of the agent architecture can be

demonstrated. This is done by using a reasoning that has a well known behaviour and then comparing this expected behaviour with the results produced using the architecture. Then the distortions, if any, imposed by the architecture can be evaluated by analysing the deviations observed in the simulation results and the expected results. For this step the standard Q-Learning was chosen as the gold standard.

The second point is to evaluate the distortion introduced by the clustering method, i.e. to verify how biased it is. This bias can be measured by comparing the results of two equivalent reasoning engine: one with and the other without the clustering. These engines are the standard Q-Learning and the modified clustered Q-Learning. If the results of these two engines are not equivalent then the clustering method can be said biased and that it introduces distortions in the reasoning of the agent.

The last point is, if the architecture does not introduce any "artefact" then no deviation is expected between the simulated and predicted results. If the clustering method also shows no or an acceptable bias, then the PT based engine can be evaluated. Unfortunately the validity of the PT can only be established with real data, which in this case is sparse. To tackle this the data from [SSC+05, Chm05] is used to evaluate the PT based Q-Learning, but this experiment does not give any clue about the working of the PT in traffic and which are the critical conditions for it. Therefore, some extended experiments were done to show how the PT based Q-Learning operates and in which conditions it deviates from rationality.

Two scenarios are proposed, the commuter scenario with two routes, which is inspired by the El Farol Bar [Art94] Problem, and the city of Burgdorf (Switzerland). The commuter scenario is chosen because some real data is available in [SSC+05, Chm05] and also because it is simple enough to be understood (Sec.8.8 and 8.9). Using this scenario all analysis are made, investigating several different set-ups to verify the influence of the PT in traffic. Besides this scenario, the city of Burgdorf is also simulated (Sec. 8.10) because some data is available and its size can show that the simulation framework do scale. However, the main objective of simulating Burgdorf is to show that the agent architecture can scale to real-world scenario, not to fit the data.¹

8.3 Different Algorithm For System 2

The different variations of the Q-Learning used in the experiments are: STD referring to the standard Q-Learning, CluEUT to the cluster modified Q-Learning, and CluPT to the clustered PT based Q-Learning. A compact view of the different modifications in the Q-Learning algorithm is presented in Eq. 8.5. There $V_n(s)$ is the decision for state s at step $n, a \in \mathcal{A}$ is a route, α_n is the learning factor for step n, r_n is the travel-time for the last route, $d(r_n)$ is the reasoning selection function, and $Cluster(\bullet)$ corresponds to Algo. 5.1.

$$eut(\mathbf{C}) = \sum_{\mathbf{c}\in\mathbf{C}} c_{\mathbf{c}} \frac{a_{\mathbf{c}}}{\sum_{\mathbf{c}\in\mathbf{C}} a_{\mathbf{c}}}$$
(8.1)

$$pt(\mathbf{C}) = \sum_{\mathbf{c}\in\mathbf{C}} v(c_{\mathbf{c}})\pi\left(\frac{a_{\mathbf{c}}}{\sum_{\mathbf{c}\in\mathbf{C}} a_{\mathbf{c}}}\right)$$
(8.2)

$$d(r_n) = \begin{cases} r_n & \text{standard } \mathcal{Q}\text{-Learning} & STD \\ eut(\text{Cluster}(\mathbf{C}_{n-1}, r_n)) & \text{EUT based } \mathcal{Q}\text{-Learning} & CluEUT \\ pt(\text{Cluster}(\mathbf{C}_{n-1}, r_n)) & \text{PT based } \mathcal{Q}\text{-Learning} & CluPT \end{cases}$$
(8.3)

$$Q_n(s,a) = (1 - \alpha_n)Q_{n-1}(s,a)$$
(8.4)

$$+\alpha_{n} \left[-d(r_{n}) + \gamma \sum_{s' \in \mathcal{S}} V_{n-1}(s') \right]$$
$$V_{n}(s) \equiv \operatorname*{argmax}_{a \in \mathcal{A}} \left[Q_{n}(s, a) \right]$$
(8.5)

¹ When this scenario is presented, it will be explained why it cannot be used for the PT evaluation (Sec. 8.10).

In the reasoning selection function $d(r_n)$ (Eq. 8.3) some clarification is necessary. First it is has a "minus" in Eq. 8.4 because the reward is actually a cost (travel-time) and therefore the effective reward is the negative of the cost (Sec. 6.4). Second, this function triggers the reward manipulation according to the reasoning engine being in use (*STD*, *CluEUT*, or *CluPT*). The first instance, the *STD*, the working is the one corresponding to standard *Q*-Learning where the reward is used directly. The second case, the *CluEUT*, is the clustered version of the *STD*. This modification works by, as the *CluPT*, accumulating the rewards received into the cluster structure **C** and uses the same function to perform the clustering: the function $Cluster(\bullet)$ (Sec. 5.5). But the difference between *CluEUT* and *CluPT* is that the former calculates the rational utility derived from the clustering structure (the function $eut(\bullet)$ from Eq. 3.3 in Sec. 3.3) and the *CluPT* calculates the non-rational utility from the PT (Eq. 4.1).

The function of CluEUT is to verify if the clustering method is biased and by how much. This is so because CluEUT is supposed to yields the same results as STD, i.e. to behave rational.

8.4 Common Parameters Across The Experiments

The parameters used for all experiments are presented in Tab. 8.1, where the reasoning engines correspond to the ones presented in Eq. 8.3.

Parameter	Value
Short-term memory capacity	7
Short-term memory maximum age	7
Long-term memory list capacity	∞
Reasoning engines	Eq. 8.3: $\{STD, CluEUT, CluPT\}$
Q -Learning α	Eq. 5.5: $\alpha_{init} = 1$ and $\alpha_{final} = 0.1$
Q-Learning exploration	Eq. 5.5: $\alpha_{init} = 0.2$ and $\alpha_{final} = 0.1$
Q -Learning γ	0.9
\mathcal{Q} -Learning $V_0(\bullet)$	0.0
\mathcal{Q} -Learning $Q_0(\bullet)$	0.0
Cost function	Eq. 6.3: $bpr(\bullet)$
Reward function	Eq. 6.2: $-bpr(\bullet)$

Table 8.1: Simulation parameters

In all experiments, before any data were collected, one last simulation step was run, where no learning or exploring occur. This means that the agents, on this one last iteration, are only exploiting their knowledge. This strategy is used to eliminate any stochastic "artefact" that does not come from the agent's accumulated knowledge, i.e. that does not come from the individual $V_n(s)$ (Eq. 8.5).

8.5 Calibration

An equilibrium can be defined as the existence of an attractor [Mil85] and the probability of reaching this attractor is 1. This assumes the existence of a stable state space towards which the function being calibrated inexorably moves. The problem using MASim is that this equilibrium is usually not a specific point in the search space, it is likelier to have a multi-point or a steady state equilibrium. This is an issue for most of the standard calibration methods available because they expect a single point for each parameter configuration set. To make it more clear, standard calibration methods expect the fitness function to be an actual mathematical function, i.e. only one point in the codomain for each parameter set. This is not true for MASim, where a parameter set yields potentially multiple points in the codomain, one for each simulation.

The problem with the equilibrium has no easy solution for MASim. One alternative, which may not be acceptable, is to increase the parameter's step to the point where all steady states are disjointed. This makes the multi-point equilibrium behaves as a single point because of the "disjointedness" of it. In this case any gradient oriented procedure will not get "confused" about which configuration set is better than the other.

The previous method has two main disadvantages. The first is that it may miss the optimal solution because it allows only disjointed configuration sets. Second, it assumes that the step size for "disjointedness" is known beforehand, which can be hard to estimate.

Another way to bypass the hysteresis issue (multiple equilibrium points) is to aggregate several simulation runs for the same configuration set and let this aggregated point be the result. This, however, makes the calibration process even more cumbersome, bringing two other problems: how to aggregate and how much simulations are enough. Unfortunately no general valid answer can be used. In some cases, for instance, the mean value suits the problem and in others the median is more appropriated. For the amount of simulations the answer is easy but not simple: to make as much repetitions as necessary so that the aggregation points are disjoint.

The calibration of MASim is the subject in [FKP04, FKP06] where a top-down approach with feedback is adopted. The idea is to break the fitness function into multiple sub-functions (each tackling a facet or region of the search space) and then calibrating the agents in groups according to these sub-goals. As the sub-goals are optimised/calibrated so is the macro level (the final calibration). This method greatly reduces the search space but not all scenarios can be optimised using this procedure, as in traffic assignment.

One way to break the calibration task into sub-goals is to approach the calibration problem by calibrating each OD pair as a sub-task. The problem with it is that the resulting calibration has a multitude of parameters, potentially one set for each OD pair, which greatly increases the scenario analysis complexity. This approach is however an improvement over using the traditional methods in the way described before, which requires several simulations or making the parameters discrete enough.

8.6 Difficulties In Comparing With Microeconomics

Besides the theoretical fundamentals for the validity of this discrete choice modelling framework one could speculate about its experimental validity. An obvious method to support the validity with evidences would be to use this framework (with the rational reasoning engine) and compare its results for the same problem modelled by a microeconomic model (Sec. 3.4). The appropriated model to use is a Multi Nominal Logit (MNL) model (Sec. 3.4), which does not model the correlation among the options. Then, assuming similar results, it can be said that the frameworks are equivalent.

The two main issues of this methodology are: first, the option generation and, second, the calibration methods. The latter is related to the calibration problems when using Multi-Agent Simulations (MASim) (Sec. 8.5) but it can be bypassed with the appropriated adjustments. Yet both models (microeconomics and MASim) must be calibrated under the same techniques and configurations, i.e. same fitness function and parameters' variation steps. This is more or less a matter of attention to the details. The first issue is rather more complicated.

In one sentence, the option generation for both models must be equivalent. This means that they must use the same algorithm under the same metrics. This imposes a practical difficulty that is on transferring the used infrastructure (the super-network, Sec. 6.5 and A.2) and the manipulation algorithms (the route generator, Sec. 6.5 and A.5) to another implemented framework (such as the BIOGEME² [Bie03, Bie08] or R³ [CNZ99]). Therefore this comparison is let out of the scope of this work, since microeconomics are models for rational choice as well as the standard Q-Learning [Wat89, WD92], which is equivalent to the MNL.

 $^{^2}$ http://biogeme.epfl.ch/

³ http://www.r-project.org/



Figure 8.1: El Farol traffic scenario

8.7 Results Analysis

The concern is on the investigation about the applicability of a non-rational model for traffic assignment. This means that the user equilibrium state from both rational and non-rational models are compared.

The objective is to compare the user equilibrium, or Wardrop's Equilibrium [War52], from both models (rational and non-rational). If both have the same equilibrium point this means that both are most likely to be equivalent and therefore the use of non-rationality through the PT does not give an alternative to the EUT and so no need for this framework (since it is more complex than an MNL model). But, on the other hand, if the equilibrium points differ it means that the PT is indeed an option worth of investigation. To answer these questions several experiments (under different configurations) were made to identify the conditions necessary for the PT to deviate from EUT, i.e. when the PT is relevant (Sec. 8.8). To verify the algorithms' practical applicability, the Selten scenario (Sec. 8.9) using the data from [SSC⁺05, Chm05] was simulated. Finally, the scalability of the framework is analysed in Sec. 8.10.

Notice that for the El Farol Bar inspired scenario (Sec. 8.8) and for the Burgdorf scenario (Sec. 8.10) the travel-time is calculated by $bpr(\bullet)$ function (Sec. 6.4.1). But in the Selten scenario (Sec. 8.9) a different travel-time function is used, which is presented and explained in the corresponding section.

8.8 El Farol Bar Inspired Scenario

The El Farol Bar [Art94] problem is a minority game, i.e. several agents must choose among two options, but the more agents choose the same choice the worse they will perform doing so. In the original problem, two choices are given to the agents: go to the bar (the El Farol bar) or stay at home. The bar had a limited capacity, so if the amount of agents in the bar is higher than the ideal capacity, then the agents that decided to stay at home win. If, on the other hand, the capacity was not exceeded then the agents at the bar win.

To transfer for the traffic is rather simple and already made in $[BBA^+00]$, there two routes were given and one route had a higher capacity than the other. Then, if the agent is in a route whose capacity has been exceeded it experiences congestion and is penalised with a higher travel-time than the agents on the other route, whose capacity has not been exceeded. This scenario is a binary choice where agents must achieve equilibrium and the rational choice says that the agents tend to reach the perfect split, where both routes have the same travel-time. The hypothesis in this experiment is that the rational behaviour (EUT) reaches this equilibrium but the non-rational does not.

A schematic view of the scenario is in Fig. 8.1, where the main route is identified by Main and the secondary by *Secondary*. The origin is identified by O and the destination by D.

8.8.1 Clustering Bias

It will be evident in the results that the clustering made by function $Cluster(\bullet)$ (Algo.5.1) is biased. The bias is towards the *Main* route, meaning that it tends to penalise the *Secondary* route. The reason for this is in its "memory" of bad results. The *Secondary* route is two times

more sensible than the *Main* route because the first has half the capacity of the second. Then if x agents exceed the perfect split in *Secondary* they are more severely penalise than if they where in the *Main* route. This discrepancy associated with the clustering "memory" tend to "poison" the lottery/prospect associated with the *Secondary* route.

However, the bias does not vary too much and is fairly homogeneous in the experiments. Therefore, in the results an extra-column was include (μ_{CluPT}^b) that represent the deviation of the CluPT, excluding the bias observed in CluEUT. This is calculated by Eq. 8.6, where from the values in CluPT the bias observed in CluEUT is removed (μ_{CluEUT}^b) . The bias in the CluEUT (μ_{CluEUT}^b) is the difference between the expected value (μ_{EUT}) and the actual obtained value (μ_{CluEUT}) . To overcome the bias it is necessary, as explicitly shown in the next section, several simulation steps, 1000 or more.

$$\mu^b_{CluPT} = \mu_{CluPT} - \mu^b_{CluEUT} \tag{8.6}$$

$$\mu_{CluEUT}^{o} = \mu_{CluEUT} - \mu_{EUT} \tag{8.7}$$

The same effect does not happen in the STD because the learning factor α_n guarantees that "bad" scores are attenuated as the agent get experienced with the routes and these "bad" scores are rather seldom. This, however, is not the case with the clustering function. There the scores receive no special treatment if they are old or new, they are all "remembered" the same. No simple solution was found for this issue and therefore left as future work.

To make the bias evident, in all results the error is given in field "err" by each reasoning engine. This error is calculated by: $err = \mu_{\bullet} - \mu_{EUT}$.

8.8.2 Scenario Experiments

This simple scenario was used to investigate the influence of the different parameters. The parameters were: simulation horizon, amount of agents, and target density at equilibrium. For all experiments the theoretical expected value is given, identified by μ_{EUT} .

The horizon experiments were designed to verify which is the minimal simulation horizon necessary to reach behavioural stability, i.e. the minimum horizon where this variable (the horizon) does not bias the agent's behaviour. It was also verified if the agents still remain in the same behaviour when this horizon is stretched beyond the minimum.

For this scenario, when the perfect split (2/3 for the Main and the remaining 1/3 for the Secondary route) is reached all agents experience the same travel-time regardless the chosen route. It also follows that both routes have the same density.⁴

The next experiment varied the agent amount and was designed to evaluate if the amount of agents has any influence on the agents' behaviour, i.e. if the deviations from rationality are attenuated or accentuated by increasing the amount of agents.

The so called target density at equilibrium experiments were made modifying the routes length according to the amount of agents to yield the "target density", when the perfect split is reached. These experiments are helpful to evaluate the agent's behaviour under free-flow, regime, and congestion conditions, i.e. how they respond (according to the reasoning engine used) to the different stages that the traffic conditions can assume. It can also show when the PT based Q-Learning deviates from the rational behaviour.

All simulations were repeated 100 times and the results aggregated into the mean occupation (μ_{\bullet}) and the standard deviation (σ_{\bullet}) . When a parameter is varied all others are kept fixed. For each experiment set the complete configuration is presented. The parameter that are not explicitly given are the same as in Tab. 8.1 from Sec. 8.4, such as the travel-time cost function, which is the $bpr(\bullet)$ (Eq. 6.3 in Sec. 6.4.1). In all experiments both routes have the same length but the *Main* route has two lanes and the *Secondary* only one. Another point is why, excepting for Sec. 8.8.2, the

⁴ Density means that each route can accommodate a limited amount of vehicles and therefore the density says how much of this amount is occupied. If a route have space for, say, a maximum of 100 vehicles and 49 are using this route in a given moment, then the corresponding density is 49/100 = 0.49.

target density is 0.3. This value was chosen because it is located right after the congestion region of the travel-time function $bpr(\bullet)$ (Sec. 6.4.1). This can be appreciated by looking at Fig. 6.2 where 0.3 is right after the curve's "knee" and this value seems to highlight the PT behaviour.

For the clustering method the value was adjusted according to Eq. 8.8, where sp is the shortest path and length(sp) returns the length, in meters, of the sp. The maximum speed is $v_{max} = 120km/h \sim 33.3m/s$. Because the route length depends on the amount of agents, given that target density is fixed, so does the clustering parameter. For 100 agents it corresponds to 33.3s for a target density of 0.3.

$$\epsilon = 2 \frac{length(sp)}{v_{max}} \tag{8.8}$$

Horizon

The horizon experiments are designed to verify if the route choices are influenced by the amount of experience accumulated by the agent, i.e. if the behaviour is experience dependent. It also presents an opportunity to verify if the clustering method for building prospect can distort the decisions, comparing the standard Q-Learning with its cluster based equivalent. The fixed parameters are presented in Tab. 8.2 and the experiment variables in Tab. 8.3. The results for this experiment are shown in Tab. 8.4, where μ_{\bullet} refers to the mean occupation and σ_{\bullet} to the standard deviation.

Parameter	Value
Agent amount	100 agents
Target density	0.30
Route length	555.6m
Route capacity Main/Secondary	222/111 vehicles
Scenario capacity	333 agents
Equilibrium travel-time	19.67s
Equilibrium mean speed	28.25m/s(101.7km/h)

Table 8.2: Fixed parameters for horizon experiments

Table 8.3: Simulation parameters for horizon experiments

33.33*s*

Parameter	Values
Simulation horizon	$\{10, 50, 100, 500, 1000, 1500, 2000\}$
Reasoning engine	$\{STD, CluEUT, CluPT\}$

From the results in Tab. 8.4 it can be seen that a horizon higher than 50 is not necessary for the STD, which is the reasoning gold standard. But that is not true for STD and CluEUT, meaning that the clustering method is biased. This bias is however similar across the experiments, as argued before. Another interesting feature is that the standard deviation does not necessarily decreases with the increase of the horizon, showing that multiple simulations and aggregation is necessary to extract the real behaviour.⁵ The results also show that CluPT ($\mu_{CluPT}(\sigma_{CluPT})$) consistently deviates – for horizons higher than 50 – from its rational counterparts (STD and CluEUT). The standard deviation is fairly similar between CluEUT and CluPT (for horizons higher than 50).

The results for a horizon lower than 50 is somehow divergent, to say the least. An explanation could not be found to justify them, except that with less than 50 iterations is too low for the learning algorithm to adapt itself to the scenario, which is a known issue for learning algorithm, they need the necessary experience to start profiting from the environment.

Clustering ϵ

⁵ Note that the σ_{\bullet} does not come from the exploration rate (Sec. 5.4.1).

Horizon	$\mu_{STD}(\sigma_{STD}): err$	$\mu_{CluEUT}(\sigma_{CluEUT}): err$	$\mu_{CluPT}(\sigma_{CluPT}): err$	μ^{b}_{CluPT}
10	51.04(6.92):-15.63	51.35(6.98):-15.32	50.98(7.15):-15.67	66.30
50	65.71(16.16):-0.96	73.30(8.20):+6.63	78.0(8.65):+11.33	71.37
100	66.82(8.44):+0.15	73.93(7.91):+7.27	83.24(9.22):+16.57	75.98
200	66.80(7.07):+0.13	76.06(8.08):+9.39	83.85(9.01):+17.18	74.46
300	66.39(7.01):-0.28	75.76(8.04):+9.09	85.34(9.11):+18.67	76.25
400	66.13(6.87):-0.54	72.70(7.69):+6.03	84.77(9.10):+18.10	78.74
500	65.94(7.01):-0.73	71.41(7.52):+4.74	83.22(9.04):+16.55	78.48
1000	66.12(6.99):-0.55	67.69(7.23):+1.02	75.44(8.19):+8.77	74.41
1500	65.65(7.10):-1.01	67.06(7.13):+0.39	74.36(8.08):+7.69	73.97
2000	66.24(7.12):-0.41	66.86(6.97):+0.19	74.51(8.30):+7.84	74.32
μ_{EUT}	$66.\overline{6}$			

Table 8.4: Occupation results of the *Main* route for horizon experiments with 100 agents

However, the critical amount of simulation steps is 1000 because under this horizon the bias of the clustering method can be considered marginal. This means that at leas 1000 steps for the this two route scenario is necessary for investigation the agent behaviour, because under this limit the clustering bias is too much an influence in the results.

Agent Amount

The objective in varying the amount of agents is to verify if it plays a role and if the behaviour is consistent across different agent populations. The fixed parameters, valid for all simulations, are in Tab. 8.5 with its derived properties in Tab. 8.7. The variable parameters are shown in Tab. 8.6. Because the amount of agents implies in different route properties they are presented in Tab. 8.7. In this table is important to observe that the clustering parameter ϵ varies and this is due to Eq. 8.8, since the target density is kept constant at 0.3 but the agent amount varies (as more agents as longer the length and as higher the ϵ as well).

Table 8.5: Fixed parameters for agent population experiments

Parameter	Value
Simulation horizon	1000 steps
Target density	0.3

Table 8.6: Simulation parameters for agent population experiments

Parameter	Values
Agent amount	$\{10, 50, 100, 500\}$
Reasoning engine	$\{STD, CluEUT, CluPT\}$

The results for the different agent populations are depicted in Tab. 8.8 and are consistent with the expected results: rational behaviour near the μ_{EUT} and the non-rational PT based agents having a consistent deviation from it (μ_{CluPT}^b) . The reasons for the higher occupation in the *Main* and not in the *Secondary* route are explained by the next experiment.

Target Density

The fixed parameters are depicted in Tab. 8.9 and the variable parameters in Tab. 8.10. For convenience, the derived properties for the different densities are shown in Tab. 8.11. The density experiments (Tab. 8.12) show that it does not worth investigating scenarios where the density is

Accenta	Clustering c	Longth	Equilibrium		Capacity in veh		
Agents	Clustering ϵ	Length	Travel-time	Mean speed	Main	Sec.	Total
10	3.33 <i>s</i>	55.6m	1.97s	28.25m/s(102km/h)	22	11	33
50	16.67s	277.8m	9.83s	28.25m/s(102km/h)	111	55	166
100	33.33 <i>s</i>	555.6m	19.67s	28.25m/s(102km/h)	222	111	333
500	166.67s	2777.8m	98.33s	28.25m/s(102km/h)	1111	555	1666

Table 8.7: Derived parameters from agent amount experiments with target density 0.3

Table 8.8: Occupation results of the *Main* route for agent amount experiments and target density of 0.3

Agents	$\mu_{STD}(\sigma_{STD}): err$	$\mu_{CluEUT}(\sigma_{CluEUT}): err$	$\mu_{CluPT}(\sigma_{CluPT}): err$	μ^b_{CluPT}	μ_{EUT}
10	6.65(1.03):-0.01	6.92(1.01):+0.25	8.25(1.17):+1.59	8.00	$6.\overline{6}$
50	33.33(3.68):+0.00	34.06(3.75):+0.73	38.35(4.44):+5.02	37.62	$33.\overline{3}$
100	66.12(6.99):-0.53	67.25(7.13):+0.59	76.93(8.32):+10.26	76.35	$66.\overline{6}$
500	330.96(34.28):-2.37	336.79(34.01):+3.45	379.16(38.88):+45.83	375.70	$333.\overline{3}$

not above the limit of the free-flow region ($d \ge 0.2$). The reason, first, is that in the free-flow region not enough variability is produced and the experienced travel-times are too similar.

Table 8.9: Fixed parameters for target density experiments

Parameter	Value
Agent amount	100 agents
Simulation horizon	1000 steps

It is again important to explicitly say that the clustering parameter ϵ changes from target density to target density (Tab. 8.11). This is due to the Eq. 8.8 that depends on the route length that depends on the target density and agent amount. Since here the agent amount is kept constant at 100 but the density varies it also does the length and consequently the clustering parameter.

Another limit is to avoid high densities such as 0.9 that leaves almost no room for variations, i.e. the routes are so saturated that the experienced travel-time tends to repeat itself and that is why when the density increases the deviation (in μ^b_{CluPT}) decreases (Tab. 8.12). Because the routes are already saturated (in the congestion region) little to no difference is made in the experienced travel-time. So the best density value to appreciate the differences between the behaviours is in the region between 0.2 and 0.6 (Tab. 8.12).

A last point that needs an explanation is why the overloaded route (occupation above the equilibrium) is the *Main* and not the *Secondary* route. The explanation is in the $\pi(\bullet)$ function curvature (Fig. 4.1b), which is an inverted "S". The point is that the *Secondary* route is much more sensible to any variation in its occupation than the *Main*. This means that it is prone to have extreme travel-times (congestion region) with low frequencies, which translates into low probabilities. Because of that, the agent (through the PT evaluation method) is penalising the *Secondary* while tolerating frequent "bad" outcomes in the *Main* route.

8.9 The Selten Scenario

The Selten Scenario is a reproduction of the scenario proposed in [SSC⁺05, Chm05].⁶ There, a traffic scenario with two routes was designed, similar to the one presented in previous section. For these scenario 18 individuals were asked to decide between two routes and they were monetarily

 $^{^{6}}$ This scenario is also used in [KB04] but there the impact of giving the agents information about the traffic is evaluated and therefore not quite the application here.

Parameter	Values
Target density	$\{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$
Reasoning engine	$\{STD, CluEUT, CluPT\}$

Table 8.10: Simulation parameters for target density experiments

Table 8.11: Derived parameters from target density experiments and 100 agents

Dong Clustoning c		Longth	Equilibrium		Capacity in veh		
Dens. Clustering	Clustering ϵ	Length	Travel-time	Mean speed	Main	Sec.	Total
0.1	100.10s	1666.7m	50.050s	33.3m/s(120km/h)	667	333	1000
0.2	50.05s	833.3m	25.03s	33.3m/s(120km/h)	333	167	500
0.3	33.37s	555.6m	16.68s	33.3m/s(120km/h)	222	111	333
0.4	25.03s	416.7m	16.68s	33.3m/s(120km/h)	167	83	250
0.5	20.02s	333.3m	12.51s	33.3m/s(120km/h)	133	67	200
0.6	16.68s	277.8m	10.01s	33.3m/s(120km/h)	111	56	167
0.7	14.30s	238.1m	8.34s	27.0m/s(97km/h)	95	48	143
0.8	12.51s	208.3m	8.83s	16.0m/s(57km/h)	83	42	125
0.9	11.12s	185.2m	26.28s	7.0m/s(25km/h)	74	37	111

rewarded for choosing the fastest route (shortest travel-time). The 18 individuals were competing with each other and 100 rounds of decisions were performed.

The scenario is the same as in Fig. 8.1 but the Eq. 8.9 and 8.10 were used for calculating the travel-time (in minutes) for the *Main* and *Secondary* route respectively. There t_m is the travel-time in minutes for the *Main* route, n_m is the amount of individuals that chose the *Main* route, t_s is the travel-time in minutes for the *Secondary* route, and n_s is the amount of individuals in the *Secondary* route.

$$t_m = 6 + 2n_m \tag{8.9}$$

$$s = 12 + 3n_s$$
 (8.10)

The travel-time functions are depicted in Fig. 8.2 and Fig. 8.3, where Fig. 8.2 shows them in a regular plot and in Fig. 8.3 modifying t_m to $t_m = 6 + 2(18 - n_s)$. The rational equilibrium for this scenario is when 6 individuals choose the *Secondary* and 12 the *Main* route, which yields a 30 minutes travel-time for both routes. This is the same ratio used in the previous scenario (2/3 and 1/3).

t

For this scenario four experiments were made but from them only two are of interest here. These two are the following: first the individuals were only informed to make their decisions and to try to choose the fastest route. In the second experiment they were told that the *Main* route has higher capacity than the *Secondary*. The data collected from both experiments are reproduced in Tab. 8.13, which was extracted from [SSC⁺05, Chm05], in Tab. 18 from page 68 from the rows identified by "Variation I" and "Variation II" respectively.⁷ In Tab. 8.13 μ_s and σ_s correspond to the mean amount of individuals in the *Secondary* route and the standard deviation. The column "Deviation" corresponds to the deviation from the expected rational behaviour. There the line "Not-informed" corresponds to the first experiment, where the individuals were not informed about the differences in the routes' capacities. The "Informed" line corresponds the second experiment, where the individuals were aware of the routes' capacities. The last line, " μ_{EUT} " corresponds to the expected rational equilibrium.

It is interesting to observe that the data is more than one standard deviation "under" the expected rational behaviour. Replicating the same experiment using the three engines being

⁷ The thesis of Chmura is available at http://www.ub.uni-duisburg.de/ETD-db/theses/available/ duett-05152005-222337/.



Figure 8.2: Travel-time functions



Figure 8.3: Travel-time functions by individuals in the Secondary route

Density	$\mu_{STD}(\sigma_{STD}):err$	$\mu_{CluEUT}(\sigma_{CluEUT}): err$	$\mu_{CluPT}(\sigma_{CluPT}): err$	μ^{b}_{CluPT}
0.1	66.45(7.09):-0.21	67.77(7.13):+1.10	69.19(7.24):+2.53	68.09
0.2	66.15(7.04):-0.50	67.35(7.13):+0.68	75.57(8.23):+8.90	74.89
0.3	66.12(6.99):-0.53	67.25(7.13):+0.59	76.93(8.32):+10.26	76.35
0.4	66.43(7.05):-0.23	67.33(7.11):+0.66	74.05(8.28):+7.39	73.39
0.5	66.57(7.12):-0.09	67.41(7.15):+0.74	72.99(7.93):+6.32	72.25
0.6	66.25(7.11):-0.40	67.89(7.08):+1.22	73.48(7.83):+6.81	72.25
0.7	66.18(7.12):-0.47	67.55(7.10):+0.88	73.88(8.13):+7.21	73.00
0.8	65.74(7.13) : -0.92	67.66(7.12):+0.99	74.24(7.94):+7.58	73.25
0.9	65.85(6.96) : -0.81	67.77(7.08):+1.10	73.20(7.97):+6.54	72.10
μ_{EUT}	$66.\overline{6}$			

Table 8.12: Occupation results of the Main route for density experiments and 100 agents

Table 8.13: Table 18 from [Chm05] for the secondary route and 18 persons

Type	μ_s	σ_s	Deviation
Not-informed	4,50	1.38	-1.50
Informed	4,44	1.01	-1.56
μ_{EUT}	6.00		

tested here the results, in Tab. 8.14, are quite compelling. The simulate values were obtained by fixing the horizon at 1000, repeating the experiment 30 times, and setting the clustering threshold to 5 minutes. In Tab. 8.14 the column "Error" was added to show how much the mean (μ_s) of the simulated values deviate from the data.

Table 8.14: Simulation results for the secondary route and 18 agents

Enmino			Error		
Engine	μ_s	0 s	Non-Informed (4.50)	Informed (4.44)	
STD	5.61	1.47	1.11	1.17	
CluEUT	5.16	1.30	0.66	0.72	
CluPT	4.10	1.20	-0.40	-0.34	
μ_{EUT}	6.00				

The simulated results show that the CluPT is closer to the real data but only by a small margin when looking only at the error magnitude in the column "Error" (Tab. 8.14). However, when looking at where the CluPT got wrong then it is the clear "winner". The data shows that the individuals tend to sub-utilise the *Secondary* route, i.e. avoid it even when it means to get a worse travel-time. This is exactly what the CluPT agents do, they avoid the *Secondary* route and it can be said that the CluPT agents "exaggerate" the behaviour observed in the data. That is not how STD and CluEUT behave. They tend to approach the rational user equilibrium (6.00). In this case, the CluPT is indeed the best approach for modelling this scenario.

Even though the results supports the claims in this thesis the experiments performed in [SSC⁺05, Chm05] have some issues. The first issue regards the travel-time functions, they are very unlikely to have any practical validity. Second, the participants are rewarded directly with money which is not the case in traffic, where the monetary utility is a few steps away from the transportation experience. This is shown to be a very important issue regarding the human attitude towards accomplishing tasks [MA06, Ari08]. The third criticism is that the 100 decisions were made on the same day, where in reality it is done once or twice a day. However, the data and the simulations are not being disputed here, just its validity as a reasonable traffic scenario. The point being made is that this scenario shows that the PT based non-rational behaviour suits better to reproduce the behaviour observed in the data but cannot be said to correspond to a real situation in traffic.

A note on the apparent lack of bias in the clustered reasoning engines, CluEUT and CluPT. It has not vanished but has been attenuated by the linear nature of the travel-time functions (Eq. 8.9 and 8.10).

8.10 The Burgdorf Scenario

The objective in simulating the city of Burgdorf (Switzerland) was to verify if the framework scale to a real-world scenario. The data was kindly made available by Mr. Guido Rindsfüser from Emch+Berger Holding AG.⁸ The city network is in Fig. 8.4 and some figures about it in Tab. 8.15. The figures presented in Tab. 8.15 need some explanation. The values annotated with "(Network)" are the figures of the network, resulting from converting the input graph into the super-network structure (Sec. 6.5 and A.2). Where "(Target)" is given, it corresponds only to the private vehicles modus, from which the data is available, and "(Data)" is the effective amount of data available. This means that from the total 522 edges/links (corresponding to the private modus) only 122 are available to evaluation, which corresponds to roughly 23% of the target links.

Being more detailed. The vertices, edges, and maximum occupancy values with the "(Network)" label refer to values extracted from the input data after transforming it into a super-network (Sec. 6.5 and A.2). This means that if a link allows private vehicles, public transportation, and pedestrian traffic it is represented three times (one for each modus). The vertex and edge amounts are extracted from the input graph, provided with the data, and the maximum occupancy is a derived measure. This was calculated assuming that a vehicle has a mean size of 5m and then summing up all link lengths (multiplied by the corresponding lane amount) and dividing by the assumed car length.

The values labelled with "(Target)" refer to the elements, in the input data, related to the regular vehicular traffic. This means that some links do not allow private vehicles, such as pedestrian only zones or special public transportation lines. All links that does not allow private vehicles to travel through them are excluded from the figures corresponding to "(Target)".

The last label is "(Data)" and it corresponds to the elements in the "(Target)" that have an associated occupancy data. This means that from the 522 links in the "(Target)" only 122 were annotated with occupancy data, for the other 400 links no data was collected. The "Occupancy Sum(Data)" was calculated by summing up all occupancy data available and "Mean Occupancy(Data)" was calculated by dividing up the "Occupancy Sum(Data)" by the 122 links.

Property	Amount
OD pairs	1339
Agents	7357
Vertices(Network)	756
Vertices(Target)	211
Edges(Network)	2570
Edges(Target)	522
Edges(Data)	122
Maximum Occupancy(Network)	5687.93
Maximum Occupancy(Target)	1949.47
Occupancy Sum(Data)	462.03
Mean Occupancy(Data)	3.78

Table 8	3.15:	Burgdorf	scenario	figures
---------	-------	----------	----------	---------

Moreover, the load data shows that only 3 of the 122 links^9 (line "Edges(Data)" in Tab. 8.15) have a density higher than 0.3 and the conditions to observe any deviation are to have high

⁸ http://www.emchberger.ch/

⁹ To calculate this the density of all links that have occupancy data was calculated: $density = (lanes \times length)/size_{car})$, where $size_{car} = 5m$.



Figure 8.4: Burgdorf scenario

variability in the travel-time, as demonstrated by the experiments in Sec. 8.8. This means that situations near complete congestion (very high densities) and free-flow (very low densities) do not present any of the conditions where the PT based reasoning deviates from the rational behaviour. Even if the travel-time function is changed, so that the maximum flow occurs at the density of 0.2 for example, does only make a difference at those 3 links, which represent about 2.5% of data (from the 121 links) and about 0.6% of the total amount (from the 522 links). This means that the difference in the final evaluation will be barely noticed, if it can be at all. It is important to stress that this not a manipulation made here but an analysis made over the given data.

Another issue is the OD matrix. It isn't quite clear what the numbers, associated with each OD pair, represent because summing up all available space in the network it yields about 5688 vehicles at the same time (line "Maximum Occupancy(Network)" in Tab. 8.15).¹⁰ But that is the capacity of the network (assuming that it is acceptable that cars touch each others, bump to bump). If the actual counts are summed up, then it is reduced to 462 vehicles or to $462/122 \times 522 \sim 1978$, extrapolating for the other links. This means that the 7357 agents being simulated cannot possibly represent a particular moment in time, such as a rush-hour. A strategy would be to multiply, at each OD pair, the amount of vehicles/agents by a factor, such as $1978/7357 \sim 0.27$, but this would make several OD pairs disappear (some OD pairs have only 1 vehicle/agent assigned). To avoid any manipulation, that could compromise the validity of the experiments, the data was used as it was provided and for the travel-time calculations the BPR formula (Sec. 6.4.1) $bpr(\bullet)$ was used (Eq. 6.3). As expected, all links are over-saturated (densities higher than 1.0) and therefore showing no noticeable difference among the reasoning engines. The reason is simple, at oversaturation the variability is low, i.e. every agent tend to experience the same travel-time over and over again. This means that little to none stochastic variation is present in the prospects/clusters, leading to a decision that is equivalent to the rational decision. Nevertheless the results in Tab. 8.16 are presented to show that such a scenario can be simulated and that the whole framework does scale to real-world scenarios. There each simulation has an horizon of 1000 steps¹¹ and each simulation was repeated 20 times. The column "MSE" represents the final Mean Squared Error (MSE) values, " \sqrt{MSE} " the absolute error (not squared), "Expected (\sqrt{MSE})" is what the data yields as expected for the previous column, "Error" is difference between the expected and the actual values (the simulated minus the expected value), and "Error factor" represents how many times the expected value have been missed (the division between the simulated and expected values).

Table 8.16:	Burgd	orf s	scenario	MSE	resul	.ts
-------------	-------	-------	----------	-----	-------	-----

Engine	MSE	\sqrt{MSE}	\overline{occup}	Error factor	Time (h)	Mem. (MB)
STD	683.70	26.15	3.78	6.92	15.40	309.96
CluEUT	662.58	25.74	3.78	6.81	19.90	608.17
CluPT	694.34	26.35	3.78	6.97	17.07	498.51

Because of the several issues mentioned before the results can only be used to say that a complete city can be simulated using the framework presented here. Any other claim will be bogus. It is tempting to say that the CluPT is worse than the others, but this is deceiving because the "MSE" column shows the mean squared error. When looking at the absolute mean error, in column \sqrt{MSE} from Tab. 8.16, this difference drops but says that all engines are failing by a "catastrophic" margin (column "Error factor" in Tab. 8.16). The value in CluPT (26.35) says that it missed, on average, by 22.56 vehicles in each of the 122 links, whose data is available. Then taking the average occupancy in the 122 links it gives about 3.78 vehicles per link (line "Mean Occupancy(Data)" in Tab. 8.15 and "occup" in Tab. 8.16). This means that missing these 3.78 vehicles by 25.74, for the CluEUT, which is 6.81 times more vehicles (column "Error factor"

¹⁰ The calculation done was: $total = 1/car_{length} \sum_{l \in L} l_{length} l_{lanes}$, where $car_{length} = 5m$, $l \in L$ is a link in the link set (each of the 522 links), l_{length} is the link length in meters, and l_{lanes} is the amount of lanes for this link.

¹¹ This value is not a definitive value, it must be investigated but because the objective was to verify only the scalability of the framework it was not further investigated.



Figure 8.5: Average occupancy evolution

in Tab. 8.16), is almost as bad as missing it by 26.35 (for CluPT), which is 6.97 times higher.

A note on performance (the last two columns in Tab. 8.16). This the experiments were run on an Intel[®] Pentium[®] 4 3.20 GHz with 2GB Ram. The discrepancy in *CluEUT* is because the simulation was run on another machine with the same processor but with a concurrent simulation (the processor was not dedicated for the simulations). The second difference is that the machine where *CluEUT* was run has an inferior I/O performance. The I/O performance is important because data was constantly be collected and log files from the simulations were produced (to later analysis for possible errors).

8.11 Conclusion

From the results presented with different parameters it can be seen how the different parameters are relevant when considering the use of the PT for agent modelling. The first conclusion is that the amount of agents does not have a relevant influence in the agent behaviour. Another point is the influence of the density, or load level. The higher the load the more relevant the consideration of the PT is, excluding the extreme situations where it reaches over-saturation. Another relevant aspect is the shape of function $\pi(\bullet)$ that determines, in this particular case, which route will be stressed.¹²

Some practical aspects are also relevant and they are the simulation horizon and the repetition amount. The horizon must be at least as high as the minimum required, in this case under 1000 steps, to guarantee that the implementation artefacts do not greatly influence the results, for the two route scenario. It remains to be validated for more complex scenarios. As the results show the most severe artefact is the bias in the clustering method. For the repetition amount, which were

¹² For a visual impression of function $\pi(\bullet)$ with different γ values see Fig. B.1.



Figure 8.6: Mean occupancy and standard deviation for CluPT engine

fixed in 100, it is less critical. To visually observe the influence of the repetition amount please address to the Fig. 8.5 and 8.6. Where the average occupancy for the *Main* route using different reasoning engines is plotted (Fig. 8.5). Then the same plot but for only the CluPT including its standard deviation (Fig. 8.6). For these plots the horizon is 1000, target density 0.3, and agent amount 100.

Besides analysing the conditions under which the PT based agents do deviates from the rational behaviour, it is also correct to say that they behave more closely to the real data than their rational counterparts. As seen in Sec. 8.9, the CluPT agents are closer to the real data than the rational based agents. The non-rational agents not only have a lower error but also show the same behaviour, avoidance of the *Second* route, which supports the claims that non-rationality is a better approach for the human decision-making modelling. A last point is that the framework does scale up to allow simulations of real-world scenarios, as shown by the last section.

Summarising what the results show the following claims are made. First the conditions necessary for the experiments were established with the different experiments in Sec. 8.8. Then a similar scenario was used to show that the non-rational behaviour (the CluPT engine) is indeed better at reproducing real data (Sec. 8.9). The last claim is that the framework does scale for real world scenarios and that is demonstrated in Sec. 8.10.

Chapter 9

State-Of-The-Art And Related Work

As delineated by the previous chapters this thesis has a multidisciplinary scope and it can be broken into three sub-themes. The first is the utility based discrete choice modelling (covered by chapters 2, 3, and 4), the second is agent learning and architecture (in chapters 5 and 7), and the third is the traffic assignment problem (chapter 6). In this chapter these sub-fields are again approached, presenting the advancements in each area and how they relate to this thesis.

9.1 Discrete Choice Modelling

In the previous chapters only utility based choice modelling theories were presented but other theories exist that are not based on utility functions. Consequently, first the utility based theories are reviewed and then the non-utility based alternative is presented.

9.1.1 Utility Based Modelling

The models based on the utility theory [Fis70] (UT) can be subdivided between rational and non-rational. By rationality it is meant the models that are based on the work of von Neumann and Morgenstern, as presented in chapter 3. The non-rational models, on the other hand, are the models that do not follow this definition, for instance the Prospect Theory [KT79] (PT) presented in chapter 4.

Rationality Based Utility Models

The first modelling technique under this definition is the Expected Utility Theory [vN28, vNM07] (EUT), which is the base for all others. As discussed chapter 3, this theory has some drawbacks such as requiring complete knowledge about the individual's internal decision process. To cope with such restrictions, improvements known as Random Utility Models [GP06] (RUM) were proposed. The main development of RUM is to allow partial knowledge about the decision process of the individuals being modelled, because under the RUM it is admitted that some factors cannot be modelled. But these factors are still relevant because they can disrupt the ranking established by the utility calculated based on the known factors, as mentioned in Sec. 2.5.2 and 3.4.

The models mentioned and briefly presented in Sec. 3.4 are improvements over the EUT. They are also more advanced than the model presented in this thesis because they, except for Logit [Ber44, Luc59] and Multi Nominal Logit [McF74] (MNL), assume that the options need some treatment to make them iid (independent from each other). This correlation capturing is not present in this work. Here, as in the MNL, no structure is given to extract the correlation between the options. This is a point that is left as future work because it requires a complete new

study about how agents may perceive correlation and how they could learn and extract them from the options.

The complexity and refinement in the most advanced models such as the Path-Size Logit [BAB99] or Probit [Bli34a, Bli34b] are beyond the scope of this thesis, but "better" models.¹ In a recent improvement [FB09] it is proposed that instead of using additive errors the use of multiplicative errors is evaluated. This means that instead of a utility as in $u = v + \varepsilon$ (Eq.2.5) it multiplies the error: $u = v\varepsilon$. They also report that the multiplicative model represents the data better than the additive model (in this case an MNL model). This does not mean that their improvements may not be used in the non-rational model presented here, but that they must be carefully studied and brought to light in relation to learning algorithms.

Non-Rational Utility Models

The most prominent non-rational utility based model is the Prospect Theory [KT79] (PT) and its improvements, as mentioned in chapter 4 (specially Sec. 4.7). The difference between PT and EUT is in how the utility is calculated and which assumptions are made about how u is aggregated: compare $eut(\bullet)$ (Eq. 3.3) and $pt(\bullet)$ (Eq. 4.1). Because the PT does not observe all the axioms from EUT the improvements mentioned in previous section cannot be directly incorporated into the PT, even though both theories are similar. In this regard the PT still lacks development to capture the correlation from the options to allow a broader use.² Among the improvements of the PT is the Cumulative Prospect Theory [TK92, WT93] (CPT), presented in Sec. 4.7.1. The CPT can be easily incorporated into the presented agent learning algorithm (chapter 5) and architecture (chapter 7) but no experiment was made using it because it is not clear if it is applicable for all modelling problems, as already discussed. The use of CPT into the Q-Learning is rather simple: instead of using $pt(\bullet)$ in modified Q-Learning (Sec. 5.6) the $cpt(\bullet)$ function (Eq.4.4) is used.

The PT, however, is not the only utility based non-rational theory. The Theory of Small Feedbacks [BE03] is another alternative. The problem with it is that it has some problems with large horizons decisions.³ Because of this issue no further attempts to use this theory were made.

9.1.2 Non-Utility Based Modelling

For the non-utility based modelling an example is the Fast and Frugal Way [GG96] with its extensions [GG02, HG05]. This model is based on a hierarchical decision process, where the decision problem is as follows. The decision problems are always binary, i.e. compare two options and choose the best. For these choice problems a set of binary criteria is given, where each of them can assume the values y (for yes) or n (for no), meaning that the criterion is fulfilled (case y is assigned) or not (case n is assigned). It is assumed that the decision problem can be modelled by establishing a sequence of such binary criteria, with their corresponding y and n values. This means that in a hypothetical choice problem $\mathbf{X} = \{A, B\}$ five criteria are given, say $\langle V, W, X, Y, Z \rangle$. Each of these attributes/criteria is then applied to each option (A, and B). From this evaluation a string of y and n (for yes and no) is returned by the model such as: $s_A = \langle y, n, n, y, y \rangle$ and $s_B = \langle y, n, y, n, n \rangle$. This means that: the criterion V is met by both options (the first y on each string); then the criterion X fulfilled but A has not (the third character in the sequences). In this case, according to the theory, the option B is the choice made.

The algorithm works by choosing the "winner" by comparing pairwise the two options' strings of *yes* and *no*. If both options have the same answer for the same criterion the algorithm compares the next criterion and does that until the options have different answers, winning the option with

 $^{^{1}}$ It is said "better" because these models refine the structure of the option set more carefully than it is done in this thesis.

² The PT requires the same assumption as in EUT, i.e. all options are independent from each other.

 $^{^{3}}$ The model has been implemented and tested but it "degenerates" when the iteration horizon is extrapolated beyond of what is used in the original article [BE03]. The authors of [BE03] were notified [Bar07] about the problem, including the evidences, but no answer has been received.

an y. The problem with this theory is that it is only possible to have binary choices (as the Logit model) and only binary criteria are allowed (no numerical comparison is considered).

9.2 Traffic Assignment

As explained in chapter 6, the traffic assignment step is the most complex in the four-step modelling (FSM) concept. Considering the context of this work it is necessary to mention the traditional approach and its state-of-the art as well as the work in the artificial intelligence (AI) field. Under what is here called the traditional approach it is meant the microeconomics approach, i.e. the use of microeconomic methods to model the fourth step in the FSM. However, these techniques are macroscopic. In the AI field the focus was on the multi-agent related work, to which this thesis is more closely related.

9.2.1 Microeconomics

The developments in the microeconomics discrete choice modelling are close related to the use of microeconomics for traffic modelling. When the MNL [McF74] model was developed, the objective was to model modus choice (the third step in the FSM) for shopping behaviour (mainly to model the split between transit and private vehicle choice). Later the Nested-Logit [BA73] model was developed. This model was a significant improvement for both econometrics and traffic modelling and is the central point of the book [BAL85], considered one of the milestones in traffic assignment modelling using microeconomics. Other subsequent models as Cross-Nested-Logit [Vov97] and Path-Size Logit [BAB99, BF05] were also developed for modelling traffic assignment.

The common points in them are that they are all macroscopic and stochastic models, i.e. the route distribution is given as a result for a given β vector (Sec. 2.5). Recalling the chapter 6, such models return how many individuals, in each origin destination pair (OD pair), took each of the available routes (connecting O to D). The differences are, as explained in Sec. 9.1.1, in how the correlations among the routes are explicitly represented. The advantage in using microeconomic models is that the influence of each parameter (looking at the β^* , Sec. 2.5) is explicit and can give answers to questions like: "how relevant is the price of a bus ticket when compared with the travel-time?"⁴ Those models can then be used in testing hypothetical scenarios, where the bus ticket price is increased or decreased and the impact of this change on link load can be evaluated. The main disadvantage of these models is that all of them assume the rational behaviour (EUT) as a model for human behaviour.

Some models for non-rational traffic assignment models were also proposed but for simple and synthetic scenarios, with the exception of [SK04]. There the PT was used to model the departure time choice for the commuter scenario for the Otsu city in Japan. This is not exactly traffic assignment and can be said to be more adequately classified as part of a day-activity planning strategy. In [Avi06] on the other hand, the traffic assignment problem for a synthetic binary route choice was modelled using CPT and compared with EUT. The objective was to show that the equilibria are not compatible, i.e. CPT does not have the same equilibrium point as EUT does and therefore it is worth investigation (the use of CPT for traffic assignment).

The most closely related work to this thesis is the work done by Connors and Sumalee [CS09] where CPT is used to model traffic assignment. Their work assumes a variation on the interpretation of the Wardrop [War52] equilibrium, as explained in chapter 6. The main differences between their approach and the one here is that there the travel-time calculation is similar to the one used here, as already mentioned, and a stochastic component is attached to it. This stochastic element is a normally distributed "uncertainty". Another great difference is that it is a entirely macroscopic approach and its utility function has some drawbacks: it can return values that oscillates

⁴ Assuming that the price value used in the utility function is normalized with the travel-time value (both are on the same range), then comparing the β_{price} with $\beta_{travel-time}$ gives the relevance directly.

about the origin, i.e. it may mix positive and negative utilities, whose behaviour is not very clear.⁵

The decision to adopt CPT in [CS09] seems to be entirely based on the use of a continuous distribution function to describe the prospects and it appears easier to derive the equations for the CPT than for the PT. Since in [CS09] no real-world scenario is modelled the assumption of CPT model does seem to be based on evidences. This is the same situation in [Avi06], which also uses CPT, and does not justify its adoption by using evidences as well. Therefore it is fair to say that the choice between CPT and PT remains to be investigated when modelling route choice.

9.2.2 Artificial Intelligence

The use of heuristics and specially learning for traffic assignment is an active research field. In [AP05] individuals were invited to participate in a binary route choice problem. For the hypothetical scenario the individuals suppose to choose the route that perform better. The travel-time functions were specified by a Probability Distribution Function (PDF) for each route, so that the travel-time experienced at each decision "round" varied. After collecting the decision data, several models where built to check which of them better reproduces the data – the best on fitting the data of the real decision-makers. The compared family of models were EUT, RUM, CPT, Bayesian Learning, and Reinforcement Learning, being the latest the best approach. However, the experiments can be said to be biased because in some of the tested scenarios the initial assumptions were changed along the experiment (the travel-time PDFs changed during the experiment), which could be the reason why learning performed better than the other approaches.

In $[BBA^+00]$ a synthetic commuter scenario⁶ with two routes was designed to verify the performance of different decision-making strategies, being these strategies based on the past two experiences. The objective is always to choose the route that has the lowest amount of agents (a minority game). The commuter scenario is also the one used in [KB02, BK03] where the optimal distribution is to have 2/3 of the agents on the main route and the remaining 1/3 on the secondary route. In the experiments performed two types of reinforcement based heuristics were used: the first based on own experience and the other taking into account a forecast given by an external source. The focus, on the first experiment, was to evaluate the best learning frequency, i.e. how often should the agent re-evaluate the route scores (using the reinforcement heuristic). On the second experiment the agent's task was extended to also evaluate the quality of the forecast, i.e. how reliable it was. The findings show that, first, a rather frequent re-evaluation is better than a seldom and, second, that the forecast reliability is inversely correlated with the amount of agents receiving this information. Then, again exploring the commuter scenario, in [KBW03] a different set-up was used, where both routes had the same capacity, i.e. the agents had to choose between physically equivalent routes. The difference between [KBW03] and [BK03] is that the agents willing to use a forecast information can shop among the different types of information. For each agent four types of information where made available, from a very simple binary information about the presence or not of congestion on the routes to the very elaborated traffic density for each route. The agents must then pay for the information with different prices according to the complexity of the information. The objective was to evaluate which would be the most efficient type of information for the agents, i.e. the type of information with the best cost/benefit rate. The surprising finding was that mean speed turned out to be a misleading information, where the agents using it performed worse than the ones using no information.

In [SSC⁺05, Chm05] a scenario similar to [BK03, KB04] were used to evaluate how good a reinforcement learning reproduces the data produced by the following set-up. Eighteen individuals participated in the experiment to choose the fastest route, given that one route is "naturally" faster than the other. The experiment had an horizon of 200 choices and different scenarios were proposed. First no information about the routes' nature was given, i.e. the individuals were not

 $^{^{5}}$ The problem is not that no mathematical treatment exist but that it depends on empirical evidences to support its application, which is the argument in [LF91, Sch03].

 $^{^{6}}$ Commuter scenario means that for a given familiar OD pair (usually from home to work and vice-versa) the individuals are experienced with the available options and must decide which route to take. The choice is normally to minimise the travel-time.

explicitly informed that one route was faster than the other. Then in a second experiment this information was given. The third was similar to the first experiment but a travel-time disturbance is introduced, simulating the situation where a construction site appears interchangeably in one of the routes. The last is similar to the third but with the additional information about which route is usually the faster. It is important to say that the individuals where monetarily rewarded for their correct decisions. The real data for the experiment without disturbance is the most interesting and closely related to this work. There the distribution is not the one of the rational equilibrium. The real data yields mean occupation, in the secondary route, of 4.50^7 (no information about the different route times) and 4.44 (full disclosure about the nature of each route) and at the rational equilibrium 6.00 is expected. Even though the results in [SSC⁺05, Chm05] are similar to the results found here in chapter 8 (for the two-route scenario) a critic must be made. It cannot be assumed that persons behave in traffic as in the computer simulations. The first reason is the time between the experiences: in traffic it is usually once a day and in the computer are 200 simulations a day. Second, in the computer simulations the participants are monetarily rewarded for their decisions, which is not the case in traffic.⁸

In another work [BK05] the Braess [Bra68, BNW05] paradox was investigated and how effective a control system would be in avoiding its emergence. To test the hypothesis, a Braess scenario was simulated where two types of drivers where present: informed and non-informed drivers. By informed it is meant that the informed drivers received traffic information from the control system about the network state. The results were that when drivers receive the information about the network state, the network state gets closer to the global optimum, but even closer if the received information is manipulated (inducing the drivers to take different routes). This work only marginally related to one presented here because no Braess condition is investigated here, but remains as future work.

The influence of a control system is also the subject in [Wah02], where a two-route scenario, along with other scenarios, is used for investigating the use of Advanced Traveller Information Systems (ATIS). There the impact of the use of ATIS is tested using Multi-Agent Simulations (MASim) but the findings are strictly hypothetical. The objective was to verify how can information about the traffic state influence the drivers behaviour. One remark, which is also present in [KBW03], is that when the amount of informed drivers (that use ATIS) grows the system performance degrades. The main reason is because when ATIS gives a forecast about future conditions it does not take into account the reaction of the drivers to this very given forecast. In this case what happens is that the more users receive a forecast, and decide accordingly, the lower the reliability is for the forecast. The relation of the work in [Wah02] and the one here resides in the use of the two-route scenario and the use of MASim.

Yet all these approaches are still based on the fundamentals of game-theory [vN28, vNM07] (the rational behaviour). One exception is the Small-Feedbacks [BE03] model which is not based on rationality, but showed some problems (as discussed early, in Sec. 9.1.1).

⁷ This data corresponds to the mean occupation for the secondary route in [Chm05], Tab. 18 at page 68 in the rows "Variation I" and "Variation II". The thesis of Chmura is available at http://www.ub.uni-duisburg.de/ ETD-db/theses/available/duett-05152005-222337/.

⁸ In traffic the monetary reward is at least one step away (considering gasoline consumption) or several (considering vehicle depreciation, maintenance, insurance, etc). It is known that human behaviour changes when money is not directly handled (see [MA06, Ari08] for further information) and as further way the money reward/cost from the action is as less aware people are of it.

9.2. TRAFFIC ASSIGNMENT

Chapter 10 Conclusion And Future Work

This thesis proposes a novel approach for modelling human decision-makers using Multi-Agent Simulations (MASim) with a non-rational behaviour. This non-rational behaviour is here based in the Prospect Theory [KT79] (PT), chapter 4, which was compared to the rational behaviour in the Expected Utility Theory [vNM07] (EUT), chapter 3. This model was used to design a modified Q-Learning algorithm (chapter 5). The PT based Q-Learning was then integrated into an agent architecture (chapter 7).

The designed architecture with the different reasoning engines (Q-Learning variations) where then used to simulate traffic assignment (chapter 6). The results (chapter 8) show that, as expected, the non-rational behaviour manifests itself only under some conditions. Theoretically these conditions are the existence of variability in the outcomes received by the agent (Sec. 4.5.2), which builds an outcome probability distribution called lottery/prospect. The variability makes it possible to have outcomes located at one or both "bumps" of the function $\pi(\bullet)$ (Eq. 4.3 and Fig. 4.1b). Some experiments, organised in sets, were performed to evaluate the proposed agent reasoning and architecture. The first set (Sec. 8.8) was designed to show which are the theoretical conditions for the occurrence of divergences between rational and non-rational behaviour in a practical scenario, in this case traffic. It also aimed at showing the possible problems with the proposed agent architecture and which are simulation parameters necessary to minimise the influence of the bias imposed by the simulation framework. From this first set it was shown that the proposed agent architecture does have some limitations (requiring a minimal amount of agents and a minimum value for simulation horizon). But it also shows that once these limitations are circumvent, by a high enough agent amount and high enough simulation horizon, the simulation framework works as expected. In the specific case of traffic the main conclusion is that the PT based reasoning deviates from the rational behaviour only when the agents experience congestions. This means that in scenarios where no traffic congestion is present no difference is expected between a rational (EUT based) and a non-rational (PT based) decision-making behaviour. On the other hand, it is also clear that only when congestions appear it is worth investigating the traffic conditions because if no congestion occurs no intervention is necessary.

The first experiments show that the congestion tends to be concentrated where most of the agents are, i.e. if a route has a higher capacity than another it also tends to be used more intensively. Consequently, "good" routes tend to be considered too "good" and to be overused, even though rationally seen it is a "bad" judgement. But this is only speculative and hypothetical. Therefore a second set of experiments was performed (Sec. 8.9) with a similar scenario but that provided real data [SSC+05, Chm05]. The data shows exactly the same tendency, i.e. to overestimate how good a "good" route is. This behaviour is the same as the data, which is the same for the non-rational agents. Not only the PT based Q-Learning shows the same tendency as it fits a slightly better the data than its rational counterparts (Tab. 8.14).

Finally, the third experiment set (Sec. 8.10) shows that the proposed agent reasoning, inserted in the agent architecture, can be used for simulating real-world scenarios. There, the city of Burgdorf in Switzerland was simulated, whose data was kindly provided by Guido Rindsfüser from Emch+Berger.¹ Unfortunately, the data provided does not seem to be suitable for the simulated scenario, as extensive discussed in Sec. 8.10. Nonetheless, it serves the purpose of showing the scalability of the proposed architecture.

10.1 Future Work

This work is just the beginning. Several aspects of the non-rational decision-making are left for investigation and some of them, the ones considered more important, are discussed here.

The first and most important is to improve the clustering method so that less simulation runs are necessary to overcome its bias. A second point is to include an explicit structure for expressing the correlation among the options in the choice set This is important because usually options are correlated (specially in traffic) and this correlation must be modelled by the agent because the PT as the EUT require the options to be iid (Sec. 3.4). To tackle this point the Support Theory [TK94, RT97, Nar04] seems to be the most advanced model. But it may be too complex for a realistic use in MASim and therefore an approach similar of to the one used in microeconomics (Sec. 3.4) seems reasonable.

Another deficiency in the PT is that it is calibrated for monetary outcomes and its use with other types of outcomes remains unknown. In traffic, it means to collect real route decision data, without the monetary reward used in [SSC⁺05, Chm05], and then to calibrate the PT parameters, the α , β , λ , and γ values (Eq. 4.1). Again in traffic, the correct use of the *status quo* must be investigated. In other words, how people set their internal *status quo* value: is it the absolute 0, the last experienced travel-time, or the estimated travel-time. Moreover, is it route dependant or is it a common value for all routes going from A to B. Those questions can only be answered investigating the attitudes of individuals fulfilling their transportation needs.

Advancing one more step into the traffic modelling, how individuals make decision about the transportation modus: why to use the car or bus or bicycle or even deciding to walk. The structure supporting this was already done here using the super-network structure (Sec. A.2), which includes the necessary algorithms for navigating in such structure (Sec. A.4). But the decision model is not yet available.

In another direction, how can this agent architecture be adapted for other decision problems, such as stock markets and consumption behaviour. If it can, is this a better model than the rational model, does the data express deviations from the rational choice that corresponds to the PT behaviour. Moreover, are other psychological characteristics necessary for a better modelling. As explained in [Kah02], some of them are: framing, anchoring, accessibility, and attribute substitution. Framing corresponds to the "wording" of the problem presentation, anchoring is the presence of an external reference point for the choice. By accessibility it is meant how easy the solution can be recovered from memory and attribute substitution is a consequence of the accessibility, where instead of using the necessary attributes to evaluate an option some correlated attribute is used that is more easily accessible.

As one can see, much research is still ahead and this work is just scratching the surface. Hopefully it will bring awareness to the use of non-rational models instead of the rational ones, for modelling human decision-makers, as well as to show that the non-rational models are indeed better in modelling persons than the rational models.

¹ http://www.emchberger.ch/

Bibliography

- [ACaERSMM93] José Augusto Azevedo, Maria Emília O. Santos Costa, Joaquim Jo ao E. R. Silvestre Madeira, and Ernesto Q. Vieira Martins. An algorithm for the ranking of shortest paths. European Journal of Operational Research, 69(1):97–106, August 1993.
- [AF00] I. Ajzen and M. Fishbein. Attitudes and the attitude-behavior relation: Reasoned and automatic processes. *European review of social psychology*, 11:1–33, 2000.
- [AH79] M. Allais and G. M. Hagen. Expected Utility Hypotheses and the Allais Paradox: Contemporary Discussions of the Decisions Under Uncertainty with Allais' Rejoinder, volume 21 of Theory and Decision Library. Kluwer Academic, September 1979.
- [AH02] J. Aycock and R.N. Horspool. Pratical earley parsing. *The Computer Journal*, 45(6):620–630, 2002.
- [All53] M. Allais. Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'École américaine. *Econometrica*, 21(4):503–546, October 1953. Translated and reprinted in [AH79].
- [ALSS95] Rakesh Agrawal, King-Ip Lin, Harpreet S. Sawhney, and Kyuseok Shim. Fast similarity search in the presence of noise, scaling, and translation in time-series databases. In VLDB '95: Proceedings of the 21th International Conference on Very Large Data Bases, pages 490–501, San Francisco, CA, USA, 1995. Morgan Kaufmann Publishers Inc.
- [ALSU06] Alfred V. Aho, Monica S. Lam, Ravi Sethi, and Jeffrey D. Ullman. *Compilers: Principles, Techniques, and Tools.* Addison Wesley, 2nd edition, August 2006.
- [AOR05] Dan Ariely, Axel Ockenfels, and Alvin Roth. An experimental analysis of ending rules in internet auctions. *RAND Journal of Economics*, 36(4):890–907, Winter 2005.
- [AP05] Erel Avineri and Joseph N. Prashker. Sensitivity to travel time variability: Travelers' learning perspective. *Transportation Research Part C: Emerging Technologies*, 13(2):157–183, April 2005.
- [Ari08] Dan Ariely. Predictably Irrational: The Hidden Forces That Shape Our Decisions. Harper Collins Publishers, 77–85 Fulham Palace Road, Hammersmith, London W6 8JB, March 2008.
- [Art94] W. Brian Arthur. Inductive reasoning and bounded rationality. The American Economic Review, 84(2):406–411, May 1994. Papers and Proceedings of the Hundred and Sixth Annual Meeting of the American Economic Association.

- [Avi06] Erel Avineri. The effect of reference point on stochastic network equilibrium. *Transportation Science*, 40(4):409–420, 2006.
- [BA73] M. Ben-Akiva. The structure of travel demand models. PhD thesis, MIT, 1973.
- [BAB99] M. Ben-Akiva and M. Bierlaire. Discrete choice methods and their applications to short-term travel decisions. In Randolph W. Hall, editor, *Handbook* of Transportation Science, Kluwer's International Series Operations Research Management Science, book 2, pages 5–34. Kluwer, 1999.
- [BAL85] Moshe Ben-Akiva and S. Lerman. Discrete Choice Analysis Theory and Application to Travel Demand. MIT Press, Cambridge, MA., 1985.
- [BAN06] M. Balmer, K. W. Axhausen, and K. Nagel. A demand generation framework for large scale micro simulations. In 85th Annual Meeting of the Transportation Research Board, TRB, Washington, D.C., USA, January 2006.
- [Bar07] Greg Barron. personal communication, September 12th 2007.
- [Bat07] John Bates. History of demand modelling. In David A Hensher and Kenneth J Button, editors, *Handbook of Transport Modelling*, volume 1, chapter 2, pages 11–34. Emerald Group Publishing Limited, 2nd edition, October 2007.
- [Bay63] Thomas Bayes. An essay towards solving a problem in the doctrine of chances. by the late rev. mr. bayes, f. r. s. communicated by mr. price, in a letter to john canton, a. m. f. r. s. *Philosophical Transactions of the Royal Society of London*, 53:370–418, 1763.
- [BBA⁺00] Ana L.C. Bazzan, Rafael H. Bordini, Gustavo K. Andriotti, Rosa Vicari, and Joachim Wahle. Wayward agents in a commuting scenario (personalities in the minority game). In Edmund Durfee, editor, Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS'2000),10–12 July, Boston, pages 55–62, Los Alamitos, CA, 2000. IEEE Computer Society.
- [BBAR06] Shlomo Bekhor, Moshe E. Ben-Akiva, and M. Scott Ramming. Evaluation of choice set generation algorithms for route choice models. *Annals of Operations Research*, 144(1):235–247, May 2006.
- [BE03] Greg Barron and Ido Erev. Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, 16(3):215–233, July 2003.
- [Bel57a] Richard Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, dover paperback (2003) edition, 1957.
- [Bel57b] Richard Bellman. A markovian decision process. Journal of Mathematics and Mechanics, 6, 1957.
- [Ber38] Daniel Bernoulli. Specimen theoriae novae de mensura sortis. Commentarii Academiae Scientiarum Imperiales Petropolitanae, 5:175–192, 1738. Translated and reprinted in [Ber54].
- [Ber44] Joseph Berkson. Application to the logistic function to bio-assay. Journal of the American Statistical Association, 39(227):357–365, September 1944.
- [Ber54] Daniel Bernoulli. Exposition of a new theory on the measurement of risk. *Econometrica*, 22(1):23–36, January 1954.
- [Ber93] James O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Springer, 2nd edition, March 1993.

[BF05]	Michel Bierlaire and Emma Frejinger. Route choice models with subpath components. In <i>Proceedings of Swiss Transport Research Conference (STRC) 2005</i> , Monte Verità, Ascona, Switzerland, March 2005.
[BHL05]	Piet H. L. Bovy and Sascha Hoogendoorn-Lanser. Modelling route choice behaviour in multi-modal transport networks. <i>Transportation</i> , 32(4):341–368, July 2005.
[Bie03]	Michel Bierlaire. BIOGEME: a free package for the estimation of discrete choice models. In <i>Proceedings of the 3rd Swiss Transport Research Conference</i> , Monte Verita, Ascona, Switzerland, 2003.
[Bie05]	Herman J. Bierens. Introduction to the Mathematical and Statistical Foundations of Econometrics. Cambridge University Press, digital edition, January 2005.
[Bie08]	Michel Bierlaire. An introduction to BIOGEME Version 1.7. EPFL, digital edition, August 2008. Version 1.7, last visited at 27.11.2008.
[BJM98]	Christopher L. Barrett, Riko Jacob, and Madhav V. Marathe. Formal language constrained path problems. In S. Arnborg and L. Ivansson, editors, <i>Algorithm Theory SWAT'98</i> , volume LNCS 1432, pages 234–245. Springer Berlin / Heidelberg, 1998.
[BK03]	A. L. C. Bazzan and F. Klügl. Route decision behaviour in a commuting scenario: Simple heuristics adaptation and effect of traffic forecast. In <i>Proceedings of the</i> <i>Euroworkshop on Behavioural Responses to ITS</i> , Eindhoven, 2003.
[BK05]	Ana L.C. Bazzan and Franziska Klügl. Case studies on the braess paradox: Simulating route recommendation and learning in abstract and microscopic models. <i>Transportation Research Part C: Emerging Technologies</i> , 13(4):299–319, August 2005.
[Bli34a]	C. I. Bliss. The method of probits. <i>Science</i> , 79(2037):28–39, January 1934. Corrected in [?].
[Bli34b]	C. I. Bliss. The method of probits – a correction. Science, $79(2037)$:409–410, January 1934.
[BMR ⁺ 08]	M. Balmer, K. Meister, M. Rieser, K. Nagel, and K.W. Axhausen. Agent- based simulation of travel demand: Structure and computational performance of matsim-t. In 2nd TRB Conference on Innovations in Travel Modeling, Portlan, USA, June 2008.
[BNW05]	Dietrich Braess, Anna Nagurney, and Tina Wakolbinger. On a paradox of traffic planning. <i>Transportation Science</i> , 39(4):446–450, November 2005. joint translation of "Über ein Paradoxon der Verkehrsplanung" [Bra68] with Anna Nagurney, Tina Wakolbinger.
[Bra68]	Dietrich Braess. Über ein paradoxon der verkehrsplanung. Un- ternehmensforschung, 12:258–268, March 1968.
[BRV ⁺ 05]	M. Balmer, M. Rieser, A. Vogel, K.W. Axhausen, and K. Nagel. Generating day plans based on origin-destination matrices. a comparison between visum and matsim based on kanton zurich data. In <i>Proceedings of Swiss Transport Research Conference (STRC) 2005</i> , Monte Verità, Ascona, Switzerland, March 2005.
[BS90]	Piet H.L. Bovy and Eliahu Stern. <i>Route Choice: Wayfinding in Transport Networks</i> . Kluwer Academic Publishers, August 1990.

[CFCLB03]	Kristof Carlier, Stella Fiorenzo-Catalano, Charles Lindveld, and Piet Bovy. A supernetwork approach towards multimodal travel modeling. In <i>In Proceedings of the 82nd TRB annual meeting</i> , pages 1–16, Washington D.C., 2003. National Academy Press.
[CH94]	Colin F. Camerer and Teck-Hua Ho. Violations of the betweenness axiom and nonlinearity in probability. <i>Journal of Risk and Uncertainty</i> , 8(2):167–196, March 1994.
[Chm05]	Thorsten Chmura. Analyse, Modellierung und Simulationen von Routenwahlver- halten. PhD thesis, Universität Duisburg-Essen, Physik, Astronomie - Universi- taet Duisburg-Essen, April 2005.
[Cho54]	Gustave Choquet. Theory of capacities. Annales de l'institut Fourier, 5:131–295, 1954.
[Cho56]	Noam Chomsky. Three models for the description of language. <i>IRE Transactions on Information Theory</i> , 2:113–124, 1956.
[CLRS01]	Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. Introduction to Algorithms. The MIT Press, second edition, September 2001.
[CNZ99]	Francisco Cribari-Neto and Spyros G. Zarkos. R: Yet another econometric programming environment. <i>Journal of Applied Econometrics</i> , 14(3):319–329, May/June 1999.
[CS70]	John Cocke and Jacob T. Schwartz. Programming languages and their com- pilers: Preliminary notes. Technical report, Courant Institute of Mathematical Sciences, New York University, 1970.
[CS09]	Richard D. Connors and Agachai Sumalee. A network equilibrium model with travellers' perception of stochastic travel times. <i>Transportation Research Part B: Methodological</i> , In Press, Corrected Proof, 2009.
[DG08]	Carlos F. Daganzo and Nikolas Geroliminis. An analytical approximation for the macroscopic fundamental diagram of urban traffic. <i>Transportation Research Part B: Methodological</i> , 42(9):771–781, 2008.
[Dia71]	Robert B. Dial. A probabilistic multipath traffic assignment model which obviates path enumeration. <i>Transportation Research</i> , $5(2)$:83–111, June 1971.
[Dij55]	E. W. Dijkstra. A note on two problems in connexion with graphs. In <i>Numerische Mathematik</i> , volume I, pages 269–271, 1955.
[DMKSD06]	Benedetto De Martino, Dharshan Kumaran, Ben Seymour, and Raymond J. Dolan. Frames, biases, and rational decision-making in the human brain. <i>Science</i> , 313(5787):684–687, 2006.
[DS05]	G.B. Davies and S.E. Satchell. Continuous cumulative prospect theory and individual asset allocation. Cambridge Working Papers in Economics CWPE 0467, Faculty of Economics (formerly DAE), University of Cambridge, April 2005. available at http://www.dectech.org/People_Greg.html.
[Ear68]	Jay Earley. An Efficient Context-Free Parsing Algorithm. PhD thesis, Carnegie-Mellon University, Pittsburgh PA, Dept. Of Computer Science, August 1968. Accession Number: AD0677685.
[Ear70]	J. Earley. An efficient context-free parsing algorithm. Communications of the ACM , $13(2):94-102$, 1970.

[Edw92]	A. W. F. Edwards. <i>Likelihood</i> . The Johns Hopkins University Press, Baltimore, expanded edition, October 1992.
[EH99]	N. Christofides E. Hadjiconstantinou. An efficient implementation of an algorithm for finding k shortest simple paths. <i>Networks</i> , 34(2):88–101, 1999.
[Ell61]	Daniel Ellsberg. Risk, ambiguity, and the savage axioms. <i>The Quarterly Journal of Economics</i> , 75(4):643–669, November 1961.
[Epp98]	David Eppstein. Finding the k shortest paths. SIAM Journal on Computing, 28(2):652–673, 1998.
[Eps94]	Seymour Epstein. Integration of the cognitive and the psychodynamic uncon- scious. <i>American Psychologist</i> , 49(8):709–724, August 1994.
[ER07]	Ido Erev and Alvin E. Roth. Multi-agent learning and the descriptive value of simple models. <i>Artificial Intelligence</i> , 171(7):423–428, May 2007.
[FB09]	M. Fosgerau and M. Bierlaire. Discrete choice models with multiplicative error terms. <i>Transportation Research Part B: Methodological</i> , 43(5):494–505, 2009.
[FCHLvN03]	Stella Fiorenzo-Catalano, Sascha Hoogendoorn-Lanser, and Rob van Nes. Choice set composition modelling in multi-modal travelling. In 10th International Conference on Travel Behaviour Research, Lucerne, August 2003.
[FCvNB04]	Stella Fiorenzo-Catalano, Rob van Nes, and Piet H.L Bovy. Choice set gen- eration for multi-modal travel analysis. <i>European Journal of Transport and</i> <i>Infrastructure Research</i> , 4(2):195–209, 2004.
[Fis22]	Ronald Aylmer Fisher. On the mathematical foundations of theoretical statis- tics. <i>Philosophical Transactions of the Royal Society of London</i> , 222(A):309–368, 1922.
[Fis70]	Peter C. Fishburn. <i>Utility Theory For Decision Making</i> . Operations Research Society of America. Publications in operations research. Wiley, 1970.
[Fis81]	Peter C. Fishburn. Subjective expected utility: A review of normative theories. <i>Theory and Decision</i> , 13(2):139–199, June 1981.
[FKP04]	Manuel Fehler, Franziska Klügl, and Frank Puppe. Techniques for analysis and calibration of multi-agent simulations. In <i>Engineering Societies in the Agents World V</i> , number 3451 in Lecture Notes in Computer Science, pages 305–321. Springer Berlin / Heidelberg, 2004.
[FKP06]	Manuel Fehler, Franziska Klügl, and Frank Puppe. Approaches for resolving the dilemma between model structure refinement and parameter calibration in agent-based simulations. In AAMAS '06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems, pages 120–122, New York, NY, USA, May 2006. ACM.
[FT87]	Michael L. Fredman and Robert Endre Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. <i>Journal of the ACM (JACM)</i> , 34(3):596–615, July 1987.
[FW97]	H.P. Fennema and P.P. Wakker. Original and cumulative prospect theory: A discussion of empirical differences. <i>Journal of Behavioral Decision Making</i> , 10(1):53–64, 1997.
[GG96]	Gerd Gigerenzer and Daniel G. Goldstein. Reasoning the fast and frugal way: Models of bounded rationality. <i>Psychological Review</i> , 103(4):650–669, 1996.

[GG02]	Daniel G. Goldstein and Gerd Gigerenzer. Models of ecological rationality: The recognition heuristic. <i>Psychological Review</i> , 109(1):75–90, January 2002.
[Goo98]	P. Goodwin. The end of equilibrium. In T. Gärling, T. Laitila, and K. Westin, editors, <i>Theoretical Foundations of Travel Choice Modeling</i> , pages 185–195. Pergamon, 1st edition, June 1998. 978-0080430621.
[GP06]	Faruk Gul and Wolfgang Pesendorfer. Random expected utility. $Econometrica,$ 74(1):121–146, 2006.
[GSP ⁺ 89]	Gerd Gigerenzer, Zeno Swijtink, Theodore Porter, Lorraine Daston, John Beatty, and Lorenz Kruger. <i>The Empire of Chance: How Probability Changed Science and Everyday Life.</i> Cambridge University Press, The Pitt Building, Trumpington Stree, Cambridge CB2 1RP, UK, 1989. Reprint in 1990.
[HB07]	David A Hensher and Kenneth J Button, editors. <i>Handbook of Transport Modelling</i> , volume 1. Emerald Group Publishing Limited, 2nd edition, October 2007.
[HBZ04]	E. A. Hansen, D. S. Bernstein, and S. Zilberstein. Dynamic programming for partially observable stochastic games. In <i>Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI-04)</i> , pages 709–715, San Jose, California, July 2004. Association for the Advancement of Artificial Intelligence, AAAI Press.
[HG05]	John M.C. Hutchinson and Gerd Gigerenzer. Simple heuristics and rules of thumb: Where psychologists and behavioural biologists might meet. <i>Behavioural Processes</i> , 69(2):97–124, May 2005. Proceedings of the meeting of the Society for the Quantitative Analyses of Behavior (SQAB 2004).
[HK05]	Jiawei Han and Micheline Kamber. <i>Data Mining: Concepts and Techniques</i> . Morgan Kaufmann, San Francisco, CA, USA, second edition, November 2005.
[HMS03]	John E. Hershberger, Matthew Maxel, and Subhash Suri. Finding the k shortest simple paths: a new algorithm and its implementation. In <i>Proceedings 5th Workshop Algorithm Engineering & Experiments (ALENEX)</i> , pages 26–36. SIAM, January 2003.
[HMU01]	John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. Introduction to Automata Theory, Languages, and Computation. Addison Wesley, second edition, 2001.
[Hoc05]	H.H. Hochmair. Investigating the effectiveness of the least-angle strategy for wayfinding in unknown street networks. <i>Environment and Planning B, Planning and Design</i> , $32(5):673-691$, 2005.
[Hor84]	Joel L. Horowitz. The stability of stochastic equilibrium in a two-link transportation network. <i>Transportation Research Part B: Methodological</i> , 18(1):13–28, February 1984.
[HSB07]	John Hershberger, Subhash Suri, and Amit Bhosle. On the difficulty of some shortest path problems. <i>ACM Trans. Algorithms</i> , 3(1):5, 2007.
[JM03]	Víctor M. Jiménez and Andrés Marzal. A lazy version of eppstein's k shortest paths algorithm. In <i>Proceedings Experimental and Efficient Algorithms: Second International Workshop, WEA 2003</i> , pages 179–190, Ascona, Switzerland, May 2003.
[Kah02]	Daniel Kahneman. Maps of bounded rationality: A perspective on intuitive judgment and choice. In Tore Frängsmyr, editor, <i>The Nobel Prizes 2002</i> , pages 449–489, Aula Magna, Stockholm University, December 2002. Nobel Foundation. presented by Professor Torsten Persson, Chairman of the Prize Committee.
---------	---
[Kas65]	T. Kasami. An efficient recognition and syntax-analysis algorithm for context-free languages. Technical report, Air Force Cambridge Research Lab, Bedford, MA, 1965. Scientific report AFCRL-65-758.
[KB02]	F. Klügl and A. L. C. Bazzan. Simulation of adaptive agents: Learning heuris- tics for route choice in a commuter scenario. In <i>Proceedings of the first in-</i> <i>ternational joint conference on Autonomous Agents and Multi-Agent Systems,</i> <i>AAMAS</i> , volume 1, pages 217–218, Bologna, Italy, July 2002. New York: ACM Press. Extended Abstract.
[KB04]	Franziska Klügl and Ana L. C. Bazzan. Route decision behaviour in a commuting scenario: Simple heuristics adaptation and effect of traffic forecast. <i>Journal of Artificial Societies and Social Simulation</i> , 7(1), January 2004.
[KBW03]	F. Klügl, A. L. C. Bazzan, and J. Wahle. Selection of information types based on personal utility - a testbed for traffic information markets. In <i>Proceedings of the</i> <i>second international joint conference on Autonomous Agents and Multi-Agent</i> <i>Systems, AAMAS</i> , Melbourne, Australia, July 2003. New York: ACM Press.
[KF02]	D. Kahneman and S. Frederick. Representativeness revisited: Attribute sub- stitution in intuitive judgment. In T. Gilovich, D. Griffin, and D. Kahneman, editors, <i>Heuristics and Biases: The Psychology of Intuitive Judgment</i> , chapter 2, pages 49–81. Cambridge University Press, New York, 2002.
[KF05]	D. Kahneman and S. Frederick. A model of heuristic judgment. In K.J. Holyoak and R.G. Morrison, editors, <i>The Cambridge Handbook of Thinking and Reasoning</i> , chapter 12, pages 267–293. Cambridge University Press, 2005.
[KF06]	D. Kahneman and S. Frederick. Frames and brains: elicitation and control of response tendencies. <i>Trends in Cognitive Sciences</i> , 11(2):45–46, February 2006.
[Kni21]	Frank H. Knight. <i>Risk, Uncertainty, and Profit.</i> Number 31 in Hart, Schaffner, and Marx Prize Essays. Houghton Mifflin Company, Boston and New York, 1921.
[Knu98]	Donald E. Knuth. Art of Computer Programming, volume 3. Addison-Wesley Professional, 3rd edition, April 1998. Sorting and Searching.
[KR88]	Brian W. Kernighan and Dennis M. Ritchie. <i>C Programming Language</i> . Prentice Hall Software. Prentice Hall PTR, 2nd edition, April 1988.
[KR00]	John H. Kagel and Alvin E. Roth. The dynamics of reorganization in matching markets: A laboratory experiment motivated by a natural experiment. <i>The Quarterly Journal of Economics</i> , 115(1):201–235, February 2000.
[KT79]	Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. <i>Econometrica</i> , 47(2):263–292, March 1979.
[LF91]	R. Duncan Luce and Peter C. Fishburn. Rank- and sign-dependent linear utility models for finite first-order gambles. <i>Journal of Risk and Uncertainty</i> , 4(1):29–59, January 1991.
[Lit94]	Michael L. Littman. Markov games as a framework for multi-agent reinforce- ment learning. In <i>Proceedings of the 11th International Conference on Machine Learning</i> , pages 157–163, New Brunswick, NJ, 1994. Morgan Kaufmann.

[Lov91]	William S. Lovejoy. Computationally feasible bounds for partially ob- served markov decision processes. <i>Operations Research</i> , 39(1):162–175, Jan- uary/February 1991.
[Luc59]	R. Duncan Luce. Individual Choice Behavior: A Theoretical Analysis. Dover Publications, New York, 1959.
[MA06]	Nina Mazar and Dan Ariely. Dishonesty in everyday life and its policy implica- tions. Journal of Public Policy and Marketing, 25(1):117–126, Spring 2006.
[Man77]	Charles F. Manski. The structure of random utility models. <i>Theory and Decision</i> , 8(3):229–254, July 1977.
[Mar60]	J. Marschak. Binary choice constraints on random utility indications. In K. Arrow, editor, <i>Stanford Symposium on Mathematical Methods in the Social Sciences</i> , pages 312–329, Stanford, CA, 1960. Stanford University Press.
[Mar71]	A. A. Markov. Extension of the limit theorems of probability theory to a sum of variables connected in a chain. In R. Howard, editor, <i>Dynamic Probabilistic Systems</i> , chapter Appendix B. John Wiley and Sons, August 1971. reprinted.
[Mar04]	Robert Martin. The st. petersburg paradox. Online encyclopedia, Stanford University, http://plato.stanford.edu/entries/paradox-stpetersburg/, July 2004. Stanford Encyclopedia of Philosophy.
[McF74]	D. McFadden. Conditional logit analysis of qualitative choice behavior. <i>Frontiers of Econometrics</i> , 1974.
[McF99]	Daniel McFadden. Rationality for economists? Journal of Risk and Uncertainty, 19(1–3):73–105, December 1999.
[McN07]	Michale G. McNally. The four step model. In David A Hensher and Kenneth J Button, editors, <i>Handbook of Transport Modelling</i> , volume 1, chapter 3, pages 35–41. Emerald Group Publishing Limited, 2nd edition, October 2007.
[Mil56]	George A. Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. <i>Psychological Review</i> , 63:81–97, 1956.
[Mil85]	John Milnor. On the concept of attractor. Communications in Mathematical Physics, 99(2):177–195, July 1985.
[MM02]	Gurmeet Singh Manku and Rajeev Motwani. Approximate frequency counts over data streams. In VLDB '02: Proceedings of the 28th international conference on Very Large Data Bases, pages 346–357. VLDB Endowment, 2002.
[MMS90]	M. Migliore, V. Martorana, and F. Sciortino. An algorithm to find all paths between two nodes in a graph. J. Comput. Phys., 87(1):231–236, 1990.
[MT00]	D. McFadden and K. Train. Mixed mnl models of discrete response. <i>Journal of Applied Econometrics</i> , 15:447–470, 2000.
[Nar04]	Louis Narens. A new foundation for support theory. Technical Report MBS 04-04, Mathematical Behavioral Sciences at UC Irvine, University of California, Irvine, January 2004.
[NCK ⁺ 06]	Ehren L. Newman, Jeremy B. Caplan, Matthew P. Kirschen, Igor O. Korolev, Robert Sekuler, and Michael J. Kahana. Learning your way around town: How virtual taxicab drivers learn to use both layout and landmark information. <i>Cog-</i> <i>nition</i> , 2006. In Press, Corrected Proof, Available online 1 August 2006.

[Ned99]	Zhivko Prodanov Nedev. Finding an even simple path in a directed planar graph. SIAM Journal on Computing, $29(2)$:685–695, 1999.
[Osb03]	Martin J. Osborne. An Introduction to Game Theory. Oxford University Press, New York, New York, USA, August 2003.
[Ram02]	Michael Scott Ramming. <i>Network Knowledge and Route Choice</i> . PhD thesis, Massachusetts Institute of Technology, February 2002.
[Reb89]	Arthur S. Reber. Implicit learning and tacit knowledge. <i>Journal of Experimental Psychology: General</i> , 118(3):219–235, September 1989.
[Ric59]	F. J. Richards. A flexible growth function for empirical use. <i>Journal of Experimental Botany</i> , 10(2):290–301, 1959.
[RPM04]	Roger P. Roess, Elena S. Prassas, and William R. McShane. <i>Traffic Engineer-</i> <i>ing.</i> Pearson Education, Inc., Upper Saddle River, New Jersey 07458, third (international) edition, 2004.
[RT97]	Yuval Rottenstreich and Amos Tversky. Unpacking, repacking, and anchoring: Advances in support theory. <i>Psychological Review</i> , 104(2):406–415, April 1997.
[Rub07]	Ariel Rubinstein. Instinctive and cognitive reasoning: A study of response times. <i>Economic Journal</i> , 117(523):1243–1259, October 2007.
[RW02]	M. Raubal and S. Winter. Enriching wayfinding instructions with local land- marks. In Max J. Egenhofer and David M. Mark, editors, <i>Geographic Informa-</i> <i>tion Science</i> , volume 2478 of <i>Lecture Notes in Computer Science</i> , pages 243–259. Springer, Berlin, 2002.
[RW08]	Marc Oliver Rieger and Mei Wang. Prospect theory for continuous distributions. <i>Journal of Risk and Uncertainty</i> , 2008. to appear.
[Sch03]	Ulrich Schmidt. Reference dependence in cumulative prospect theory. <i>Journal of Mathematical Psychology</i> , 47(2):122–131, 2003.
[Sha53]	L. S. Shapley. Stochastic games. Proceedings of the National Academy of Sciences of the United States of America, 39(10):1095–1100, October 1953.
[Sim55]	Herbert A. Simon. A behavioral model of rational choice. <i>The Quarterly Journal of Economics</i> , 69(1):99–118, February 1955.
[Sim56]	Herbert A. Simon. Rational choice and the structure of the environment. <i>Psychological Review</i> , 63:129–138, 1956.
[Sim79]	Herbert A. Simon. Rational decision making in business organizations. <i>The American Economic Review</i> , 69(4):493–513, September 1979.
[Sim82]	Herbert Simon. Models of Bounded Rationality. MIT Press, 1982.
[Sim86]	Herber A. Simon. Rationality in psycghology and economics. <i>The Journal of Business</i> , 59(4):209–224, October 1986.
[Sim92]	Herbert A. Simon. Rational decision-making in business organizations. In Assar Lindbeck, editor, <i>Nobel Lectures, Economics 1969-1980</i> , pages 343–371, Singapore, December 1992. Nobel Foundation, World Scientific Publishing Co. Nobel Prize in Economics 1978.
[Sip05]	Michael Sipser. Introduction to the Theory of Computation. Course Technology, second edition, February 2005.

[SK04]	Metin Senbil and Ryuichi Kitamura. Reference points in commuter departure time choice: A prospect theoretic test of alternative decision frames. <i>Journal of Intelligent Transportation Systems</i> , 8(1):19–31, January 2004.
[SLB08]	Yoav Shoham and Kevin Leyton-Brown. <i>Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations.</i> Cambridge University Press, New York, New York, USA, December 2008.
[Slo96]	Steven A. Sloman. The empirical case for two systems of reasoning. <i>Psychological Bulletin</i> , 119(1):3–22, January 1996.
[Slo02]	Steven A. Sloman. Two systems of reasoning. In T. Gilovich, D. Griffin, and D. Kahneman, editors, <i>Heuristics and Biases: the Psychology of Intuitive Judg-</i> ment, pages 379–396. Cambridge University Press, Cambridge, 2002.
[SPG03]	Y. Shoham, R. Powers, and T. Grenager. Multi-agent reinforcement learning: a critical survey. Technical report, Stanford University, 2003.
[SPG04]	Yoav Shoham, Robert Powers, and Trond Grenager. On the agenda(s) of research on multi-agent learning. In AAAI 2004 Fall Symposium on Artificial Multi-Agent Learning. AAAI Press, 2004.
[SPG07]	Yoav Shoham, Robert Powers, and Trond Grenager. If multi-agent learning is the answer, what is the question? <i>Artificial Intelligence</i> , 171(7):365–377, 2007. special issue on Foundations of Multi-Agent Learning.
[SSC+05]	R. Selten, M. Schreckenberg, T. Chmura, T. Pitz, S. Kube, S.F. Hafstein, R. Chrobok, A. Pottmeier, and J. Wahle. Experimental investigation of day- to-day route-choice behavior and network simulations of autobahn traffic in north rhine-westphalia. In M. Schreckenberg and R. Selten, editors, <i>Human</i> <i>Behaviour and Traffic Networks</i> , pages 1–21. Springer, Heidelberg, June 2005. 978-3540212201.
[Ste46]	S. S. Stevens. On the theory of scales of measurement. Science, $103(2684){:}677{-}680,$ June 1946.
[Sto07]	Peter Stone. Multiagent learning is not the answer. it is the question. Artificial Intelligence, 171(7):402–405, May 2007. Foundations of Multi-Agent Learning.
[Str00]	Bjarne Stroustrup. The $C++$ Programming Language. Addison Wesley, 201 W. 103rd Street, Indianapolis, IN 46290, special edition, 2000.
[Sut84]	Richard S. Sutton. <i>Temporal credit assignment in reinforcement learning</i> . PhD thesis, University of Massachusetts Amherst, July 1984.
[Sut88]	Richard S. Sutton. Learning to predict by the methods of temporal differences. <i>Machine Learning</i> , 3(1):9–44, August 1988.
[SV00]	Peter Stone and Manuela Veloso. Multiagent systems: A survey from a machine learning perspective. <i>Autonomous Robots</i> , 8(3):345–383, June 2000.
[TK83]	Amos Tversky and Daniel Kahneman. Extensional versus intuitive reason- ing: The conjunction fallacy in probability judgment. <i>Psychological Review</i> , 90(4):293–315, October 1983.
[TK92]	Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. <i>Journal of Risk And Uncertainty</i> , 5(4):297–323, October 1992.

[TK94]	Amos Tversky and Derek J. Koehler. Support theory: A nonextensional repre- sentation of subjective probability. <i>Psychological Review</i> , 101(4):547–567, Octo- ber 1994.
[Tra00]	Transportation Research Board. <i>Highway Capacity Manual: 2000.</i> Transportation Research Board, Washington, D.C., lslf edition, December 2000.
[vdZC05]	N.J. van der Zijpp and S. Fiorenzo Catalano. Path enumeration by finding the constrained k-shortest paths. <i>Transportation Research Part B: Methodological</i> , 39(6):545–563, July 2005.
[vN28]	John von Neumann. Zur theorie der gesellschaftspiele. Mathematische Annalen, $100(1)$:295–320, December 1928.
[vNM07]	John von Neumann and Oskar Morgenstern. <i>Theory of Games and Economic Behavior</i> . Princeton University Press, Princeton University, 60th anniversary edition edition, March 2007. 1st Edition: 1944.
[Vov97]	Peter Vovsha. The cross-nested logit model: Application to mode choice in the tel-aviv metropolitan area. <i>Transportation Research Record</i> , 1607:6–15, 1997.
[Wah02]	Joachim Wahle. Information in Intelligent Transportation Systems - Information in Intelligenten Transportsystemen. PhD thesis, Gerhard-Mercator-Universität, Physik, Astronomie - Gerhard-Mercator-Universität, January 2002.
[War52]	John Glen Wardrop. Some theoretical aspects of road traffic research. In <i>Proceedings of Institution of Civil Engineers, Part II</i> , volume 1, pages 325–378, London, 1952.
[Wat89]	Christopher J. C. H. Watkins. <i>Learning with Delayed Rewards</i> . PhD thesis, Cambridge University, 1989. Psychology Department.
[WD92]	Christopher J. C. H. Watkins and Peter Dayan. Technical note: Q-learning. <i>Machine Learning</i> , 8(3):279–292, 5 1992.
[WG96]	George Wu and Richard Gonzalez. Curvature of the probability weighting function. Management Science, $42(12)$:1676–1690, 1996.
[Wim76]	W. C. Wimsatt. Reductionism, levels of organization and the mind-body prob- lem. In G. G. Globus, G. Maxwell, and I. Savodnik, editors, <i>Consciousness and</i> <i>the Brain: A Scientific and Philosophic Inquiry</i> , pages 199–267. Plenum Press, New York, 1976.
[Woo03]	Jeffrey Wooldridge. Introductory Econometrics - A Modern Approach. South-Western, 2nd edition, 2003.
[WT93]	Peter Wakker and Amos Tversky. An axiomatization of cumulative prospect theory. <i>Journal of Risk and Uncertainty</i> , 7(2):147–175, October 1993.
[You67]	Daniel H. Younger. Recognition and parsing of context-free languages in time n^3 . Information and Control, 10(2):189–208, 1967.

Appendix A Technologies And Algorithms

This chapter is an appendix because it is only marginally related to the subject of this thesis, even though it is important for its infrastructure. Three technologies are here presented, namely: supernetworks; Context-Free Grammars (CFG), for navigating in a super-network; and path generation algorithms for the super-network. A super-network can be seen as an edge-labelled weighted directed graph (digraph) that is generated from another (compacter) edge-labelled weighted digraph. The other technologies are modifications of existing algorithms for using this special graph (the super-network).

The motivation for adopting the super-networks in the traffic network representation is because it presents an efficient alternative for multi-modal route generation. This concept was informally presented in [CFCLB03, FCHLvN03, FCvNB04] and based on this structure it is here formalised and an efficient way to "navigate" in it is also developed (using a modified Dijkstra [Dij55] algorithm). The advantage in using the methods described here is twofold: performance and flexibility. In [vdZC05] the authors use a restricted form of super-networks and from it they generate several routes, eliminating the invalid ones afterwards – the method shown here does not generate any invalid route. In another approach to the multi-modal route choice [BHL05], the "multi" is restricted to a maximum of three modi combination (Access + Main + Egress) – the new approach shown here has not such restriction. Expanding the amount of modi, in [BJM98] it is suggested the use of a grammar; but a prohibitive approach is used there by adopting a bottom-up parser, which needs the whole word to give any evaluation about its validity (almost the same procedure as in [vdZC05]). All these issues are not present in the solution presented here.

A.1 Notation, Typography, And Conventions

Before presenting the algorithms it is necessary to establish a common notation that crosses several fields. These definitions for the set theory are:

- Upper case Roman letters, as in A, B, and C, represent sets.
- Lower case Roman letters, such as a, b, and c, represent elements in sets.
- Greek letters have a context dependant meaning, i.e. their meaning will be explained as they are presented (can be sets, usually in upper case Greek letters as Σ , or elements, usually in lower case as α).
- Upper case Roman letters in different typefaces are used to represent data-structures or set of sets, as in N (a set of networks); but the type faces as in 𝔅 are used exclusively for sets of sets. In this particular case N is one of the possible sets in 𝔅 − 𝔅 is the universal set of N.
- |A| corresponds to the amount of elements in the set A.

- A^B represents a map from B into A, i.e. being A and B sets then a structure exists that maps elements of B into elements of A.
- Get : $\mathfrak{X} \times B \mapsto A$ is the function that recovers the A element mapped by the given B element, or *NIL* if no A element is associated with the B element. The association is stored into a map $X \in \mathfrak{X}$, where \mathfrak{X} is the set of all possible maps.
- Put: X × B × A ↦ {} is the function to maps an element of B into an element of A, stored in X ∈ X.
- A^n is an *n*-ary Cartesian product, i.e. it represents an ordered tuple of *n* elements from the set/type *A*, with $n \in \mathbb{N}^*$.
- Get: Aⁿ × [1, n] → A is the function that returns the *i*th element (second argument) of a tuple (first argument) it is assumed that the first element's index is 1 and the last n.
- $A \setminus B$ represents the set A without the elements present in B: $A \setminus B = A B$.

For the graph theory, the definitions are:

- G(V, E) represents a graph with vertex set V and edge set E. It is also possible to refer to these elements by subscriptions: V_G is the vertex set of graph G and E_G its edge set. If "G(V, E)" is in the right-side of an attribution, as in $G' \leftarrow G(V, E)$, it means that a new graph is created with vertex set V and edge set E.
- \mathfrak{G} is the set of all graphs. This notation is used to specify domain and codomain of functions. The same logic is valid for \mathfrak{V} and \mathfrak{E} to denote the set of all sets of vertices and all sets of edges.
- Source : $E \mapsto V$ returns for a digraph the corresponding source vertex for the given edge.
- Target : $E \mapsto V$ as in Source() but returns the target vertex.
- In: V → € is a function that returns for a digraph all edges that have the given vertex as target: In(v) = E' and E' = {∀e ∈ E' | Target(e) = v}.
- Out : $V \mapsto \mathfrak{E}$ as in In() but returns all edges that have the given vertex as source: $\operatorname{Out}(v) = E'$ and $E' = \{ \forall e \in E' \mid \operatorname{Source}(e) = v \}.$

A.2 Super-Network

A super-network, after [CFCLB03, FCHLvN03, FCvNB04], is a way to structure a weighted digraph specially designed for representing a traffic network. The principle is to stratify a graph according to its modi, i.e. each stratus corresponds to a network that contains only the edges associated with the modus it represents. Then several interconnected networks form a supernetwork. In Fig. A.1 this process (in a schematic view) can be seen, where G is the original graph – representing a traffic network – and S is the resulting super-network, with its several networks $(N_1, N_2, \text{ and } N_3)$.

For simplicity the word network is used to refer to a sub-graph/network of a super-network and this last (super-network) refers to the resulting structure of a graph after this "segmentation" process. Unfortunately neither a formal definition of the super-network structure nor the transformation algorithms were given. The guidelines are to have each modus represented in a network in the super-network. Therefore here the definition of super-networks is made (in a precise mathematical definition) as well as the super-network generation algorithm (derived from the definition).



Figure A.1: From a digraph G to a super-network \mathcal{S} .

A.2.1 Super-network: Formalisation

To formalise the concept of super-networks some definitions are necessary. First it is assumed that a super-network S is derived or generated from an initial labelled digraph G. Moreover, each edge $e \in E_G$ has a set of labels $L_e \neq \{\}$ and each label supposedly corresponds to a modus – therefore from now on the label concept instead of the modus concept is used in the context of networks and super-networks. The set of all labels in a graph is: $L_G = \bigcup_{e \in E_G} L_e$.

A super-network for its turn is defined by $\mathcal{S}(G, \mathcal{N}, S_0, C, \mathcal{L})$, where G is the original digraph, \mathcal{N} the network set, $S_0 \in \mathcal{N}$ is a special network (the start network), C the connecting edge set (explain latter), and \mathcal{L} the label set. It is also expected that a function Labels : $E_{\mathcal{S}} \mapsto \mathfrak{L}$ is given that returns the labels associated with a given edge, where \mathfrak{L} is the set of label sets and E_G is the edge set from G. Then $E_{\mathcal{S}} = E_G \cup E_{\mathcal{N}} \cup C$, where $E_{\mathcal{N}} = \bigcup_{N \in \mathcal{N}} E_N$ is the set of all edges of all networks in the network set \mathcal{N} (E_N is the edge set of the network N). The same semantic has also Labels : $\mathfrak{E} \mapsto \mathfrak{L}$ but for a set of edges, returning $L = \bigcup_{e \in E} \mathtt{Labels}(e)$.

A network $N(V, E) \in \mathcal{N}$ is also a digraph and resents a modus/label from the traffic network, given by G. Therefore a function is necessary: Equiv : $V_{\mathcal{N}} \cup E_{\mathcal{N}} \mapsto V_G \cup E_G$. This function returns the reference element of G that is being represented by the given network element. It implies that all elements in a network have a corresponding element in G, i.e.: $\exists x \in V_{\mathcal{N}} \cup E_{\mathcal{N}} \mid \text{Equiv}(x) =$ NIL. With this function it can also be defined IsEquiv : $V_{\mathcal{N}} \cup E_{\mathcal{N}} \times V_G \cup E_G \mapsto \{true, false\}$ that asserts if two elements are semantically equivalent, i.e. if the first argument semantically corresponds to the second element (in G). This function is formally defined as: IsEquiv $(x, y) \equiv$ (Equiv $(x) = y) \land$ (Equiv $(x) \neq NIL$).

Let a super-network S has an initial graph G, a set of networks \mathcal{N} , a set of labels \mathcal{L} , and a set of connecting edges C, then all statements below hold.

- Every network maps all vertices from G exactly once.
- If an edge from G is represented/mapped in a network it is the only representation in that particular network, maintaining its semantic.
- Each network has only one label and no two networks have the same label, i.e. exactly one network per label.

The above mentioned statements are formally expressed in Eq. A.1.

 $\begin{array}{l} \forall N \in \mathcal{N} \Rightarrow |V_N| = |V_G| \\ \forall N \in \mathcal{N}, \forall v_N \in V_N, \exists v_G \in V_G, \not\exists v'_G \in V_G \mid \\ \mathbf{IsEquiv}(v_N, v_G) = true \land \mathbf{IsEquiv}(v_N, v'_G) = true \land v_G \equiv v'_G \\ \forall N \in \mathcal{N} \Rightarrow \forall v_N \in V_N, \not\exists v'_N \in V_N \mid \mathbf{Equiv}(v_N) = \mathbf{Equiv}(v'_N) \land v_N \neq v'_N \\ \forall N \in \mathcal{N}, \forall e_N \in E_N \Rightarrow \mathbf{Equiv}(\mathbf{Source}(e_N)) = \mathbf{Source}(\mathbf{Equiv}(v_N) \\ \land \mathbf{Equiv}(\mathbf{Target}(e_N)) = \mathbf{Target}(\mathbf{Equiv}(v_N) \\ \forall N \in \mathcal{N} \Rightarrow \forall e_N, e'_N \in E_N, \exists e_G \in E_G \mid e_N \neq e'_N \land \mathbf{IsEquiv}(e_N, e_G) \\ \oplus \mathbf{IsEquiv}(e'_N, e_G) \\ \forall N \in \mathcal{N}, N \neq de_N, e'_N \in E_N \mid \mathbf{Labels}(e_N) = \mathbf{Labels}(e'_N) \land |\mathbf{Labels}(e_N)| = 1 \\ \forall N, M \in \mathcal{N}, N \notin M \Rightarrow \forall e_N \in E_N \not\exists e_M \in E_M \mid \mathbf{Labels}(e_N) = \mathbf{Labels}(e_M) \end{array}$

Another special element of the super-network structure is the start network S_0 . This network must also respect all rules in Eq. A.1 but with one extra restriction: $E_{S_0} = \{\}$, i.e. it has no edges and it is the only network with an empty edge set: $\exists N \in \mathcal{N} \mid E_N = \{\} \land N \neq S_0$.

To completely specify a super-network the sets C and \mathcal{L} must be defined. The set C is a special edge set that have no correspondence with any edge in G, they are the connecting edges. These edges have labels too and their set is L_C and $L_C \subset \mathcal{L} \wedge L_C \cap L_G = \{\}$, meaning that L_C is disjoint from the regular labels (L_G) . Thus the set C is defined below and formally in Eq. A.2.

- Each network is connected to all other networks only through equivalent vertices and just one connecting edge $c \in C$ connects any two vertices
- The edges in C have no equivalent representation in G, another formalisation would be: $\forall c \in C \mid \texttt{Equiv}(c) = NIL.$
- The label $\alpha \in L_C$ is reserved for connecting the start network S_0 to all others.
- Similar to α , ω is also reserved for connecting the other networks back to the start network.
- The edges in C do not belong to any network: $C \cap E_{\mathcal{N}} = \{\}$.

$$\forall N, M \in \mathcal{N} \exists c \in C \mid \text{Source}(c) \in V_N \land \text{Target}(c) \in V_M \land N \neq M \\ \land \text{Equiv}(\text{Source}(c)) = \text{Equiv}(\text{Target}(c)) \\ \forall c \in C \quad \exists c' \in C \mid \text{Source}(c) = \text{Source}(c') \land \text{Target}(c) = \text{Target}(c') \\ \forall c \in C \quad \exists e_G \in G \mid \text{IsEquiv}(c, e_G) = true \\ \forall v \in V_{S_0} \exists c, c' \in C \mid \text{Source}(c) = v \land \text{Labels}(c) = \{\alpha\} \\ \land \text{Target}(c') = v \land \text{Labels}(c') = \{\omega\} \\ \exists c \in C \mid (\text{Labels}(c) = \{\alpha\} \land \text{Source}(c) \notin V_{S_0}) \\ \lor (\text{Labels}(c) = \{\omega\} \land \text{Target}(c) \notin V_{S_0}) \\ \end{cases}$$
(A.2)

It also implies that $|L_C| \geq 3$, since all edges have labels and $\alpha \in L_C$ as well as $\omega \in L_C$.

The super-network definition can be optimised by adding the rule in Eq. A.3 for excluding unnecessary vertices in \mathcal{N} . It is assumed that the super-network was already generated and the function Remove : $V_{\mathcal{N}} \mapsto \mathfrak{S}$ is provided with the following behaviour (where \mathfrak{S} the set of all super-networks): the function removes the given vertex as well as any edge connected to it, returning a new super-network without these elements.

$$\forall N \in \mathcal{N} \setminus \{S_0\} \; \forall v \in V_N \mid \operatorname{In}(v) \subset C \land \operatorname{Out}(v) \subset C \Rightarrow \operatorname{Remove}(v) \tag{A.3}$$

What Eq. A.3 does is to select all vertices that have only connecting edges and eliminate them. They would be like the vertices in the super-network S from Fig. A.1 that have no full lines reaching them.

A.2.2 Super-network: Generation Algorithm

In this section the implemented algorithm for generating a super-network that attains to the rules in Sec. A.2.1 (including the optimisation step of Eq. A.3), is presented. For the algorithms extra functions are necessary:

- Copy : $V_G \mapsto V_N$ It creates a semantically equivalent copy from the given original vertex (from graph G) to be used in a network $N \in \mathcal{N}$.
- Copy : $E_G \times V_N \times V_N \times \mathcal{L} \mapsto E_N$ As the previous function it also makes a semantically equivalent copy but from an edges in G. The difference is that the edge is copied but its source set to be the second argument (a vertex resulting from the previous function), the target the third argument, and the label set by the fourth argument.
- Connect : $V_N \times V_N \times L_C \mapsto C$ This function connects two vertices assigning a label and returning a connecting edge it uses the first argument as source, the second as the target, and the third as the label.
- ConnLabel : $L_G \times L_G \mapsto L_C$ For the connecting edges it is necessary a label $l \in L_C$ and this label is only predefined for the edges connecting the S_0 to and from other networks α and ω , respectively. To recover the other labels this function is necessary. Because each network has only one label (except for the start network S_0) it is enough to specify as arguments the source label (first argument) and the target label (second).

With these above defined functions, a super-network is generated by the function Supernet(G) from Algo. A.1, where G is the input digraph, from which a super-network is wanted. This algorithm generates a super-network that obeys all rules in Sec. A.2.1, including Eq. A.3.

A.2.3 Super-network: Example

To give an idea of how a super-network looks like a series of images are here presented, actually one of the many possible visual representations. To avoid visual "pollution" in the illustrative figures neither arrows (to indicate edges' directions) nor multiple edges for the same vertex pair are depicted. This means that in Fig. A.2 each edge should be "seen" as actually two, one for each possible direction.

The process of creating a super-network starts with an input labelled digraph (as in Fig. A.2). In this figure each edge carries its labels, a subset from $L_G = \{b, c, w\}$. From this graph (in Fig. A.2) the creation process "extracts" the different networks, one for each label. The network set \mathcal{N} (before including the connecting edges) is in Fig. A.3, where the start network S_0 called *start* is already present. For convenience each network is marked with its corresponding label.

The next step is to connect all networks, the creation of the edge set C. In Fig. A.4 an aesthetic view of the final super-network S can be seen, which was generated from the graph in Fig. A.2. There, the dashed lines represent the connecting edges in a "reduced" form, to keep it visually clear (the edges in C). The dashed lines (in Fig. A.4) are a reminder of what they really are, depicted in Fig. A.5. There, all connections between any two networks, including the start network S_0 , are depicted (observe that each edge represents a pair of directed edges, one for each direction).

To make it even clearer, in Fig. A.6 a possible path in this super-network is shown. The path "jumps" over networks using the connecting edges (in this representations all other edges where omitted to avoid confusion). It can be noticed that the path starts and ends in the start network.

A.2.4 Navigation

For the navigation in the super-network some prerequisites are necessary:

• All paths must start and end in the start network S_0 .

Algorithm A.1: Supernet(•)

Data: G the input graph **Result**: A new super-network \mathcal{S} 1 $\mathcal{N} \leftarrow \{\}$; /* creates an empty network set */ /* creates the start network S_0 */ **2** $S_0 \leftarrow \{\};$ s forall $v \in V_G$ do 4 | $S_0 \leftarrow S_0 \cup \{ \text{Copy}(v) \};$ 5 end /* creates the regular networks */ 6 $H \leftarrow \{\}$; /* creates an empty map ($\mathcal{N} \mapsto \mathcal{L}$), with manipulation functions Get and Put */ 7 forall $lbl \in L_G$ do $N \leftarrow \text{Network}(G, lbl);$ 8 $\mathcal{N} \leftarrow \mathcal{N} \cup \{N\} \; ; \textit{/* see Algo. A.2 */}$ 9 Put(H, N, lbl); /* associates the key N to the value lbl in map H */1011 end 12 $C \leftarrow \{\}$; /* creates an empty edge set */ /* connects all regular networks */ 13 $\mathcal{L} \leftarrow L_G;$ 14 forall $N \in \mathcal{N}$ do $l_N \leftarrow \texttt{Get}(H, N)$; /* gets the label associated with the source network N15*/ $\mathbf{16}$ forall $M \in \mathcal{N} \setminus \{N\}$ do $l_M \leftarrow \texttt{Get}(H, M)$; /* gets the label associated with the target network 17 M * / $l_C \leftarrow \text{ConnLabel}(l_N, l_M);$ 18 19 $\mathcal{L} \leftarrow \mathcal{L} \cup \{l_C\};$ $C \leftarrow C \cup \texttt{ConnectNet}(N, M, l_C)$; /* see Algo. A.3 */ 20 $\mathbf{21}$ end 22 end /* connects all regular networks to the start network */ **23** $\mathcal{L} \leftarrow \mathcal{L} \cup \{\alpha, \omega\};$ 24 forall $N \in \mathcal{N}$ do 25 | $C \leftarrow C \cup \text{ConnectNet}(S_0, N, \alpha) \cup \text{ConnectNet}(N, S_0, \omega);$ 26 end 27 $\mathcal{N} \leftarrow \mathcal{N} \cup \{S_0\};$ 28 $\mathcal{S} \leftarrow \mathcal{S}(G, \mathcal{N}, S_0, C, \mathcal{L})$; /* creates a new super-network */ 29 return S;

Algorithm A.2: Network(•)

Data: G the input graph **Data**: lbl the corresponding label **Result**: A new network N 1 $V_N \leftarrow \text{Vertices}(G, lbl)$; /* see Algo. A.4 */ 2 $E_N \leftarrow \text{Edges}(G, lbl)$; /* see Algo. A.5 */ 3 $N \leftarrow N(V_N, E_N)$; /* creates a new network */ 4 return N;

Algorithm A.3: $ConnectNet(\bullet)$
Data : N the source network
Data : M the target network
Data : <i>lbl</i> the label to assign to the edges
Result : Connecting edge set C
1 $C \leftarrow \{\}$; /* creates an empty edge set */
<pre>/* for all possible source vertices */</pre>
2 forall $s \in V_N$ do
<pre>/*searches the appropriated target */</pre>
$3 \mathbf{forall} \ t \in V_M \ \mathbf{do}$
<pre>/*finding the appropriated target */</pre>
4 if Equiv(s) = Equiv(t) then
/*connects both */
5 $C \leftarrow C \cup \{ \text{Connect}(s, t, lbl) \};$
6 end
7 end
8 end
9 return C ;

Algorithm A.4: Vertices(•)
Data : G the input graph
Data : <i>lbl</i> the corresponding label
Result : A vertex set V
1 $V \leftarrow \{\}$; /* creates an empty vertex set */
2 forall $v \in V_G$ do
<pre>/* gets all labels reaching the node */</pre>
3 $L \leftarrow \texttt{Labels}(\texttt{In}(v)) \cup \texttt{Labels}(\texttt{Out}(v));$
/* just nodes that have edges with the given label lbl are worth of
creation */
4 if $lbl \in L$ then
5 $V \leftarrow V \cup \{ Copy(v) \};$
6 end
7 end
s return V :

Algorithm A.5: Edg	ges(●))
--------------------	--------	---

```
Data: G the input graph
   Data: lbl the corresponding label
   Result: A edge set E
 1 E \leftarrow \{\}; /* creates an empty edge set */
   /* for all possible source vertices s ... */
 2 forall s \in V_N do
       /* ...and all possible target vertices t ... */
       forall t \in V_N do
 3
           F \leftarrow \text{Out}(\text{Equiv}(s)) \cap \text{In}(\text{Equiv}(t));
 4
           /* ...and all edges between the given pair (s,t) ... */
          forall e \in F do
 \mathbf{5}
              if lbl \in Labels(e) then
 6
                /* ...copies the edges */
               E \leftarrow E \cup \{ Copy(e, s, t, lbl) \};
 7
 8
              \mathbf{end}
          \mathbf{end}
 9
       \mathbf{end}
10
11 end
12 return E;
```



Figure A.2: Labelled graph



Figure A.3: Network view



Figure A.4: Super-network in "aesthetic" view



Figure A.5: One column "real" view from connecting edges



Figure A.6: A possible path in the super-network

- A grammar must be provided for path validation, which is based on the alphabet \mathcal{L} .
- The final label sequence must be a valid word according to the given grammar. This means that the sequence of labels (built by concatenating the labels from each edge in the path) must be a valid word according to the provided grammar.

A.3 Context-Free Grammars

In this section the use of Context-Free Grammars [HMU01] (CFG)¹ for the super-network navigation and path validation will be discussed. The use of grammars for navigating in graphs was already proposed by Nedev [Ned99], where CFGs were used for solving the *k* disjoint paths problem for directed planar graphs – here, nevertheless, another approach is proposed. The parsing process of a "label" word (hereafter called just word) demands three elements: the word itself, a grammar, and a parser. The parser takes the rules encoded in the grammar and evaluates if the word belongs to (can be generated by) the grammar. In case the grammar can be used to generate the word then the word is said to be accepted/valid. For the approach presented here only CFGs are allowed and the Earley parser [Ear68, Ear70] is adopted – actually an optimised version [AH02].

A CFG is defined as $R(V, \Sigma, P, \delta)^2$ where V is the finite non-terminal (or variable) set, Σ the finite terminal set (or alphabet), P the finite production rule set, and $\delta \in V$ the start symbol. Another property is that $V \cap \Sigma = \{\}$, i.e. no symbol is shared between Σ and V. Specific to the CFGs, the rules in P are in the form: $[V \to (V \cup \Sigma)^*]$. This means that in the left-side of the production rule must have exactly one non-terminal symbol and in the right-side none or any sequence of terminals and/or non-terminals.

The advantage of a CFG is in the trade-off between expressiveness and parser efficiency. A CFG is a powerful formalism and allows among other things to build syntax parsers for computer programming languages, such as C [KR88], C++ [Str00], or Java.³ It has nevertheless restrictions as well. A CFG cannot restrict the total amount of a specific terminal – in traffic such rule would be something like: "no route is allowed to have more than 10 walk segments". Even though this particular case could be implemented in the Earley parser, without complexity increase. In general, any rule that requires a back-tracking, i.e. to go back and re-analyse the terminal sequence, cannot be implemented using a CFG. This implies that a CFG is not a general purpose constraint expressing formalism. It means that if a CFG is not enough for expressing the wanted constrains some post-generation filtering will be necessary.

A.3.1 Context-Free Grammars: Earley Parser

The parser chosen for the task is the Earley parser [Ear68, Ear70, AH02] that is a top-down parser, i.e. it evaluates prefixes and builds its evaluation (acceptance or rejection) based on the read terminal and the previous read word (previous terminal sequence). A counter-example would be the Cocke-Younger-Kasami [CS70, Kas65, You67] (CYK) algorithm which is bottom-up (needs the whole word before evaluating it) and also needs the grammar to be in the Chomsky normal form [Sip05, HMU01] (CNF) (CFGs can be automatic transformed to conform the CNF). The advantages of the Earley are overwhelming to even consider any other parser. First, it can parse any CFG in any form (just the original parser, but even for the optimal the modifications are minimal in comparison to CNF requirements). Second, its complexity for any CFG is $O(n^3)$, for non-ambiguous CFGs is $O(n^2)$, and O(n) for most LR(K) grammars – the CYK algorithm has complexity $\Omega(n^3)$ for all sub-cases of CFGs. A derived advantage is: if the grammar is in fact

 $^{^1}$ In Chomsky [Cho56] classification system, CFGs are type 2 languages. More about grammars can be found in [Sip05, HMU01].

 $^{^{2}}$ In the praxis the letter G is used to symbolise a grammar but to avoid conflict with the input graph, also referred as G, the letter R, from rules, is used.

³http://java.sun.com/

a sub-case of a CFG the parser will perform accordingly – for more information about LR(K) grammars and other sub-cases, please refer to [ALSU06].

The Earley algorithm is a look-ahead algorithm too, it first tries to predict which terminal comes next and when it reads it, the algorithm eliminates the invalid hypothesis. For the optimised version [AH02] it is assumed that the grammar was modified to include the following production: $[\delta' \rightarrow \delta]$ and δ' included in V as well as made the start symbol. In the parser, the basic structure is the so called Earley set of dot-productions (or Earley's dot notation); it can also be called Earley state. A dot-production has the aspect in Eq. A.4, where $A \in V$ is a non-terminal; $\rho, \tau \in (V \cup \Sigma)^*$ are sequences of terminals and/or non-terminals; j the set where the dot-production was created; and • the current parsing position for the rule (in Eq. A.4 it means that ρ was already parsed and τ is still to be parsed). The notation used is: upper-case letters to represent non-terminals in V; lower-case letters for terminals in Σ ; and Greek letters for sequences of terminals, non-terminals, or a mixture of both.

$$[A \to \rho \bullet \tau, j] \tag{A.4}$$

Each Earley set has a sequence and they are identified by S_j , meaning that it refers to the *j*th set in the sequence. A word is accepted if after reading the last terminal the corresponding set has a dot-production as $[\delta' \to \delta \bullet, 0]$ (the start production has completed its parsing process). The initial set S_0 is always created before scanning any terminal and it starts by including the dot-production $[\delta' \to \bullet \delta, 0]$. After that the three parsing procedures below are executed (in order). Assuming that the input word is x and x_i corresponds to the terminal at position i and x_1 is the first terminal:

Scanner If $[A \to \rho \bullet a\tau, j]$ is in S_i and $x_{i+1} = a$, then add $[A \to \rho a \bullet \tau, j]$ to S_{i+1} .

Predictor If $[A \to \rho \bullet B\tau, j]$ is in S_i , then add $[B \to \bullet\rho, i]$ to S_{i+1} .

Completer If $[A \to \rho \bullet, j]$ is in S_i , then add $[B \to \tau A \bullet, k]$ to S_i for all dot-productions $[B \to \mu \bullet A\varphi, k]$ in S_j .

For a visual impression of this parsing process, assume the example grammar in Eq. A.5 that is modified into the grammar in Eq. A.6.

The original grammar	(A.5)
The non-terminal set	
The terminal set	
Production rules	
Start symbol E	
The modified grammar	(A.6)
The non-terminal set	
The terminal set	
Production rules	
Start symbol $S^{'}$	
	The original grammar The non-terminal set The terminal set Production rules Start symbol E The modified grammar The non-terminal set The terminal set Production rules Start symbol S'

The word to parse is x = n + n and the parsing process is shown in Tab. A.1. There, the first column refers to the dot-production's number (used just for clarification); then comes the dot-product followed by a code, where the letters refer to: S = Scanner, P = Predictor, and C = Completer procedure. The letters (indicating the procedure) are also followed by a reference, such as "P $S_0(0)$ " which means that the resulting dot-production was created by the Predictor procedure run over the first dot-production of set S_0 . As already explained, the first set S_0 is initialised with the inclusion of the first dot-production $[S' \to \bullet E, 0]$, indicated in Tab. A.1a by "Added". In this case when the parser reads the first terminal (Tab. A.1b) it can be seen that just n is a valid word and therefore it is accepted (indicated by the presence of the $[S' \to E \bullet, 0]$ dot-product in the second item from set S_1). Because the input wasn't consumed the algorithm keeps parsing until it reaches the end of word (set S_3 in Tab. A.1d) and as a result it is a valid word (this is again confirmed by the presence of the dot-production $[S' \to E \bullet, 0]$, in S_3).

Table A.1: Parsing of x = n + n

		(a) $S_0 : \bullet n +$	n				(b) $S_1: n \bullet +$	$\cdot n$	
(0)	$S^{'}$	\rightarrow	$\bullet E$, 0	Added	(0)	E	\rightarrow	$n \bullet$,0	$] S S_0(2)$
(1)	E	\rightarrow	$\bullet E + E$,0	P $S_0(0)$	(1)	$\mid S'$	\rightarrow	$E \bullet$,0	C $S_0(0)$
(2)	E	\rightarrow	ullet n	,0	P $S_0(0)$	(2)	E	\rightarrow	$E \bullet + E$,0	C $S_0(1)$
		(0	c) $S_2: n + \bullet$	n				(d	l) $S_3: n+n$	•	
(0)	E	\rightarrow	$E + \bullet E$,0	$S S_1(2)$	(0)	E	\rightarrow	$n \bullet$, 2	$S S_2(2)$
(1)	E	\rightarrow	$\bullet E + E$, 2	P $S_2(0)$	(1)	E	\rightarrow	$E + E \bullet$,0	C $S_2(0)$
(2)	E	\rightarrow	$\bullet n$, 2	P $S_2(0)$	(2)	E	\rightarrow	$E \bullet + E$,2	C $S_2(1)$
						(3)	S'	\rightarrow	$E \bullet$,0	C $S_0(0)$

A.3.2 Context-Free Grammars: Procedures For Automatic Grammar Generators

It is strongly suggested the use of grammar automatic generators. In this section the procedures for generating grammars automatically are presented. For the procedures a function that was informally presented in Sec. A.2.2 is necessary: ConnLabel(•). Additionally to this, the functions below are also necessary. There the symbol \Re represents the set of all grammars; \Re the set of all production sets; \mathfrak{X} the set of all terminal and non-terminal sets; and X a terminal or a non-terminal.

- AddProd : $\mathfrak{R} \times \mathfrak{P} \mapsto \mathfrak{R}$ This function adds a production rule (second argument) to the given grammar (first argument), from which a new derived grammar is returned. This modified grammar has the production added to P_R with all necessary non-terminal symbols that are not already present in V_R (either in the left- or right-side of the production).
- NewVar : $\mathfrak{X} \mapsto X$ With this function it is possible to generate a new symbol, i.e. a symbol that is not present in the given symbol set (argument). The objective is to use the returned symbol as a non-terminal. A typical function call would be: NewVar $(V_R \cup \Sigma_R)$.
- NewProd : $V \times X^n \mapsto P$ It creates a new production having a non-terminal (first argument) as the production's left-side and a sequence of terminals and non-terminals (second argument) as the right-side.
- Arrangements : $\mathfrak{X} \times \mathbb{N}^* \mapsto (X^n)^n$ This function returns all possible Arrangements for the input symbol set (first argument) with the given size (second argument). An example would be:

To generate a simple grammar see Algo. A.6 but because the resulting grammar is not optimised and has some ambiguities it is recommended the use of an automatic grammar simplifier [HMU01].

A	lgorithm A.6: Grammar(•)
	Data : Σ the regular terminal set
	Data : Σ' all connecting terminals
	Data : α the special terminal from start network to all others
	Data : ω the special terminal from all others networks to the start
	Data: $\mathcal{C}: (\Sigma^n)^n$ terminal combinations
	Result : The grammar <i>R</i>
1	$N \leftarrow \Sigma ~\cup~ \Sigma^{'} ~\cup~ \{lpha, \omega\} \;;$ /* all symbols (all terminals) */
2	$S \leftarrow \texttt{NewVar}(N)$; /* the start symbol */
3	$N \leftarrow N \cup \{S\};$
4	$R \leftarrow R(\{\}, \Sigma ~\cup~ \Sigma^{'} ~\cup~ \{lpha, \omega\}, \{\}, S) ~;$ /* creates a new grammar, with empty
	production set */
	/* one production rule for each combination */
5	forall $c \in \mathcal{C}$ do
6	$v \leftarrow \text{NewVar}(N);$
7	$N \leftarrow N \cup \{v\};$
8	$P \leftarrow LabelComb(R, c, v); /* see Algo. A. / */$
0	for all p could close */
9 10	$ R \leftarrow \operatorname{AddProd}(R P) $
11	end
	/* adds combination production to start symbol: $[S \rightarrow v] */$
12	$R \leftarrow \text{AddProd}(R, \text{NewProd}(S, (v)));$
13	end
14	$N \leftarrow \Sigma_R \cup V_R;$
	/* creates a new start symbol to include the connection with the start
	network */
15	$S' \leftarrow \texttt{NewVar}(N);$
16	$N \leftarrow N \cup \{S'\};$
17	$R \leftarrow \texttt{AddProd}(R, \texttt{NewProd}(S', (\alpha, S, \omega)))$; /* encapsulation: $[S' \rightarrow \alpha S \omega]$ */
	/* creates a new start symbol to conform the prerequisites of the parser */
18	$S^{''} \leftarrow \texttt{NewVar}(N) \hspace{0.1 cm} ; \hspace{0.1 cm} / * \hspace{0.1 cm} [S^{''} ightarrow S^{'}] \hspace{0.1 cm} * /$
19	$R \leftarrow \texttt{AddProd}(R, \texttt{NewProd}(S^{''}, (S^{'})));$
20	$R \leftarrow R(V_R, \Sigma_R, P_R, S^{''});$
21	return R;

Algorithm A.7: LabelComb(•)

```
Data: R the grammar
   Data: \Sigma the terminal set to combine
   Data: S the left-side symbol
   Result: A production set P
   /* creates the recursive production rules for each terminal */
 1 N \leftarrow V_R \ \cup \ \Sigma ; /* holds the already used symbols */
 2 L \leftarrow \text{RecurLbls}(\Sigma, N, S); /* see Algo. A.9 */
 3 P \leftarrow P_L;
 4 N \leftarrow N \cup N_L;
 5 M \leftarrow M_L;
 6 E \leftarrow E_L;
   /* non-terminals for all combinations */
 7 A \leftarrow \text{VarsComb}(\Sigma, N, S, E); /* see Algo. A.8 */
 8 P \leftarrow P \cup P_A;
 9 C \leftarrow C_L;
10 N \leftarrow N \cup N_A;
11 F \leftarrow F_L;
12 E \leftarrow E \cup E_A;
   /* adds the production rules to the combination non-terminals */
13 forall c \in C do
      r \leftarrow (); /* creates an empty production right-side symbol sequence */
14
       for i \rightarrow 1 to |c| - 1 do
15
          r \leftarrow r + \text{Get}(M, \text{Get}(c, i)); /* \text{ adds to the right-side the corresponding}
16
           simple recursive production */
          r \leftarrow r + \text{ConnLabel}(\text{Get}(c, i), \text{Get}(c, i+1)); /* \text{ adds the necessary}
\mathbf{17}
          transition terminal from current terminal to the next in the
         combination */
\mathbf{18}
       end
       /* the last symbol in the combination is generated by any non-terminal
           whose production starts with that terminal */
       forall e \in \text{Get}(E, \text{Get}(c, |c|)) do
19
           /* creates the production corresponding to the non-terminal associated
               with the combination */
           P \rightarrow P \cup \{ \text{NewProd}(\text{Get}(F, c), (r, e)) \};
20
       end
21
22 end
23 return P;
```

Algor	ithm A.8: VarsComb(•)
Data	a : Σ the terminal set
Data	a: N already used symbols
Data	a : S the left-side symbol
Data	a: map E from terminals to non-terminal set, whose productions start with the
	generation of terminal set as key
Res	ult : A production set P
\mathbf{Res}	ult : C combination set, with all possible terminal combinations
$\mathbf{Res}\iota$	ult : N updated
Res	ult : map F that maps terminal combinations to corresponding non-terminal
Resi	ult : map E updated
/* g	enerates the non-terminals for all possible combinations */
$1 C \leftarrow$	$\{\}$; /* creates an empty set $((\Sigma^n)^n)^*$ /
2 $F \leftarrow$	$\{\}$: /* creates an empty map (N^{Σ^N}) */
з for i	$\leftarrow 2 \text{ to } \Sigma \text{ do}$
4 fo	$\mathbf{rall} \ c \in Arrangements(\Sigma, i) \ \mathbf{do}$
5	$t \leftarrow \texttt{Get}(c, \theta);$
	/* creates a new non-terminal to hold the combination generator symbol
	*/
6	$v \leftarrow \texttt{NewVar}(N);$
7	$N \leftarrow N \cup \{v\};$
8	$C \leftarrow C \cup \{c\};$
	/* adds v as a non-terminal that generates a production that starts
	generating the terminal $t */$
9	Put $(E, t, \text{Get}(E, t) \cup \{v\});$
	/* associates the non-terminal v with the combination c */
10	Put(F, v, c);
	/* adds a production to generates the specific combination */
11	$P \leftarrow P \cup \{ \text{NewProd}(S, (v)) \};$
12 e	nd
13 end	
14 retu	$\mathbf{rn} (P, C, N, F, E);$

Algorithm A.9: RecurLbls(•)
Data : Σ the terminal set
Data : N already used symbols
Data : S the left-side symbol
Result : A production set P
Result : N updated
Result : map M from terminal to non-terminal, whose production is the simple recursion production
Besult: map E from terminals to non-terminal set, whose productions start with the
result. map D from terminals to non-terminal set, whose productions start with the
1 $P \leftarrow \{\}$; /* creates an empty production set */
2 $M \leftarrow \{\} \; ; \; /* \; ext{creates an empty map} \; (N^{\Sigma}) \; */$
3 $E \leftarrow \{\}$; /* creates an empty map ($(N^n)^\Sigma$) */
/* creates the recursive production rules for each terminal $*/$
4 forall $t \in \Sigma$ do
5 $v \leftarrow \texttt{NewVar}(N)$; /* creates a new non-terminal to hold the recursion
generator symbol */
$6 N \leftarrow N \cup \{v\};$
7 Put(M, t, v); /* associates the terminal t with the non-terminal v */
/* add v as a non-terminal that generates a production that starts
generating the terminal $t * /$
8 Put($E, t, \text{Get}(E, t) \cup \{v\}$);
9 $P \leftarrow P \cup \texttt{RecurProd}(v, t)$; /* see Algo. A.10 */
<pre>/* adds a simple production to permit single symbol strings */</pre>
10 $P \leftarrow P \cup \{ \text{NewProd}(S, (v)) \};$
11 end
12 return $(P, N, M, E);$

Algorithm A.10: RecurProd(•)

Data: N left-side non-terminal **Data:** t the terminal **Result:** A production set P 1 $P \leftarrow \{\}$; /* creates an empty production set */ 2 $P \leftarrow P \cup \{ \text{NewProd}(N, (t)) \}$; /* simple production: $N \rightarrow t$ */ 3 $P \leftarrow P \cup \{ \text{NewProd}(N, (N, t)) \}$; /* recursive production: $N \rightarrow Nt$ */ 4 return P; /* $N \rightarrow Nt|t$ */

A.3.3 Context-Free Grammars: Earley Parser Validity States

Using the Earley parser presented, two functions are built to test the validity of any state, i.e. in any state the validity of the already read word can be checked. This property is explored by the shortest-path algorithm presented next. The two functions are: IsValid : $\mathfrak{S} \mapsto \{true, false\}$, where \mathfrak{S} is the set of all Earley sets. The algorithm for this function is simple and to explain it, it is assumed that π is the validity criterion (in the example of Eq. A.6 it would be $\pi = [S' \to E \bullet, 0]$); then IsValid $(S_j) = true \equiv \pi \in S_j$, other else it returns false.

For the second function the optimised version [AH02] of Earley parser is needed. In this version another restriction is necessary: elimination of ϵ -productions (a step in the grammar optimisation routine) and elimination of ϵ completely. This means that grammars that accept empty strings are not allowed (this is easily overcome if the ϵ -production elimination checks if the empty word is accepted and then "flags" the grammar to accept empty strings, this is the approach adopted). The parser is below presented:⁴

Scanner If $[A \to \rho \bullet a\tau, j]$ is in S_i and $x_{i+1} = a$, then add $[A \to \rho a \bullet \tau, j]$ to S_{i+1} .

Completer If $[A \to \rho \bullet, j]$ is in S_i , then add $[B \to \tau A \bullet, k]$ to S_i for all dot-productions $[B \to \mu \bullet A\varphi, m]$ in S_j .

Predictor If $[A \to \rho \bullet B\tau, j]$ is in S_i , then add $[B \to \bullet \rho, i]$ to S_i .

The second function MayBeValid: $\mathfrak{S} \mapsto \{true, false\}$ is then even simpler: MayBeValid $(S) = true \equiv S \neq \{\}$. It works because the only procedure that adds items to the Earley sets is the scanner. The scanner will only add an item if it finds a dot-production that has the read terminal right before the dot position. This means that a production sequence exists (following the dot-productions backwards in the set sequence) that could generate the current sequence, so the word may still be valid. If this is not true, then no dot-production will be added to the next set (the other procedures work only on the current set).

A.3.4 Context-Free Grammars: Earley Parser Complexity

This modified Earley parser has the same complexity of the original and for each function its complexity is given. It is assumed that the word to parse is much larger than the amount of productions in the given grammar.

Scanner O(n) because it only analyses the dot-productions on the current set and since in the worst case $|S_i| = k|x|^5$ and the procedure only depends on the amount of items in the set, so its complexity is O(n).

Predictor O(n) for the same reason of the previous procedure.

Completer $O(n^2)$ because it is the only procedure that does a backtracking, searching in all previous sets (contribution with O(n), since a parsing process has |x| + 1 sets) and in the worst case it may need to analyse kn productions, then its complexity is $O(n^2)$ – actually, $kn(0 + 1 + 2 + \cdots + n)$, being k the amount of productions in the grammar and n the size of the word.

This means that the complexity of the parsing process for a single terminal can be said $O(n^2)$ (in the worst case). For the proof of correctness please refer to [Ear68, Ear70, AH02].

⁴ For the proof of validity please refer to [AH02].

⁵ Notice that since the grammar production amount is fixed and assumed smaller than the word size, then in the worst case when a terminal is read all productions are added, so for each state a maximum of $|P_G|$ productions are added, which is still proportional to |x|.

A.4 Modified Dijkstra

For the path finding algorithms a shortest-path method is necessary. For this task the Dijkstra Algo. [Dij55] was chosen and the original algorithm is here modified to include the label parsing procedure explained before – the Earley parser. It is assumed that exactly one label is associated with each edge in a super-network, then the function GetLabel : $E_{\mathcal{S}} \mapsto \mathcal{L}$ returns this label. This function can be seen as GetLabel($e \in E_{\mathcal{S}}$) = GetAny(Labels(e)), where GetAny(\bullet) returns an element from a set and because in the super-network the label set associated with any edge has exactly one element it returns the correct element.

Before presenting the modified version, the standard Dijkstra algorithm (only modified to support multiple edges per vertex pair) is shown in Algo. A.11. However, to completely understand the algorithm, some structures and functions must be defined. A Fibonacci Heap [FT87] is the basic heap structure used here, which is modified to store a tuple instead of a pair. The tuple is as follow: $t : \mathbb{R}^+ \times V \times H \times E \times \{true, false\}$. Where the arguments are: weight $w \ge 0$; corresponding vertex; previous tuple entry that lead to this entry; used edge to reach the vertex; and a boolean value to flag the entry as the true minimum. The tuple can be seen as a structure with the following field names: t : (weight, vertex, previous, edge, minimun). To facilitate the manipulation, some auxiliary functions are defined as Weight(t) which returns the corresponding field on the structure and Weight(t, w) which adjust the field weight in t to w. The other functions are $\texttt{Vertex}(\bullet)$ for manipulating the vertex field; $\texttt{Edge}(\bullet)$ for the edge field manipulation; and so on.

In this algorithm the function Min(H) supposes that H is a Fibonacci Heap [FT87] and it returns the tuple with the lowest weight, assuming the values in the field *weight*. Another function that manipulates the heap is Decrease(H, t, w), which decreases the key/weight for entry t to w for the heap H. The function $EdgeWeight(\bullet) : E \to \mathbb{R}^+$ returns the corresponding weight for the edge informed as argument. The Algo. A.11 preserves the original optimal complexity $O(m+n \log(n))$. The proof will be omitted here because it can be derived from the proof for the modified version. To keep the explanations simple the operation $SaveTuple(\bullet)$ stores the corresponding tuple on the corresponding vertex. This is latter changed and the complexity impact analysed as well.

A.4.1 Modified relaxation function

The modifications to include the Earley parser are shown in Algo. A.13, which replaces the original function in Algo. A.12. The stored tuple for the Fibonacci Heap must be expanded, it must include an entry for the corresponding Earley set, letting t : (weight, vertex, previous, edge, minimun, state). Then additionally to the already existing functions the Parse : $\mathfrak{E} \times \Sigma \mapsto \mathfrak{E}$ that corresponds to parsing a single terminal (second argument) and generating the next Earley set based on the current set (first argument), where \mathfrak{E} is the set of all Earley sets.

A.4.2 Algorithmic Complexity

The difference between the relaxation functions in Algo. A.12 and A.13 is first in the weighting operation in line 2 of Algo. A.12 that expands to lines 2 through 7 in Algo. A.13. This operation corresponds to a single terminal parsing, whose complexity is $O(n^2)$ – see Sec. A.3.4. Following the edge terminal parsing operation the used edge weight is adjusted. The adjustment lets the weight unchanged if the corresponding Earley state "may be valid" and infinity otherwise, which invalidates the weight/edge. The second modification is the Earley set storage in the main **if** in Algo. A.13 (in line 11), which is not present in Algo. A.12.

The complexities of the used functions are expressed in Tab. A.2. The functions of kind $\mathtt{Set}(\bullet)$ and $\mathtt{Get}(\bullet)$ have complexity O(1) because it is just an access to an item of a structure. It is assumed that the $\forall v \in V \mid \mathtt{Out}(v) \ll |V|$ and for this reason $\mathtt{Out}(\bullet)$ has complexity O(1). The MayBeValid(\bullet) function is just a simple test: if the Earley set is empty or not, so O(1). Please refer to [FT87] for the complexity of the functions related to the heap manipulation as $\mathtt{Min}(\bullet)$, Decrease(\bullet), and tuple inclusion, whose complexity is O(1). The *n* and *m* refer to the amount Algorithm A.11: Standard Dijkstra algorithm with modified heap

```
Data: A weighted digraph G(V, E)
   Data: a source vertex o \in V
   Result: Shortest-Path spanning tree
   /* H is an initially empty Fibonacci heap */
 1 H \leftarrow \{\};
   /* initialisation: insert the initial tuple, where ? means undefined */
 2 H \leftarrow (0.0, o, ?, ?, false);
   /* resets the vertex stored tuple */
 3 forall v \in V do
 4 | SaveTuple(v, ?);
 5 end
   /* the main loop. */
 6 while H \neq \{\} do
      /* removes the entry with the lowest weight in H */
 7
      u \leftarrow \operatorname{Min}(H);
      Minimum(u, true);
 8
      v \leftarrow \texttt{Vertex}(u);
 9
10
      if Tuple(v) =? then
          SaveTuple(v, u);
11
          forall e \in Out(v) do
12
              /* relaxation, see Algo. A.12 */
             H \leftarrow \text{Relax}(H, u, \text{Target}(e), e);
13
14
          end
      \mathbf{end}
15
16 end
```

Algorithm A.12: Relax(•)

Data: *H* tuple heap Data: *t* the target vertex Data: *u* the source tuple Data: *e* the edge being evaluated that connects *u* to *t* Result: *H* updated /* Relaxation function */ 1 $m \leftarrow \text{Tuple}(t)$; 2 $w \leftarrow \text{Weight}(u) + \text{EdgeWeight}(e)$; 3 if $m \neq ? \land \text{Minimum}(m) = false \land \text{Weight}(m) > w$ then 4 | Weight(m, w); 5 | Edge(m, e);

```
6 Decrease(H, m, w);
```

```
7 else
8 | H \leftarrow (w, t, u, e, false);
```

```
9 end
```

```
10 return H;
```

Algorithm A.13: ModRelax(•)

```
Data: H tuple heap
   Data: t the target vertex
   Data: u the source tuple
   Data: e the edge being evaluated that connects u to t
   Result: H updated
   /* Modified relaxation function.
                                                 */
 1 m \leftarrow \text{Tuple}(t);
 2 s \leftarrow \text{Parse(State}(u), \text{GetLabel}(e));
 3 if MayBeValid(s) then
      w \leftarrow \text{Weight}(u) + \text{EdgeWeight}(e);
 4
 5 else
 6
    w \leftarrow \infty;
 7 end
 s if m \neq ?\land Minimum(m) = false \land Weight(m) > w then
       Weight(m, w);
 9
       Edge(m, e);
10
11
       State(m, s);
       Decrease(H, m, w);
12
13 else
    H \leftarrow (w, t, u, e, false, s);
\mathbf{14}
15 end
16 return H;
```

of vertices and edges of the graph and it is assumed that m > n. The edge data access functions $Weight(\bullet)$, $GetLabel(\bullet)$ and $Target(\bullet)$ have complexity O(1) because it is supposed to be an access to the data structure already present in the edge.

Table A.2: Function complexity table

Function	Complexity	Function	Complexity	Function	Complexity
Minimum(ullet)	O(1)	GetLabel(ullet)	O(1)	Tuple(ullet)	O(1)
Weight(ullet)	O(1)	Target(ullet)	O(1)	MayBeValid(ullet)	O(1)
$\mathtt{State}(ullet)$	O(1)	$\mathtt{Out}(ullet)$	O(1)	Parse(ullet)	$O(m^2)$
Edge(ullet)	O(1)	Weight(ullet)	O(1)	Min(ullet)	O(log(n))
$\mathtt{Vertex}(ullet)$	O(1)	$\texttt{SaveTuple}(\bullet)$	O(1)	Decrease(ullet)	O(1)

The complexity of the initialisation process in the Algo. A.11 is dominated by the tuple initialisation for all vertices (between lines 3 and 5, which is O(n)). Then for the modified relaxation function in Algo. A.13 the most time consuming operation is the Parse(\bullet) function, whose complexity is $O(m^2)$. The **forall** loop (line 12 in Algo. A.11) combined with the **while** (line 6 in Algo. A.11) will run at most over all edges, this means that the Relax(\bullet) function will be executed m times. Therefore the relaxation function has complexity $O(m^3)$. This allows the exclusion of the **forall** (line 12 in Algo. A.11) loop from the remainder complexity calculation. This simplified remaining **while** loop (line 6 in Algo. A.11) has its complexity dominated by the Min(\bullet) function, which is executed for all vertices. This means that it contributes with O(n*log(n)), then summing up all terms leaves a final complexity of $O(m^3 + n*log(n))$, where m = |E| and n = |V|.

This is however the theoretical complexity. This standard approach is not appropriated because it modifies the input data (graph and starting vertex). The modifications are related to the function $\texttt{SaveTuple}(\bullet)$, which alters the vertex structure to include the corresponding tuple. This makes this algorithm not optimal for an environment where the Dijkstra is in constant use and may have

concurrent access. Therefore it is modified to let the input data intact. This requires an additional data structure to store the tuples corresponding to the vertices and it may imply in complexity increase.

The adopted data structure is a hash table [Knu98], which has an access time ranging from O(1) to O(n). For further information as also complexity proof please refer to [Knu98]. The complexity increase, that may occur, is not in the relaxation function from Algo. A.15 (which is still dominated by the Parse(•) function); but in the main while loop (line 4 in Algo. A.14). The complexity change may occur if the hash table degrades to O(n) for insertion. This increases the simplified inner loop complexity from O(log(n)) (the complexity of Min(•)) to O(n). For the worst case the final complexity will be $O(m^3 + n^2)$. To prevent this problem the hash table is initialised with a capacity much higher than the amount of vertices. The hash table resize is supposed not necessary as well (it is made beforehand). The complexity of the implemented Dijkstra is then $O(m^3 + n * log(n))$.

```
Algorithm A.14: Implemented Dijkstra algorithm
   Data: A weighted digraph G(V, E)
   Data: a source vertex o \in V
   Data: a starting Earley state s
   Result: Shortest-Path spanning tree
   /* H is an initially empty Fibonacci heap */
 1 H \leftarrow \{\};
   /* initialisation: insert the initial tuple, where ? means undefined */
   /* notice that the initial Earley state s is included in the tuple */
 2 H \leftarrow (0.0, o, ?, ?, false, s);
   /* SPT is a hash table, which has a vertex as key and a tuple as value */
 3 SPT \leftarrow \{\};
   /* the main loop. */
 4 while H \neq \{\} do
      /* removes the entry with the lowest weight in H */
      u \leftarrow \operatorname{Min}(H);
 5
      Minimum(u, true);
 6
      v \leftarrow \operatorname{Vertex}(u);
 7
      if Tuple(v) =? then
 8
          SPT \leftarrow u;
 9
10
          forall e \in Out(v) do
              /* see Algo. A.15 */
             H \leftarrow \text{ImpRelax}(H, u, \text{Target}(e), e);
11
          end
12
      end
13
14 end
15 return SPT;
```

The complexity contribution of the relaxation function is said to be $O(m^3)$ but not completely discussed. The interpretation is that each edge, adopting the super-network, has exactly one label. On the worst case all edges are "parsed", i.e. their terminals/labels are parsed. In this situation the parsing process will have the complexity of parsing a word with size m and therefore $O(m^3)$.

What was not explained is that the final result is the table SPT associated with the graph G and source vertex o. With SPT is possible to reconstruct any path, remember that in the tuple the leading edge is stored. The validity criterion was yet left open because just the "may be valid" state is tested by the Dijkstra. The validity is implicit on the grammar and on the super-network structure. Recall that all paths must start and end on the access layer and the special edges from and to this layer have special labels. These labels are latter used to "encapsulate" the grammar

```
Algorithm A.15: ImpRelax(•)
```

```
Data: H tuple heap
   Data: t the target vertex
   Data: u the source tuple
   Data: e the edge being evaluated that connects u to t
   Result: H updated
   /* Implemented relaxation function */
  /* the entry is NOT removed from the hash table */
 1 m \leftarrow \text{Get}(SPT, t);
 2 s \leftarrow \text{Parse(State(u), GetLabel(e))};
   /* Algo. A.16 */
 3 w \leftarrow ModifiedEdgeWeight(s, u, e);
   /* Minimum(m) = false means that the entry still exists in H, i.e. it was
      not removed by Min() and the functions inside the if change the entry in
      H and SPT as well \ast/
 4 if m \neq ?\land Minimum(m) = false \land Weight(m) > w then
      Weight(m, w);
 \mathbf{5}
      Edge(m, e);
 6
 \mathbf{7}
      State(m, s);
      Decrease(H, m, w);
 8
 9 else
10 | H \leftarrow (w, t, u, e, false, s);
11 end
12 return H;
```

Algorithm A.16: ModifiedEdgeWeight(•)
Data : s resulting Earley state from parsing the edge e
Data : u the source tuple
Data : e the edge to be evaluated
Result : modified weight for the edge e
<pre>/* Modified grammar aware edge weighting function */</pre>
1 $w \leftarrow 0;$
2 if MayBeValid(s) = true then
$w \leftarrow \texttt{Weight}(u) + \texttt{EdgeWeight}(e);$
4 else
5 $w \leftarrow \infty;$
6 end
7 return w ;

start symbol. Then just valid paths reach the access layer and "may be valid" states that are also "invalid" will never be "valid" by adding the final edge (back to the access layer). This guarantees that paths searched within the access layer using the implemented Dijkstra are always grammatically correct, if they exist.

A.5 Dijkstra Based Path Generators

In the previous section a modified version of the Dijkstra algorithm was presented that incorporates the Earley parser. A shortest-path (SP) algorithm can nevertheless give only one path for a given OD pair. To overcome this shortage in paths some algorithms were proposed and among them are: k-shortest-path in [Epp98, JM03] or in [EH99, HMS03, HSB07] (k-SP); Dial in [Dia71]; Constrained k-SP in [vdZC05]; SP with several metrics in [BBAR06]; Way-finding in [NCK⁺06] or in [BS90, RW02, Hoc05], which is a SP with a different metric. All algorithms cited above are based on a SP algorithm to generate their multitude of paths. For some of them the use of the modified Dijkstra is transparent, i.e. no additional change in the algorithm in necessary. This is the case of SP with several metrics and Way-finding. For the others the algorithm changes are necessary to take into account the use of the grammar. For example the Dial algorithm needs the inverse of the grammar (it can be automatic generated) to calculate reverse shortest-paths (from a destination to an origin).

The most simple of these multi-path generators are the ones using different metrics. These are followed by the Dial, which was implemented but showed itself too time consuming for a realistic use and therefore not adopted. The algorithms based on the k-SP are even more complex and have a major drawback: they generate paths that are too similar to each other. This would require the generation of many paths and to discard some of them (implying in a post-generation filtering procedure), which increases even more the complexity.

The drawbacks of the multi-metric SP is that the amount of paths generated are restricted to the amount of metrics used. Therefore the adopted algorithm is a Link-Elimination Shortest-Path (LESP) that is simple and generates several paths on-demand. This algorithm is based on [ACaERSMM93] but here just one link is eliminated by each run.

The basic algorithm is in Algo. A.17. There, the function $\mathtt{Dijkstra}(\bullet)$ (Algo. A.14) is modified to returns only the shortest-path between o and d. Besides this, it also checks if the edge being evaluated is in W (the last argument that works as a taboo list) and if it is then it invalidates the edge ($\mathtt{EdgeWeight}(e) = \infty$). The complexity is not affected by the use of W because the inclusion has complexity ranging from O(1) to O(m) if implemented as a hash-table. Then the path generation is still dominated by the $\mathtt{Dijkstra}(\bullet)$ call. The function $\mathtt{Random}(\bullet)$ has also no impact because its complexity is at most O(m) – it selects an edge among the edges in the already calculated paths, excluding the ones in W. The last issue is the query of W inside the function $\mathtt{ModifiedEdgeWeight}(\bullet)$ (Algo. A.16).

This complexity could degenerate to O(m) for consulting W, therefore the complexity of ModifiedEdgeWeight(•) increases from O(1) to O(n). However, this is a sub-step in the Relax(•) (Algo. A.15) that is still dominated by the Parse(•) function (whose complexity is $O(m^2)$), therefore no complexity increase is made by the inclusion of these extra steps. The final complexity of the LESP algorithm is $kO(m^3 + n \log(n))$, being: k the amount of paths desired; m the amount of edges in the super-network; and n the amount of vertices.

Algorithm A.17: Link-Elimination Shortest-Path

Data: A weighted digraph G(V, E)**Data**: source vertex $o \in V$ **Data**: target vertex $d \in V$ **Data**: k the amount of paths **Result**: path set P1 $P \leftarrow \{\}$; /* path set */ 2 $W \leftarrow \{\}$; /* an edge set */ 3 forall $i \leftarrow 1$ to k do $sp \leftarrow \texttt{Dijkstra}(G, o, d, W)$; /* generates a new path */ $\mathbf{4}$ $P \leftarrow P \cup \{sp\};$ $\mathbf{5}$ /* selects a random edge among all edges in ${\cal P},$ excluding the edges present in W */ 6 $e \leftarrow \text{Random}(E_P, W);$ 7 | $W \leftarrow W \cup \{e\};$ s end 9 return P;

A.5. GENERATORS

Appendix B Extra Figures

In this appendix are included the graphics that would be too cumbersome to be included in the main text or that give extra understanding but are not closely related with the text in question.

B.1 Prospect Theory $\pi(\bullet)$ Function

In Fig. B.1 are depicted the behaviour of the $\pi(\bullet)$ function with different γ parameters, as mentioned in Sec. 8.11.



(a) Function $\pi(\bullet)$ for the used γ values (b) Fu

(b) Function $\pi(\bullet)$ for several γ values

Figure B.1: Function $\pi(\bullet)$ for different γ values

B.2 Exponential Learning Parameter α_0 Comparison

Here the comparison between two different initial exponential learning parameter α_0 is depicted. The Fig. B.2 shows the α_n evolution for 1,000,000 steps, has commented in Sec. 5.4.1.

B.3 Richard's Function With Different Parameters

It was mentioned in Sec. 5.4.1 that the parameters in the Richard's function (Eq. 5.5) are flexible. Therefore, in Fig. B.3 this function is depicted varying each of the parameters: Q, B, M, and ν .



Figure B.2: Exponential $\alpha(i)$ for a horizon of 1000000



Figure B.3: Richards' function with different parameter value