

The London School of Economics and Political Science

***Beyond Lucky: Measuring and
Modelling the Impact of 'Probability
Control' on Risky Choice***

Shweta Agarwal

*A thesis submitted to the Department of Management of the London School of
Economics for the degree of Doctor of Philosophy in Management Science.*

London, 2014

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without my prior written consent. I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of **63,607** words.

Statement of Conjoint Work

I confirm that Paper II (Chapter 3) was jointly co-authored with Dr. Gilberto Montibeller and Prof. Alec Morton; and Paper III (Chapter 4) was jointly co-authored with Dr. Gilberto Montibeller. I contributed 80% and 90% of the work for these papers, respectively. I contributed 100% of the work for Paper I (Chapter 2). The breakdown of contribution is summarised below:

	Shweta Agarwal (Candidate)	Dr. Gilberto Montibeller (1st Supervisor)	Professor. Alec Morton (2nd Supervisor)
Chapters 1 and 5	Determined the contents. Compiled and wrote the chapters.	Reviewed draft and suggested directions for improvements.	
Chapter 2 (Paper 1)	Identified and reviewed literature. Designed study and set up hypothesis. Programmed study on Qualtrics and collected data. Compiled and wrote the paper.	Reviewed study instructions and tasks. Assisted with data collection (emailed students). Reviewed arguments of the paper and suggested improvements.	Discussion to verify the rigour of experimental design.
Chapter 3 (Paper 2)	Identified and reviewed literature. Prepared initial scope of the paper and contents. Derived all the technical arguments and procedures that are discussed. Compiled and wrote the paper.	Reviewed early drafts of the paper, making suitable revisions. Helped tailor writing style to academic journal. Reviewed final draft for clarity and cogency.	Several discussions on structure of paper and flow of arguments. Suggested directions for developing the normative/technical ideas and reviewed the technical arguments presented.
Chapter 4 (Paper 3)	Reviewed relevant literature. Prepared material for discussion with the organization for whom the case study was conducted. Performed data analysis presented. Prepared framework and contents of the paper. Compiled and wrote the paper.	Facilitated meetings with the organization. Reviewed initial and intermediate drafts making suitable changes. Suggested revisions and reviewed final draft of the paper.	

Abstract

Managers frequently deal with risk by considering uncertainty as an element of the decision problem over which they can exert control — for example, lobbyists trying to exert influence over regulators or managers trying to mitigate Operational Risks related to human processes. This perspective that the probabilities of uncertain events are at times ‘mutable’ — i.e. subject to one’s influence — has an important and previously under-appreciated role in decision-making under risk. The present research, structured as a series of three papers, addresses this gap between theory and practice on the topic of ‘control’ from a descriptive, theoretical and prescriptive perspective.

The descriptive paper discusses a novel empirical test of the behavioural effect of ‘control’ on risk taking. The key finding that control does not always enhance risk taking but, instead, has a *moderating effect* on attitudes to risk, extends insights from related research. Strong preference for exerting control to eliminate uncertainty is also revealed. Affective and cognitive interpretations of the findings are offered and their correspondence with managerial attitudes to risk taking is discussed.

The theoretical paper builds on methods in Decision Analysis and Philosophy, and develops a new probability revision rule for modelling control as interventions on uncertainties. This rule is shown to dramatically alleviate the judgmental burden of analysing multiple interventions. Foundational properties for probability revision rules for interventions, similar to the coherence criterion for Bayes rule, are also constructed and a proof that the proposed rule satisfies these properties is offered.

In the prescriptive paper, a real world application of the probability revision rule is illustrated in the context of Operational Risk assessment, where several uncertainties are controllable (e.g. staff strikes). It is shown how this rule can be integrated with Operational Risk calculations to explicitly incorporate the effect of managerial mitigations on loss events, thus making a useful contribution to the field.

In summary, this research explores the concept of ‘probability control’ as a way to *manage risks* in the context of Decision Sciences. It furthers our behavioural understanding of risk attitudes to better resonate with managerial perspectives on risk taking and extends the relevance of Decision Analysis methods to corporate risk management.

Acknowledgements

I would like to thank a number of people who have supported and guided me through this thesis. I thank the AXA Research Fund for the scholarship award which has been instrumental to conducting the research presented in this thesis. I would like to thank my first supervisor, Gilberto Montibeller, for helping me apply for the AXA Research Fellowship, being encouraging when the ideas were unpolished, patiently guiding me as the ideas developed and helping me cultivate the writing style that is suitable for academic journals. I would like to thank my second supervisor, Alec Morton, for the long, intellectually stimulating discussions where I was pushed to defend my ideas rigorously so that they meet the necessary academic standards. I would like to thank Barbara Fasolo for engaging in discussions with me regarding the scope and basic nature of experimental studies.

I would also like to thank my Behavioural Finance mentors and colleagues at Barclays — Greg Davies, Peter Brooks, Emily Haisley and Daniel Egan — for reviewing some early versions of the experiment presented in this paper and for engaging in discussions that helped develop some of the initial ideas. I would also like to thank James Matheson and Ian Durbach for detailed comments on Paper 2 and Paper 1, respectively. I would like to thank the AXA Research Fund for identifying an opportunity to collaborate with their Human Resource and Operational Risk team at AXA and also extend my gratitude to members of these teams for their enthusiasm to explore the potential application of my research (Paper 3). I would like to thank Goodenough College for making me a part of its vibrant community and giving me a delightful abode in Central London.

I extend my appreciation to various administrative teams at LSE — The Management Science Group Department non-teaching staff for administrative support; Corporate Research Unit for assisting with the AXA Research Fund scholarship; and the Research Degrees Unit for the paperwork that was needed from time to time.

I am grateful to my colleagues Dimitrios, Ioannis, Nadine, Ana, Sumitra, Anastasia, Camelia, Shane, Maha, Ofer, Srini and Konstantinos for the enjoyable time in the department and engaging in interesting discussions that were relevant to my research as well as helping me rehearse for conference presentations. I thank my friends Josin Chin-Sang and Joshua Mora for reviewing early drafts of my research proposal and helping me articulate my ideas coherently.

I thank Barclays and Decision Technology for employing me during the early years of my PhD and giving me the opportunity to explore the practical applications of Decision Sciences. I would like to thank Brockwood Park School and the Management Science Group at LSE for giving me the opportunity to teach on various courses and thus develop my pedagogical skills, while also strengthening my understanding of several concepts related to Management Science.

Finally, I would like to thank my family — my mother, Suman Agarwal, my father, Parmanand Agarwal and my brother, Sahil Agarwal — for their constant encouragement and support through my PhD journey.

I would like to dedicate this thesis to my mother who, due to certain circumstances, had to abandon her doctoral studies.

She has stimulated my passion for thinking deeply about problems and fuelled my determination to find solutions.

Contents

<i>List of Tables</i>	10
<i>List of Figures</i>	12
Chapter 1 <i>An Introduction to the Notion of ‘Control’ in Decision Sciences</i>	15
1.1 Research motivation and context setting	15
1.1.1 On distinguishing uncertainty based on controllability	17
1.1.2 Addressing the notion of ‘controllable’ uncertainty	18
1.2 Previous research on the topic of ‘control’ and unresolved issues	20
1.2.1 Behavioral Decision Making	20
1.2.2 Decision Analysis	21
1.2.3 Other disciplines	23
1.3 Research focus	25
1.3.1 Research aims	25
1.3.2 Research methodology	27
1.3.3 Potential research contributions	29
1.4 Structure of this thesis	30
1.5 Appendix for Chapter 1	31
1.5.1 The practical advantages of a probability revision calculus	31
1.6 References for Chapter 1	31
PRELUDE TO CHAPTER 2	37
Chapter 2 <i>The Moderating Effect of ‘Probability Control’ on Risk Taking</i>	38
ABSTRACT	38
2.1 Introduction	39
2.2 Design of experiments	42
2.3 Effect of control on risk taking	45
2.3.1 Estimation and hypotheses	45
2.3.2 Results	47
2.3.3 Discussion	50
2.4 Perceived value of control	52
2.4.1 Estimation	52
2.4.2 Results	54
2.4.3 Discussion	55
2.5 Other analysis and findings	56
2.5.1 Demographic effects	56
2.5.2 Experimental factors	57

2.5.3	Qualitative responses	58
2.6	Conclusion	58
2.7	Appendices for Chapter 2	60
	Appendix 2-1 Parameter estimation	60
	Appendix 2-2 Accompanying tables.....	67
2.8	References for Chapter 2	69
2.9	Supplementary material for Chapter 2.....	74
	S 2-1: Other tables and figures	74
	S 2-2: Experiment invitation e-mail.....	77
	S 2-3: Sample of experiment (screenshots for one instance of gains-only gambles).....	78
	S 2-4: Raffle winner announcements.....	89
	S 2-5: Experiment closed e-announcement.....	90
	S 2-6: Procedure and instruction for the raffle session	90

PRELUDE TO CHAPTER 3 **93**

Chapter 3 *Stochastic Interventions: A New Probability Revision Rule to Model the Effect of Actions on Uncertainties* **94**

ABSTRACT	94
3.1 Introduction.....	95
3.2 Definition of an influence diagram and notations used in this paper.....	99
3.3 Modelling interventions as probability revision	101
3.3.1 The generic controller	102
3.3.2 Generalizing the generic controller.....	106
3.3.3 Generalized generic controller and imaging	109
3.4 Some foundations for probability revision rules for interventions	112
3.4.1 Basic properties for probability revision rules for interventions	114
3.4.2 GGC satisfies the properties of probability revision rules for interventions	116
3.5 Summary and directions for further research.....	118
3.6 Appendices for Chapter 3	120
Appendix 3-1: Eliciting the coefficients for the generalized generic controller using the influence wheel.....	120
Appendix 3-2: Proof for Theorem 1 (representation theorem for generalized generic controller).....	121
Appendix 3-3: The Dutch book for probability revision functions which are not Bayes conditionalization preserving.....	122
Appendix 3-4: A necessary and sufficient condition for the basic properties of probability revision functions for interventions	123
3.7 References for Chapter 3	125
3.8 Supplementary material for Chapter 3.....	127
S 3-1: A discussion of the conceptual issues with non-canonical influence diagrams	127

S 3-2: The Dutch book argument for Bayes rule	129
S 3-3: Expected utility calculations in influence diagrams	130
S 3-4: Order invariance for Bayes rule	132
S 3-5: Illustrating the elicitation of the GGC through a conversation	133
S 3-6: Bayes rule and the linearity condition	136
S 3-7: Post-intervention inference in a non-canonical influence diagram	137
PRELUDE TO CHAPTER 4	139
Chapter 4 A Quantitative Method for Measuring the Value of Operational Risk Mitigations.....	140
ABSTRACT	140
4.1 Introduction.....	141
4.2 Measuring and quantifying Operational Risk.....	143
4.2.1 Loss Distribution Approach.....	144
4.2.2 Limitations of the Loss Distribution Approach.....	144
4.3 Modelling mitigations when dealing with Operational Risks.....	145
4.3.1 Modelling interventions in Decision Analysis: The generalized generic controller	148
4.3.2 The probability revision matrix.....	150
4.3.3 The probability revision matrix for OpRisk mitigations	151
4.4 Case Study: Modelling mitigations for the ‘Staff Strike’ scenario	152
4.4.1 Problem description	153
4.4.2 Estimating the loss distribution.....	154
4.4.3 Modelling risk mitigation.....	155
4.4.4 Sensitivity analysis.....	157
4.5 Summary and conclusion.....	160
4.6 Appendices for Chapter 4	162
Appendix 4-1: Proof that the probability revision matrix satisfies the linearity property	162
Appendix 4-2: Geometric interpretation of probability revision matrix	162
Appendix 4-3: Distribution fitting functions (weighted least square method) for Lognormal and Weibull distribution.....	163
4.7 References for Chapter 4	163
4.8 Supplementary Material for Chapter 4	166
S 4-1: Screenshots from the Excel tool showing the model inputs and outputs for data	166
S 4-2: Algorithm for simulating the loss distribution	167
S 4-3: Sample Questionnaire used to understand AXAs model and risk mitigation policies (not used in the paper).....	168
Chapter 5 Concluding Remarks	173
5.1. A summary of the areas of research covered in this thesis	175

5.1.1.	Understanding the behavioral effects of control	175
5.1.2.	Modelling control in Decision Analysis	176
5.1.3.	Application of Decision Analysis to areas of risk assessment and risk management ..	177
5.2.	Main contributions of this research revisited.....	177
5.3.	Research limitations.....	180
5.4.	Directions for further research	182
5.5.	References for Chapter 5	185
Chapter 6	<i>Consolidated List of References</i>	188
	<i>ANNEXURE: Brief Background of some areas of Decision Sciences</i>	199
7.1	The basic theoretical framework for evaluating decisions.....	199
7.2	Modelling the psychology of decision making	200
7.3	Decision support tools for modelling decisions in practice	203
7.4	References for Annexure	204

List of Tables

Table 2-1: List of 15 choice tasks for the DWC condition. Choice tasks for the DWOC condition were derived from these gambles using equation (2-1).....	43
Table 2-2: Ordering of tasks (as numbered in Table 2-1) based on overall probability of gains/losses.....	46
Table 2-3: Proportion of riskier choices for each choice task. Statistically significant differences in the DWC and the DWOC conditions are indicated with *.....	48
Table 2-4: Matched groups of tasks, {2, 3, 4, 5} and {6, 7, 8, 9}, for a comparison between $p_A = 0.8$ and $p_A = 0.75$. Both 0.8 and 0.75 always feature on the less risky gamble, A_{DWC} and the other gamble, B_{DWC} , is same. The task numbers correspond to the numbers in Table 2-1.	53
Table 2-5: ‘Matched’ grouping of tasks for pairwise comparison between all probability levels of the proportion of times control is exerted.....	53
Table 2-6: Proportion of times control is exerted at various probability levels when the relative risk at the level is same. The pair in each cell compares the proportion of times control is exerted on the row probability vs. on the column probability for the task groups in Table 2-5. Statistically significant differences are marked with *.....	54
Table 2-7: Odds ratios from stepwise regression on demographic variables for risk taking in the DWC and the DWOC conditions and increase in risk taking in the DWC condition.....	57
Table 2-8: The utility function and probability weighting functions that are commonly assumed for Cumulative Prospect Theory.....	61
Table 2-9: Ordering of models based on AIC (lower AIC implies better explanatory power of a model). ‘Pr’ stands for the Prelec probability weighting function and ‘GE’ for the Linear-in-Log-odds probability weighting function.....	64
Table 2-10: Parameter estimates for the CPT and the PVoC models (with the GE probability weighting function).....	64
Table 2-11: Choice patterns examined and the number of choices in each pattern.	66
Table 2-12: Pairs of gambles for the 15 choice tasks. x and y denote the payoffs for the three payoff domains. For gains-only gambles, $x = £1500$; $y = £0$; for losses-only gambles, $x = £0$; $y = -£1500$; and for mixed gambles, $x = £1500$; $y = -£1500$	67
Table 2-13: Number of respondents excluded.	67
Table 2-14: Inconsistency rates for each of the two randomly selected questions that were repeated (for both conditions).....	67
Table 2-15: Proportion of riskier choices in each condition by demographics.....	68

Table 2-16: Difficulty (end of each section) and Clarity (end of practice section) of questions.....	68
Table 2-17: Effect of the order in which the two conditions appeared and the side on which gamble A appeared (logistic regression).	68
Table 2-18: Proportion of riskier choices corresponding to the order in which the two conditions appeared and the side on which gamble A appeared.	69
Table 2-19: Standard error and p values for Table 2-3.	74
Table 2-20: Gender effects for risk taking for each choice task.	75
Table 2-21: Choices of a sample participant selected for raffle (Draw/Modify from a bag).....	92
Table 3-1: Numerical comparison of the number of elicitations required for each approach.	112
Table 3-2: Transformations of the probability distribution in Box 3, where f_a is the probability revision function in Box 3 and $g_{s_{-d}^c}$ is Bayes conditioning once s_{-d}^c (competitor does not advertise) is observed.	115
Table 3-3: Bets designed by a bookie which show the payoff to the owner of the bet.....	116
Table 3-4: Number of elicitations required for each approach.	129
Table 3-5: Probabilities of market share conditional on the competitor's decision to advertise or not.	137
Table 3-6: Probabilities of competitor advertising, conditional on market share.	137
Table 3-7: Post-intervention joint distribution for $\{S_M, S_C\}$ depending on the direction of the arc between S_M and S_C	138

List of Figures

Figure 1-1: Hypothetical relationship between states which can feature in an influence diagram.....	31
Figure 2-1: The equivalence between playing two binary gambles and playing one corresponding three outcome ‘compound’ gamble.	43
Figure 2-2: Stimuli for the DWC and the DWOC conditions for task 10 in the gains-only experiment. The design for the other payoff domains was identical except for the winning amounts.	44
Figure 2-3: Mean proportion of riskier choices in the decision with control (DWC) and the decision without control (DWOC) conditions for all domains.....	48
Figure 2-4: Proportion of riskier choices at various probability levels. Shaded points indicate proportions that are significantly different from 0.5; triangles indicate proportion in the DWC condition is significantly different from proportion in the DWOC condition.....	49
Figure 2-5: Comparison of the probability weighting functions estimated for the DWOC and the DWC conditions for all three domains of gambles (based on parameter estimations in Table 2-10).....	65
Figure 2-6: Predicted vs. observed proportion of riskier choices for each choice task for the CPT and the PVoC models in Table 2-10.	76
Figure 3-1: Influence diagram for market share problem.	96
Figure 3-2: Notations for components of a hypothetical influence diagram (extended version of Figure 1).	100
Figure 3-3: An influence diagram (a) modified to represent the generic controller (b).....	102
Figure 3-4: Probability simplex for three atomic states s_o^M, s_p^M, s_f^M where s_f^M is the most desired state and the intervention has quality $q = 0.5$	103
Figure 3-5: Influence diagram for market share problem with multiple interventions.	104
Figure 3-6: Proportion of probability mass transfer illustrated for the generic controller applied to the market share example, where $q = 0.5$	105
Figure 3-7: Probability simplex for three atomic states s_o^M, s_p^M, s_f^M where the probabilities of both s_f^M and s_p^M are raised ($S^{M'} = \{s_p^M, s_f^M\}$) and the intervention has qualities $q_p = 0.3, q_f = 0.5$	107
Figure 3-8: Proportion of probability mass transfer illustrated for the EGC applied to the market share example, where $S^{M'} = s_p^M, s_f^M$ and $q_p = 0.3, q_f = 0.5$	107
Figure 3-9: Influence diagram in Figure 2 modified to include another direct predecessor for the market share node S^M	111
Figure 3-10: Revised influence diagram for the market share example (in the canonical form).	128

Figure 3-11: Evaluation of a hypothetical influence diagram (extended version of Figure 3-2 with sequential decisions). The shaded nodes correspond to the information state Ω when the decision A_j needs to be made. The unshaded nodes are the variables for which the expected utility is calculated (for fixed Ω_j) in each iteration.....	132
Figure 4-1: Generic influence diagram for OpRisk of a loss event.	146
Figure 4-2: Statistical difference between managing uncertainty (distribution of an OpRisk event) and managing the impact of an OpRisk event (consequences).	147
Figure 4-3: Each of the 4 panels shows the original simplex (solid line), the simplex corresponding to the generic controller (4-6i) (dotted dash line) and the simplex corresponding to equation (4-6ii) (dashed line) for various values of q . The vertices represent the atomic state probabilities $P(s_i) = 1$	152
Figure 4-4: Influence diagram for the OpRisk of staff strike (with mitigation on length of strike).	153
Figure 4-5: Distributions for frequency ($x_f \sim Pois(\lambda = 0.02)$) and length of strike $x_l \sim LogN(\mu = 1.06, \sigma = 0.74)$; $x_l \sim W(\alpha = 0.94, \beta = 1.49)$	155
Figure 4-6: The transformation of the simplex corresponding to various values of q in equation (4-7) (modified for at most three strike days).....	156
Figure 4-7: Severity (Lognormal) and Loss distribution without mitigation and with mitigation ($q = 0.5$ in equation (4-7)).	156
Figure 4-8: Percentage reductions in capital charge (99.5 percentile) for various qualities of intervention in equation (4-9). The solid line shows the percentage reduction if the quality is fixed over time (equation (4-8)). The grey bars indicate the extent to which percentage reduction can change depending on how quality changes over time.	157
Figure 4-9: Convergence tests for capital charge estimates and percentage reduction: The black line corresponds to a Lognormal distribution and grey line corresponds to a Weibull distribution for severity ($p < 0.01$ for all figures, based on two sample Kolmogorov-Smirnov test). For Lognormal distribution the mean capital charge was found to be €3.13 million (\pm €3750) without mitigation, and 1.13 million (\pm €1150) with mitigation. The mean percentage reduction was found to be 64%. (\pm .03%).	158
Figure 4-10: Sensitivity of capital charge estimates (without mitigation) to various inputs of the score.....	159
Figure 4-11: Capital charge estimates using parameters conditional on ‘negotiability’ score vs. unconditional parameters. (Based on rank sum test, $p \leq 0.05$ for Lognormal distribution, when negotiability < 0.25).	159
Figure A-1 Illustration of functions capturing risk aversion and risk seeking behaviour for a hypothetical choice between a bet [£60, 0.5; £0, 0.5] and sure amount of £30.....	201

Figure A-2: Illustration of probability weighting ($w(p)$) typically exhibited by individuals.	202
Figure A-3 Influence diagrams for different decision problems.....	204

Chapter 1

An Introduction to the Notion of 'Control' in Decision Sciences

Some people want it to happen, some wish it would happen, others make it happen.

-Michael Jordan

1.1 Research motivation and context setting

Decisions are integral to the lives of most people. By choosing between decisions and acting upon them individuals strive to bring about more desirable states of affairs (Von Mises, 1957). Often, making decisions involves some uncertainty about the consequences of the decisions. Examples of such decisions that involve uncertainty are deciding to carry an umbrella or not, investing in stocks vs. fixed income bonds, selling or buying a house now or later, decisions about performing life threatening medical surgeries, a basketball player contemplating taking a long distance shot, earthquake evacuation programmes, etc. Furthermore, some of the potential consequences of a decision can be *negative*, such as losing initial capital when investing in stocks, or the loss of lives if an earthquake strikes. Colloquially, such decisions are regarded as being *risky*, i.e. they bear 'the *chance* of something *bad* happening'. A formal definition of risk typically constitutes some combination of the severity of the negative consequences and how likely they are (e.g. Bedford and Cooke, 2001; Holton, 2004; Pate-Cornell, 1996).¹

People are willing to engage with risky decisions because, in the real world, risk is often positively correlated with the potential reward (positive consequences) (e.g. Loewenstein et al., 2001): a greater danger also presents a greater opportunity. For example, in the corporate world risk taking is believed to have a close bearing on an individual's and a firm's success (MacCrimmon and Wehrung, 1990, 1986); investing in stocks presents greater returns on investment than investing in bonds; entrepreneurs risk starting a business because they believe it has the potential to reap huge rewards; and the risks associated with

¹ It is noted that a universal definition of risk is subject to debate in academia and definitions of risk often relate to context specific quantitative assessments such as probability assignments to uncertainty (Knight, 1921), Value-at-Risk (Holton, 2002; Leavens, 1945; Markowitz, 1952; Roy, 1952), volatility, probability of loss, etc.

dangerous sports like bungee jumping are compensated by the thrill experienced. Dating back to the eighteenth century, developing methods for analysing and modelling the risk-reward trade-off remains an active area of research across a wide range of disciplines (Bernstein, 1998).

This thesis is centred on decisions under uncertainty and its focus is on a formal study of uncertainty management as a way to mitigate risks, in the context of Decision Analysis. The perspective adopted in this research is that, in many real world situations, voluntary actions serve as instruments that can interact with uncertainty to bring about the desirable consequences, or mitigate the undesirable consequences. Thus, rather than simply choosing an option that presents the optimal risk-reward tradeoff, individuals or organizations often seek to *reduce the* risk while retaining most of the potential rewards (see Huber, 2002, 2007).

Risks can be mitigated by controlling the impact of an unfavourable decision outcome to the decision maker (e.g. purchasing insurance) and/or by managing the uncertainty (e.g. taking measures to prevent the unfavourable event from occurring). Whether it is the impact or the uncertainty (or both) that can be controlled, this depends crucially on the nature of uncertainty that the decision maker experiences. For example, consider the various types of Operational Risks an organization may experience such as an earthquake, break down of IT systems, internal frauds, staff strikes, etc. (Basel Committee, 2003; Cummins and Embrechts, 2006). While the organization can put measures in place to control the impact of natural disasters, such as an earthquake (e.g. purchase insurance, set up alternative temporary work stations), it can do nothing to make the occurrence of an earthquake less likely. In contrast, for events corresponding to human processes, such as staff strikes, the organization can not only control the impact but also put measures in place to prevent strikes from occurring (e.g. periodically review employee satisfaction) or mitigate their severity (e.g. establish good relations with trade unions to terminate the strike quickly). Similarly, consider the decision of buying a small stake in a blue chip company (with no control over board decisions) vs. buying a large stake in a start-up (with control over board decisions). In the former case, the investor cannot influence the probability of a stock price increasing but in the latter case, the investor may be able to exert some control on improving the valuation of the company.

These examples motivate a qualitative distinction between uncertainties that are ‘controllable’ and those that are ‘uncontrollable’ based on whether or not a decision maker can influence the probabilities attached to uncertain events. Such a distinction has been proposed previously from the perspective of whether the probabilities depend on skill or chance (Brandstätter and Schwarzenberger, 2001; Cohen, 1960; Goodie, 2003). In this

research, a more general distinction is made based on the *mutability of the probability distribution* of an uncertain event. This classification of uncertainty — based on controllability — is the key theme of this research and overarches the collection of three papers (Chapters 2 – 4) that constitute this thesis. The focus of all three papers that compose the thesis is on the study of decision making when the uncertainty is ‘controllable’ and this topic is explored from both behavioural and theoretical perspectives.

1.1.1 On distinguishing uncertainty based on controllability

A dichotomous conceptualization of uncertainty, based on qualitative factors, is not new to the analysis of uncertainty. A widely discussed distinction of uncertainty is based on its source, where uncertainty is classified as arising due to lack of knowledge (epistemic) or due to random variation that cannot be explained (aleatoric — e.g. error in prediction, spinning of the roulette wheel) (Fox and Ulkumen, 2011; Chatterton, 2001; Gilboa, 1999; Pate-Cornell, 1996; Hoffman and Hammonds, 1994; Chernoff and Moses, 1959; Hacking, 1975; Keynes, 1936). Another, more controversial, distinction of uncertainty is based on its quantification and uncertainties are distinguished in terms of class probability (or frequency probability) and case probability (or subjective probabilities that represent agents’ degrees of belief) (Morgan and Henrion, 1990; Von Mises, 1949, pp. 107–115). Although the foundational basis for both these classifications remains controversial, they are useful for the purpose of modelling and understanding uncertainty in complex decision problems. From a risk management perspective, the distinctions help recognize whether the uncertainty is reducible or not (see Winkler, 1996). From a behavioural perspective, there is evidence that the nature and the sources of uncertainty affect how people respond to uncertainty, and consequently, people’s decisions and preferences. This is illustrated by the famous Ellsberg Paradox (Ellsberg, 1961; see also Chow and Sarin, 2002).² The qualitative distinction of uncertainty also underlies cognitive factors that inform probability judgements (Howell and Burnett, 1978), biases in probability judgements (e.g. hindsight bias is more salient in situations where uncertainty is epistemic (see Fox and Ulkumen, 2011)) and investment behaviours (Ulkumen et al., 2014).

² The problem consists of one urn containing 30 red balls and 60 balls which may be blue or yellow. Subjects are presented with two decisions. In the first decision they choose between Bet A where the payoff is £100 if a red ball is drawn (£0 otherwise) and Bet B where the payoff is £100 if a blue ball is drawn. In the second decision the choice is between Bet C where the payoff is £100 if a red ball or yellow ball is drawn and Bet D where the payoff is £100 if a blue or yellow ball is drawn. Individuals choose Bet A and Bet D, contrary to what principles of rationality dictate, suggesting that individuals prefer to bet on known probabilities and have a distaste for ambiguity (unknown probabilities).

On a similar score, a conceptualization of uncertainty based on controllability, as the one proposed here, is also germane to the study and analysis of uncertainty but has not received much attention. From a risk management perspective, this conceptualisation helps distinguish “luck” from “knowing” from “managing”— an agent may not merely wish to cope with or manage risks that result from the model of the real world, but may instead act towards modifying the world so as to guarantee that favourable outcomes will ensue. With respect to the analysis of decisions, this concept of control recognizes that agents are not merely prognosticators of events, but in some situations can cause the desired outcomes to happen.

This perspective of causation and control is at the heart of how managers think about decisions under uncertainty and how they decide to take risk. The alliance of ‘control’ with ‘risk’ is salient in managerial conceptions of risk; they make a sharp distinction between gambling (where the odds are exogenously determined and uncontrollable) and risk taking (where skill or information can reduce the uncertainty) (March and Shapira, 1987). A manager is considered to be taking good risks when s/he can control fate and is confident of success as opposed to gambling when s/he relies on fate. Managers also look for alternatives that can be managed to meet targets, rather than assess or accept risks (March and Shapira, 1987; Strickland et al., 1966). Courtney (2000, p. 40) makes the observation that the size of strategic bets have the potential to shift the industry structure, thus alluding towards the notion that uncertainty is perceived by managers as being ‘shapeable’. The notion of controllability of uncertainty is also relevant when understanding the willingness of entrepreneurs to bear risks (Caliendo et al., 2009).

Outside the corporate world, there is evidence that perceived controllability has a positive impact on risk taking (e.g. Cohen, 1960; Chau and Phillips, 1995; Horswill and McKenna, 1999; Brandstätter and Schwarzenberger, 2001; Goodie and Young, 2007). For example, risk acceptance is higher when the probabilities associated with the uncertainty depend on skill, and hence are ‘controllable’ (Goodie, 2003; Goodie and Young, 2007); individuals find bets on skill-related probabilities more attractive than chance-related probabilities (Brandstätter and Schwarzenberger, 2001; Young et al., 2011) and driving speeds are higher when individuals imagine themselves as drivers rather than passengers (Horswill and McKenna, 1999).

1.1.2 Addressing the notion of ‘controllable’ uncertainty

It emerges that whether or not the uncertainty is controllable has a bearing on how people make decisions when faced with risk and uncertainty. When uncertainty is controllable, the choice of action is not viewed as a resolution of the decision problem at hand but as a way to

alter the problem or environment. Such acts, which control or modify the uncertainty, can be considered to be *interventions* — “active” interactions with the system to affect the causal factors that influence the uncertain states. Furthermore, since uncertainty generates much cognitive strain, being able to mitigate uncertainty in and of itself may have value (Humphreys and Berkeley, 1985). From a Decision Analysis perspective, this means that discussions of how uncertainty can be handled ought to be incorporated explicitly in the formal analysis of decisions (Humphreys and Berkeley, 1985).

In recent times, the pressing need for more research in Decision Sciences where the role of ‘control’ is formally acknowledged in the study and analysis of risky decisions has been reflected in discussions by scholars. Shatcher (2012), for instance, advocates that modelling the effect of actions on uncertainties can facilitate analysis of strategic decisions such as understanding decision opportunities or developing new and robust strategies. Rosenzweig (2013) emphasizes that identifying controllability of outcomes is essential to how individuals approach many decisions, frequently endeavouring to “make things happen”. von Winterfeldt et al. (2012) demonstrate how being in control of decisions or not can affect the value of information analysis of uncertain variables. While the field of Decision Sciences has a wealth of models, tools and techniques that study, analyse, explain and guide real world decision making, the notion of control has been under appreciated in its methods. By incorporating control as an explicit component of decision models, methods from Decision Sciences can be applied more usefully to analyse real world decisions, especially in the corporate world.

This research looks at the analysis of decisions where agents can exert *control* over the uncertain events on which the outcomes depend, i.e. the uncertainty is ‘controllable’. The papers in this thesis constitute a descriptive study of how controllability of uncertainty is an important determinant of risk taking behaviour, develop decision analytic techniques to explicitly model the impact of decision makers’ actions on uncertainties and discuss an application of the proposed techniques to analyse real world risk assessments and mitigations of Operational Risks. Developments in this direction can, hopefully, help better tie the managerial reasoning of risky decisions & corporate risk management to the methods of Decision Sciences.

The rest of this chapter is organised as follows. In the next section, previous research on the topic of control is reviewed and some areas for further research are identified. In the subsequent section the research aims are established within the context of existing research, a brief overview of the method and approach used are presented and the potential contributions of this research are highlighted. The final section discusses the structure of the

remaining thesis. The appendix covers some detailed expositions of the research gaps that are identified.

1.2 Previous research on the topic of 'control' and unresolved issues

The notion of control has received attention in a number of disciplines that study decision making such as Behavioural Decision Making, Decision Analysis, Philosophy and Artificial Intelligence. In this research, the focus will be limited to the treatment of control in Behavioural Decision Making and Decision Analysis and only passing remarks are made on associations with other disciplines. This section briefly reviews some of the discussions in these areas of study and exposes some of the issues that arise in Decision Sciences when the notion of 'controllable uncertainty' and 'control' is explicitly introduced into the formal analysis of decisions. Research areas that can merit from further study and are covered in this thesis are also identified.

1.2.1 Behavioral Decision Making

In the field of Behavioural Decision Making, the topic of control has been studied from a number of different perspectives. Early notions of control closely correspond to skill related probabilities and studies have shown that contrary to the thesis of ambiguity aversion, agents prefer to bet on skill (imprecise probabilities) rather than on chance (precise probabilities, e.g. ascertained from a roulette wheel) (Cohen, 1960, p. 85; Cohen and Hansel, 1959; Howell, 1971). This behaviour can in part be explained by overconfidence in one's own performance (Howell, 1971) and favourable perceptions of one's competence (Heath and Tversky, 1991). Langer (1975) studied the effect of 'illusionary control' — when pertinence of skill in the task is mistaken, arising from situational factors such as familiarity, involvement, foreknowledge, reinforcement — and found a positive effect on the amount subjects were willing to bet or options they chose. Following this seminal study, a number of experiments have been conducted on the variants of illusion of control (see Presson and Benassi (1996) for a review), all confirming a positive effect of illusion of control on the attractiveness of options. A recent test of this choice-judgement discrepancy³ (Li, 2011) revealed that in some cases, preference for an option stems from a *source preference* (Chew and Sagi, 2008), rather than *illusion*, since subjects recognize that they cannot modify the probabilities of the task at hand. Goodie (2003) revisited the relation between control and

³ The choice-judgement discrepancy (Heath and Tversky, 1991) is the preference to bet on event A over B even when B is judged to be at least as likely as A.

skill and introduced a formal definition of control as 'probability alterability' which corresponds closely to the concept of control in this thesis. Subsequent studies verified that controllability of uncertainty influences risk acceptance on skill related probabilities independently of competence in tasks (Goodie and Young, 2007) and this attractiveness for bets on controllable probabilities can be captured by an elevated probability weighting function (Young et al., 2011). Huber (2002) argues that control and managing the risks is an essential ingredient of every day decision making which is absent in most laboratory studies on decision making. His experiments indicate that preference for alternatives are driven by the presence of an opportunity to mitigate risk either before or after the occurrence of a negative event (Huber, 1997, 2007).

It appears that the topic of control has been studied from a variety of behavioural perspectives and there is sufficient evidence that control is an important determinant of risk taking. In most studies, however, the manipulation of control has revolved around perceptions of control and invariably been confounded with other factors such as competence or overconfidence. A more 'pure' test of the effect of control on choices and risk taking would be desirable. Furthermore, while endeavours have been made to explain the effect of control on risk acceptance using popular theories of decision making, such as Cumulative Prospect Theory (Tversky and Kahneman, 1992), there is scope for more research in this direction. Specifically, by studying the effect of control over a wider range of probability levels and alternative decision frames, especially payoff domains which entail losses only (for example, as is the case for managers dealing with Operational Risks), its effect on the commonly observed *fourfold pattern of risk attitudes*⁴ can be captured. This would lead to a more comprehensive understanding of the behavioural effect of control on risk taking. Both these directions will be explored in one part of this research.

1.2.2 Decision Analysis

The concept of control is implicit in decision support tools such as influence diagrams (represented by an arc from a decision node to a state node) (Howard and Matheson, 1984, 2005) and decision trees (represented by asymmetric decision trees). Decisions that affect probabilities of a state are conceptualized as 'interventions' on uncertainty (e.g. Matheson and Matheson, 2005). In contrast to the analysis of information which has been explored extensively, there have been very limited discussions on the formal treatment of

⁴ The fourfold pattern of risk taking, predicted by Cumulative Prospect Theory, corresponds to risk aversion for moderate to large probability of gains, risk seeking for small probabilities of gains, risk seeking for moderate to large probabilities of losses and risk aversion for small probabilities of losses.

interventions. In the case of information, Bayes rule⁵ is well established as a probability revision rule for modelling the effect of information on probability distributions. A probability revision calculus simplifies the input information required when dealing with complex decisions, described by large graphical networks (e.g. Pearl, 2000, p. 14). It enables the analyst to obtain probability inputs independently for different uncertain states from different experts, exploit the mathematical relationship between uncertain states to reduce the judgemental burden of obtaining the inputs, as well as update the distributions when one or more variables are eliminated or introduced in the graph (illustrated in Appendix section 1.5.1 for Bayes rule).

When probabilities depend on decisions, however, Bayes rule cannot be used to modify the probabilities as it would require probabilities to be assigned to actions, which is unnatural. Typically, several probability distributions conditional on each decision are elicited. As will be demonstrated in section 3.3.3, this approach can be judgementally very burdensome especially when several interventions need to be explored and the controllable state depends on other state nodes. It is for this reason, perhaps, that the majority of research in Decision Analysis and related disciplines has focused on analysing interventions that bring about a state outcome with certainty. A treatment of control as a perfect intervention on a state node, analogous to the concept of Expected Value of Perfect Information, was first proposed by Matheson (1990) using influence diagrams. A preliminary extension of the concept of perfect control to imperfect control is offered by Matheson and Matheson (2005) within the context of influence diagrams. Their approach called the 'generic controller', describes a Bayesian procedure to obtain the distribution conditional on an intervention as a probability revision of a pre-intervention distribution (distribution corresponding to the do-nothing option). A detailed discussion of the generic controller is offered in section 3.3.1.

While the procedural simplicity of the generic controller is appealing, as will be discussed later in this thesis, it is limited in the type of interventions it can model. Nonetheless, it is an interesting step in the direction of developing a probability revision calculus for describing the effect of actions on probability distributions and mirrors the well-established use of Bayes rule as a method for modelling the effect of information on probabilities. Advancing the research on modelling the dependence of probabilities on actions in Decision Analysis as probability revision rules, can enhance its use in practice,

⁵ According to Bayes rule, the probabilities of a state $S^A = \{s_i^A\}$, given information about another state $s_k^B \in S^B = \{s_i^B\}$, is computed using the equation $P(s_i^A | s_k^B) = P(s_i^A, s_k^B) / P(s_k^B)$.

especially in areas of risk management. Developing methods to model interventions and their impact on probability distributions is therefore one of the key areas of focus in this research.

It is worth noting that one of the reasons that the field of Decision Sciences has neglected the study of decision problems where probabilities can depend on decisions is because introducing this dependence conflicts with the underlying axioms that support Expected Utility Theory. While decision theories which allow this dependence have been proposed and debated (for example variants of Conditional Expected Utility Theory (Jeffrey, 1965; Bolker, 1967; Luce and Krantz, 1971) and Causal Decision Theory (Gibbard and Harper, 1978; Lewis, 1981; Skyrms, 1982; Armendt, 1986; Joyce, 1999)), their use has been somewhat limited and a marriage of these theories with non-canonical influence diagrams remains to be exposed. A detailed discussion of this issue is beyond the scope of this thesis.

1.2.3 Other disciplines

In other disciplines, the notion of control is embedded in the concept of *causality*. Pearl (2000) has produced pioneering work on causality in the field of Artificial Intelligence. In Pearl’s framework the relationship between variables described by a graph, such as Bayes networks, acquire a functional characterization and the outcome of a state S is determined by some function $f(\cdot)$. An intervention on a state S is represented by replacing the function $f(\cdot)$ with the fixed value that represents the state brought about by the intervention and the graph is modified to eliminate all arcs into the state node intervened on. These ‘atomic’ interventions are represented by mechanism modifying operators $do(S)$ and the probability distribution of the revised graph is updated using a ‘truncated’ Bayes factorization formula (Pearl, 2000, chap. 3). Heckerman and Shachter (1995; Shachter and Heckerman, 2010) link Pearl’s model and the $do(\cdot)$ operator with decision theoretic primitives in an influence diagram — acts, states and consequences. They define causality based on the notion of ‘unresponsiveness’ of one variable to another variable in an influence diagram and analyse the dependence of uncertainties on actions using ‘mapping variables’.

One limitation of the treatment of interventions offered in the field of Artificial Intelligence is that it treats interventions as being deterministic, i.e. the intervention eliminates uncertainty entirely from the state node. It is not immediately clear how more realistic interventions, which modify the probability distribution of a specific state variable, can be represented. As suggested in section 1.2.2, developing a more general probability calculus for interventions, which is not restricted to atomic or ‘perfect’ interventions, can be useful for modelling real world decisions. Addressing this need will be a primary intent of the methods for modelling interventions that are developed in this research.

Some of the interesting characteristics of probability modifications as a result of intervention vs. obtaining new information are highlighted in decision situations when the ‘controllable’ uncertain state also depends on other uncertainties. In particular, the effect of actions on probabilities cannot be treated like the effect of information on probabilities (Pearl, 2000, chap. 4); treating decisions as evidence can lead to erroneous inferences such as “patients should avoid going to the doctor to reduce the probability that one is seriously ill” or “students should not prepare for exams lest this would prove them behind in their studies” (see Pearl, 2000, pp. 108, 242). Such paradoxical decision cases, exemplified by the Newcomb paradox⁶ (Nozick, 1969), have been discussed extensively in areas of Philosophy, leading to two variants of Expected Utility Theory — Evidential Utility Theory and Causal Decision Theory. The two theories differ in their interpretation of probabilities conditional on decisions: Evidential Decision Theory treats the conditional probabilities as ordinary Bayes conditionals ($s_i \in S$ given a is performed’) (thereby spuriously adjusting probabilities for any correlation between an action and a state, even when the probabilities of the state do not depend on the action), whereas in Causal Decision Theory the conditionals are treated as counterfactuals ($s_i \in S$ if a were performed’) and describe causal influences of actions on uncertainty. The latter treatment of conditionals is considered to be a more suitable interpretation for actions (see Pearl, 2000, p. 112). A probability transformation for dealing with counterfactuals, called *imaging*, was described by Lewis (1976) and was later generalized by Gardenfors (1982). A simplified review of imaging and its comparison with the $do(\cdot)$ operator is offered by Pearl (2010). More details on *imaging* are discussed in section in section 3.3.3.

While there appears to be some work exposing the relationship of the $do(\cdot)$ operator with the concept of control in Decision Analysis (Shachter and Heckerman, 2010) as well as with imaging (Pearl, 2010), the relationship between Decision Analysis techniques to model control and imaging does not appear to have been explored before. In this research, potentially interesting relationships between Decision Analysis methods (e.g. generic controller) and imaging will be investigated.

⁶ In this fictitious paradox, there is a being who can predict the future choices of any individual with full confidence and is known to have never erred. Suppose there are two boxes B_1 and B_2 . B_1 contains \$1000 and B_2 may contain \$0 or \$1,000,000 (\$ M). An individual is presented with the choice of taking both the boxes or only B_2 . The individual knows that if the being predicts he will take both boxes, then the being will not put \$M in B_2 . If, instead, the being predicts that the individual will take only B_2 , then he will put \$M in B_2 . The sequence of events is as follows: The being first predicts the choice the individual will make and then puts the respective amount of money in B_2 . Then the individual makes a choice. The question is: what is the optimal choice for the individual?

It is worth observing that, perhaps, one reason why no reference to imaging appears in the Decision Analysis literature is that its supporting arguments are based on topics such as causality and counterfactual reasoning, which Decision Analysis has stayed clear of (see Shachter, 2012). In comparison, the widely used belief revision rule for information — Bayes rule — is grounded in rationality arguments like Dutch books (Teller, 1973) which mesh more naturally with the Decision Analysis approach. Perhaps developing similar coherence arguments for imaging, or any probability revision rule for interventions in general, can make their use more compelling and relevant in Decision Analysis. Establishing some foundational properties for probability revision rules for interventions will therefore also be an area of focus in this research.

Finally, it is noted that another area that also deals with the concept of control is *Markov Decision Processes* (Howard, 1960). Control is modelled as a state transition function which describes the probabilities that an action will move the system from one state to another. Much of the discussion of Markov Decision Processes is centred on dynamic processes where the *same* state acquires different values and outcomes of the state are either partially or completely observable at the time of a decision. While it is recognized that there could be potential synergies between Markov Decision Processes and modelling techniques in Decision Analysis, the literature on Markov Decision Processes is not covered in this research. The next section outlines the topics on control that are studied in this research.

1.3 Research focus

The previous subsection exposed a number of areas of research in Decision Sciences where the topic of control can merit from further study. Some of these directions of research will be explored in this thesis and others are left to future research. The research directions explored in this thesis relate to expanding the prevailing behavioural understanding of the effect of control on risk taking as well as further developing the techniques for modelling interventions in Decision Analysis. Potential real world applications of the methods developed will also be explored. While an attempt will be made to link the methods to model control in Decision Analysis with belief revision rules in Philosophy, their correspondence with other disciplines, such as Markov Decision Processes and Artificial Intelligence, will not be covered here. The issue related to the normative foundations of decisions where probabilities depend on actions is also left to future research.

1.3.1 Research aims

The overarching goal of this research is to further develop the formal study of decisions in contexts which present the opportunity to modify uncertainties so that they mesh more

naturally with managerial perspectives of uncertainty and how it can be managed. It seeks to better understand the potential effects of control over uncertainty on risk taking and endeavours to expand the quantitative tools for modelling the effect of control on probability distributions, while also exploring their applications in a corporate context. In relation, this research investigates four main research questions and two secondary questions. The first two of the six research aims are formulated in the area of behavioural decision making. Research aims 3-5 relate to theoretical methods of Decision Analysis. The final aim acquires a prescriptive flavour and concerns the practical value of the formal methods for analysing control that will be developed in the research. The six research aims are:

1. In the context of behavioural decision making, a primary aim is to provide a more comprehensive understanding of the behavioural effect of control on risk taking. Specifically, an empirical test, which covers a wider range of probabilities and payoffs (than explored by extant studies) and also endows individuals with actual control over probabilities, will be designed and implemented. The purpose of such a study is to test the effect of control on the fourfold pattern of attitudes to risk that is predicted by state-of-art theories of decision making such as Cumulative Prospect Theory (Chapter 2).
2. A secondary aim of the behavioural study of control, also related to one of the theses of Cumulative Prospect Theory (certainty effect), is to examine if preferences for exerting control relate to prior probabilities of success and if exerting control is valued differently based on whether it only modifies probabilities or also eliminates uncertainty (Chapter 2).
3. From an analytic perspective, a primary aim is to expand the existing procedures of Decision Analysis to model the effect of control on probability assessments in a tractable and more general way. Specifically, the goal is to develop a procedure that can capture beliefs about interventions that are not restricted to those that bring about any one outcome of an uncertain state with certainty or improve the probability of only one state (Chapter 3).
4. Another primary aim of the analytic research is defending the suitability of existing and proposed probability revision procedures for interventions by establishing theoretical foundations for them which are similar to the coherence criterion that supports Bayes rule (Chapter 3).
5. A secondary aim of the theoretical research on control is exploring any connections between modelling techniques for control in Decision Analysis and probability revision rules in Philosophy (e.g. imaging) (Chapter 3).

6. The final primary aim of the research, in the context of practical application of the present research, is to apply the procedure developed in this thesis in a real world context, to assess its feasibility and potential usefulness (Chapter 4).

Since these research aims span a number of different sub disciplines of Decision Sciences, the methods of research employed are also distinct. The next subsection details the approaches used to address these questions.

1.3.2 Research methodology

The study of decision making has been broadly distributed between three different approaches (Bell et al., 1988; Tsoukiàs, 2008) — 1) the *normative* approach, which seeks to construct theories of how people *should* make decisions; 2) the *descriptive* approach, which offers formal explanations of how people *do* make decisions; and 3) the *prescriptive* approach, which integrates the normative and descriptive insights to devise techniques to help people make better decisions. The three papers in this research, on the role of control in the study of decision making, can be roughly distinguished based on these three approaches and correspond to different aspects of extant research on control.

The first paper (Chapter 2) is firmly rooted in the descriptive study of control and endeavours to understand how a key aspect of decision making — attitudes to risk — are different with and without control. It takes the fourfold pattern of risk attitudes described by Cumulative Prospect Theory as a starting point, which is well-established especially for decisions where subjects cannot exert control on the probabilities (Tversky and Kahneman, 1992). Through a randomized controlled within–subject study, conducted using approximately 300 human subjects, the research in this paper exposes if and how the fourfold pattern of risk attitudes is altered when individuals have an opportunity to exert control on the probabilities. The data gathered is also analysed statistically from a different perspective to check if preference for exerting control is driven by prior probabilities of ‘success’ and for the prevalence of the well-known certainty effect (Tversky and Kahneman, 1986).

The second paper (Chapter 3) has a theoretical focus and the latter part of this paper acquires a normative flavour. The first part of the paper is aimed towards building a general probability revision rule for modelling the effect of actions that extends and integrates existing approaches for modelling probability revisions described by interventions. Taking Matheson and Matheson’s (2005) generic controller as a starting point, modifications to their procedure are introduced incrementally so that a wider class of interventions can be modelled. Once a sufficient generalization of the generic controller is obtained, the equivalence of the *generalized generic controller* and probability revision rules proposed in

other disciplines (specifically ‘imaging’) are formally established. In the second part of the paper, the proposed probability revision rule is defended on normative grounds using Dutch book type arguments, similar to the ones that support Bayes rule as the ‘rational’ procedure for processing information and updating beliefs. In particular, for special instances of interventions — interventions on state nodes that depend on other state nodes — a Dutch book argument is presented for the generalized generic controller and some other desirable properties are also exposed. The discussions presented in this paper are illustrated using a fictitious example of a firm contemplating advertising to increase the probability of a higher market share.

The third paper (Chapter 4) seeks to test the practical value of the method developed in Paper 2 and demonstrates how using the generalized generic controller can help generate risk assessments that tie more easily with managerial thinking, while also saving computational effort when the assessments need to be revised. One area where the concept of ‘controllable’ uncertainty is especially relevant is Operational Risks. The paper presents an approach for operationalizing the generalized generic controller so that it is more accessible to the analyst and its algebraic properties are also made more transparent. Using this alternative approach, the generalized generic controller is adapted to represent mitigations contemplated by Operational Risk managers. The integration of the generalized generic controller with existing risk assessment models in Operations Research is illustrated using a real world case-study where the mitigation examined is actions to reduce the probability of an extended staff strike, in the event of an outage. This case study was performed for a major European insurance company. The proposed method to model interventions was built as an overlay to the existing model (in Excel) and most other inputs were borrowed from the model that was already in place. Inputs specific to the generalized generic controller were elicited using a questionnaire and phone discussions. The analysis presented, however, assumes hypothetical inputs and the relationship between the inputs and eventual cost savings is presented for this specific case.

Overall, the three papers span a wide range of research approaches and disciplines. The descriptive paper looks predominantly at the literature on behavioural decision making (e.g. Young et al., 2011; Seo et al., 2010; Fennema and Wakker, 1997; Wu and Gonzalez, 1996; Tversky and Kahneman, 1992), and some related literature in psychology (e.g. Bryant and Dunford, 2008; Crowe and Higgins, 1997; Higgins, 1996a) and behavioural agency theory (e.g. Wiseman and Gomez-Mejia, 1998). The theoretical paper mainly scopes the literature in Decision Analysis (e.g. Matheson and Matheson, 2005; Matheson, 1990; Shachter, 1986; Howard and Matheson, 1984, 2005) and areas of Philosophy which have discussed belief revision rules (e.g. Lewis, 1976; Gardenfors, 1982). The application paper

reviews some literature in Operational Risk assessment (e.g. Chaudhury, 2010; Antoine Frachot et al., 2001), in addition to the literature covered in the theoretical paper.

1.3.3 Potential research contributions

The interdisciplinary approach of this thesis, possibly, brings some novel insights to the behavioural understanding of decision making under uncertainty and also makes potential contributions to the formal modelling of decisions when uncertainties are controllable, with the intent of extending the applicability of Decision Analysis to risk management. Here a brief overview of the contributions is offered. A more detailed review is presented in the concluding chapter (Chapter 5).

In the context of Behavioural Decision Making, the findings from the present study that control has a *moderating effect* on attitudes, contributes to the prevalent understanding of how control affects risk taking. This finding is not at odds with previous research, which has found that control increases risk taking, but extends these findings to decision situations that have not been studied before. Specifically, the experimental design in this research covers cases where risk seeking behaviour persists in the absence of control (e.g. negative valued payoffs) which haven't been explored properly by previous research. An attempt is made to offer a deeper understanding of the observed effects and why control affects risk taking by evoking the affective and cognitive factors that underlie risk taking behaviour and their interaction with control. Connections are also made between these explanations and contextual variations in managerial risk attitudes. In terms of preference for exerting control, the findings from the study endorse the certainty effect: other things equal (such as relative riskiness of prospect), it is found that the control is valued more when it guarantees some gains or eliminate sure losses. While this insight is not new, it highlights the need to explicitly value 'certainty' when performing a value of control analysis so that the resulting recommendations better resonate with real world perceptions of how control is valued.

From an analytic perspective, one potential contribution of the present research is that it augments the tool kit of Decision Analysis by proposing a general probability revision rule for interventions. The proposed method builds on the existing procedures in Decision Analysis (e.g. generic controller (Matheson and Matheson, 2005)) and links it to belief revision rules in other disciplines such as Philosophy (e.g. imaging (Gärdenfors, 1982)). The potential usefulness of this development, in terms of alleviating the judgmental burden of eliciting probabilities in 'large' influence diagrams, is highlighted. A second, more theoretical, potential contribution is the development of a normative rationale for using the existing and proposed procedures for modelling interventions which mirror the coherence arguments presented for using Bayes rule as the probability revision function for

information. By showing that existing and the proposed probability revision rules for interventions satisfy some key properties that are proposed and defended in this research, the hope is that it makes the use of these probability revision rules for interventions more compelling.

From a practical perspective, the present research demonstrates how the methods developed in this thesis can contribute to the techniques of modelling and assessing Operational Risks. By demonstrating how the proposed probability revision methods for modelling interventions can be operationalized and used to model the Operational Risk mitigations, this research endeavours to integrate Decision Analysis methods with Operational Risks. It is argued that this integration can significantly alleviate the burden of both obtaining inputs for Operational Risk models and updating them periodically to reflect changes in policies and decisions in an organization. Consequently, by simplifying the process by which value of mitigations (in terms of cost savings) is quantified, it can also serve the exploration of alternative risk management policies.

1.4 Structure of this thesis

This thesis is structured as a collection of three papers which are presented as three separate chapters. These chapters acquire the style of working papers, i.e. they have the authors listed and contain an abstract. Each of the papers have been written with a target journal in mind and the presentation, length and style of the papers have been adapted to roughly meet the respective journal guidelines while respecting some uniformity for the purpose of this thesis. Additional discussions or explanations that support the contents of the papers have been included as appendices (material that will be submitted to journals) and supplementary material (material which will most likely not be submitted to journals). All the three papers are published as working papers on the Management Science website at the London School of Economics.

Chapter 2 presents the paper on the behavioural effects of control on risk attitudes. Chapter 3 is the paper which develops a probability revision rule for interventions and establishes the theoretical foundations for any probability revision rule for interventions. Chapter 4 constitutes the paper which was developed in collaboration with the Operational Risk and Human Resources team of a leading insurance company. It discusses the operationalization of the probability revision rule proposed in Chapter 3 and demonstrates its application to a specific Operational Risk scenario. The concluding chapter, Chapter 5, summarizes the key ideas of this research, discusses its limitations, the challenges experienced and directions for future research. The Annexure discusses the basic expected utility framework for analysing decisions (preliminary concept which overarches Chapters 2


and 3), its extensions to study and model behavioural attitudes to risk (relevant for Chapter 3) and a decision modelling tool, influence diagrams (used in Chapter 4 and 5).

1.5 Appendix for Chapter 1

1.5.1 The practical advantages of a probability revision calculus

Consider Figure 1-1: When a state node (e.g. S^C) depends on another state node (e.g. S^B), Bayes rule describes how the probabilities of the state node (S^C) are updated to incorporate information about the preceding state node (S^B) that it depends on. The relationship is given by $P(s_i^C | s_j^B) = P(s_i^C, s_j^B) / P(s_j^B)$. This relationship also describes the decomposition of a joint probability distribution into conditional and marginal distributions which has a number of practical advantages when analysing graphical models such as Bayes nets and influence diagrams.

Figure 1-1: Hypothetical relationship between states which can feature in an influence diagram.



Firstly, it can significantly reduce the judgemental burden of eliciting the joint probability distribution of the uncertain variables. Consider Figure 1-1 again. Suppose each of the states S^A, S^B, S^C have 3 outcomes. The joint probability distribution $P(s_k^A, s_j^B, s_i^C)$ requires 26 probability inputs. However, applying Bayes rule, we have $P(s_k^A, s_j^B, s_i^C) = P(s_k^A)P(s_j^B | s_k^A)P(s_i^C | s_j^B)$ and the right hand side of this equation needs only 14 inputs. Secondly, the decomposition of the joint distribution enables an analyst to obtain the probabilities of the various states independently from different experts. This is especially useful when working with large influence diagrams where the required information may be needed from different people. Finally, if the influence diagram is modified so that some of the dependencies are dropped (e.g. node S^A is removed) then Bayes rule can be used to compute the probabilities of the new influence diagram $P(s_j^B, s_i^C) = P(s_j^B)P(s_i^C | s_j^B)$ where $P(s_j^B) = \sum_k P(s_k^A)P(s_j^B | s_k^A)$ and the probabilities do not need to be elicited again.

1.6 References for Chapter 1

1. Armendt, B., 1986. A Foundation for Causal Decision Theory. *Topoi* 5, 3–19.
2. Basel Committee, 2003. Sound Practices for the Management and Supervision of Operational Risk.

3. Bedford, T., Cooke, R., 2001. Probabilistic Risk Analysis: Foundations and Methods. Cambridge University Press, Cambridge.
4. Bell, D., Raiffa, H., Tversky, A., 1988. Descriptive, Normative, and Prescriptive Interactions in Decision Making, in: D. Bell, H. Raiffa, A. Tversky. Decision Making: Descriptive, Normative and Prescriptive Interactions. Cambridge University Press, Cambridge.
5. Bernstein, P., 1998. Against the Gods: The Remarkable Story of Risk. John Wiley & Sons, New York.
6. Bolker, E., 1967. A simultaneous axiomatization of utility and subjective probability. *Philos. Sci.* 333–340.
7. Brandstätter, E., Schwarzenberger, H., 2001. Beyond the gambling paradigm: Internal controllability in decision-making. *Psychol. Rep.* 89, 259–266.
8. Bryant, P., Dunford, R., 2008. The Influence of Regulatory Focus on Risky Decision-Making. *Appl. Psychol.* 57, 335–359.
9. Caliendo, M., Fossen, F., Kritikos, A., 2009. Risk attitudes of nascent entrepreneurs—new evidence from an experimentally validated survey. *Small Bus. Econ.* 32, 153–167.
10. Chatterton, J., 2001. Uncertainty: Types and Sources (Weak Links project: Case Studies.), Weak Links project: Case Studies. Building Research Establishment and Association of British Insurers.
11. Chau, A., Phillips, J., 1995. Effects of perceived control upon wagering and attributions in computer blackjack. *J. Gen. Psychol.* 122, 253.
12. Chaudhury, M., 2010. A review of the key issues in operational risk capital modeling. *J. Oper. Risk* 5, 37–66.
13. Chelst, K., Bodily, S., 2000. Structured Risk Management: Filling a Gap in Decision Analysis Education. *J. Oper. Res. Soc.* 51, 1420–1432.
14. Chernoff, H., Moses, L., 1959. Elementary Decision Theory. Wiley, New York.
15. Chew, S.H., Sagi, J., 2008. Small worlds: Modeling attitudes toward sources of uncertainty. *J. Econ. Theory* 139, 1–24.
16. Chow, C., Sarin, R., 2002. Known, Unknown, and Unknowable Uncertainties. *Theory Decis.* 52, 127–138.
17. Cohen, J., 1960. Chance, Skill and Luck. Penguin books, Middlesex.
18. Cohen, J., Hansel, C.E.M., 1959. Preferences for different combinations of chance and skill in gambling.
19. Crowe, E., Higgins, E.T., 1997. Regulatory Focus and Strategic Inclinations: Promotion and Prevention in Decision-Making. *Organ. Behav. Hum. Decis. Process.* 69, 117–132.
20. Cummins, J., Embrechts, P., 2006. Introduction: Special section on operational risk. *J. Bank. Finance* 30, 2599–2604.

21. Ellsberg, D., 1961. Risk, ambiguity and the Savage axioms. *Q. J. Econ.* 75, 643–669.
22. Fennema, H., Wakker, P., 1997. Original and cumulative prospect theory: a discussion of empirical differences. *J. Behav. Decis. Mak.* 10, 53–64.
23. Fox, C., Ulkumen, G., 2011. Distinguishing two dimensions of uncertainty.
24. Frachot, A., Georges, P., Roncalli, T., 2001. Loss distribution approach for operational risk. Available SSRN 1032523.
25. Gardenfors, P., 1982. Imaging and Conditionalization. *J. Philos.* 79, 747–760.
26. Gibbard, A., Harper, W., 1978. Counterfactuals and two kinds of expected utility, in: *Ifs*, (Eds.). Reidel, Dordrecht.
27. Gilboa, I., 1999. Can free choice be known?, in: *Bicchieri et Al., The Logic of Strategy*. Oxford University Press, Oxford, pp. 163–174.
28. Goodie, A., 2003. The effects of control on betting: Paradoxical betting on items of high confidence with low value. *J. Exp. Psychol. Learn. Mem. Cogn.* 29, 598–610.
29. Goodie, A., Young, D., 2007. The skill element in decision making under uncertainty: Control or competence? *Judgm. Decis. Mak.* 2, 189–203.
30. Hacking, J., 1975. *The Emergence of Probability*. Cambridge University Press, Cambridge.
31. Heath, C., Tversky, A., 1991. Preference and belief: Ambiguity and competence in choice under uncertainty. *J. Risk Uncertain.* 4, 5–28.
32. Heckerman, D., Shachter, R., 1995. Decision Theoretic Foundations for Causal Reasoning. *J. Artif. Intell. Res.* 3, 405–430.
33. Higgins, E.T., 1996. Emotional Experiences: The Pains and Pleasures of Distinct Regulatory Systems. *Emot. Interdiscip. Perspect.* 203.
34. Hoffman, O., Hammonds, J., 1994. Propagation of uncertainty in risk assessments: The need to distinguish between uncertainty due to lack of knowledge and uncertainty due to variability. *Risk Anal.* 14.
35. Holton, G., 2002. *History of Value-at-Risk*. Working paper. Contingency Analysis, Boston.
36. Holton, G., 2004. Defining risk. *Financ. Anal. J.* 60.
37. Horswill, M., McKenna, F., 1999. The Effect of Perceived Control on Risk Taking. *J. Appl. Soc. Psychol.* 29, 377–391.
38. Howard, R., 1960. *Dynamic Programming and Markov Processes*. MIT Press.
39. Howard, R., Matheson, J., 1984. Influence Diagrams, in: R.A. Howard and J.E. Matheson (Eds.), *The Principles and Applications of Decision Analysis*. Strategic Decisions Group, Menlo Park, CA.
40. Howard, R., Matheson, J., 2005. Influence diagrams. *Decis. Anal.* 2, 127–143.
41. Howell, W., 1971. Uncertainty from internal and external sources: A clear case of overconfidence. *J. Exp. Psychol.* 89, 240.

42. Howell, W., Burnett, S., 1978. Uncertainty measurement: A cognitive taxonomy. *Organ. Behav. Hum. Perform.* 22, 45–68.
43. Huber, O., 1997. Beyond gambles and lotteries: Naturalistic risky decisions. *Decis. Mak. Cogn. Models Explan.* 145–162.
44. Huber, O., 2002. Beyond gambles and lotteries. *Decis. Mak. Cogn. Models Explan.* 145.
45. Huber, O., 2007. Behavior in risky decisions: Focus on risk defusing, in: *Uncertainty and Risk*. Springer, pp. 291–306.
46. Humphreys, P., Berkeley, D., 1985. Handling uncertainty: levels of analysis of decision problems, in: G. Wright, *Behavioral Decision Making*.
47. Jeffrey, R., 1965. *The Logic of Decision*, 2nd ed. University of Chicago Press, Chicago.
48. Joyce, J., 1999. *The Foundations of Causal Decision Theory*. Cambridge University Press, Cambridge Books Online.
49. Keynes, J., 1936. *The General Theory of Employment, Interest and Money*. Macmillan Cambridge University Press, Cambridge.
50. Knight, F., 1921. *Risk, Uncertainty and Profit*. Houghton Mifflin Co., Boston.
51. Langer, E., 1975. The illusion of control. *J. Pers. Soc. Psychol.* 32, 311–328.
52. Leavens, D., 1945. Diversification of investments. *Trusts Estates* 80, 469–473.
53. Lewis, D., 1976. Probabilities of Conditionals and Conditional Probabilities. *Philos. Rev.* 85, 297–315.
54. Lewis, D., 1981. Causal decision theory. *Australas. J. Philos.* 59, 5–30.
55. Li, K., 2011. Preference towards control in risk taking: Control, no control, or randomize? *J. Risk Uncertain.* 43, 39–63.
56. Loewenstein, G., Weber, E., Hsee, C., Welch, N., 2001. Risk as feelings. *Psychol. Bull.* 127, 267–286.
57. Luce, D., Krantz, D., 1971. Conditional Expected Utility. *Econometrica* 39, 253–271.
58. MacCrimmon, K., Wehrung, D., 1986. *Taking Risks: The Management of Uncertainty*. Macmillan, London.
59. MacCrimmon, K., Wehrung, D., 1990. Characteristics of Risk Taking Executives. *Manag. Sci.* 36, 422–435.
60. March, J., Shapira, Z., 1987. Managerial perspectives on risk and risk taking. *Manag. Sci.* 33, 1404–1418.
61. Markowitz, H., 1952. Portfolio Selection. *J. Finance* 7, 77–91.
62. Matheson, D., Matheson, J., 2005. Describing and Valuing Interventions That Observe or Control Decision Situations. *Decis. Anal.* 2, 165–81.

63. Matheson, J., 1990. Using influence diagrams to value information and control, in: R.M.Oliver & J.Q. Smith, *Influence Diagrams, Belief Nets and Decision Analysis*. John Wiley & Sons Ltd, pp. 25–48.
64. Morgan, M., Henrion, M., 1990. The Nature and Sources of Uncertainty, in: *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge University Press, Cambridge, pp. 47–56.
65. Nozick, R., 1969. Newcomb’s Problem and Two Principles of Choice, in: N. Rescher (Ed.), *Essays in Honor of Carl G. Hempel: A Tribute on His Sixty-Fifth Birthday*. D.Reidel, Dordrecht-Holland, pp. 114–146.
66. Pate-Cornell, M.E., 1996. Uncertainties in risk analysis: Six levels of treatment. *Reliab. Eng. Syst. Saf.* 54, 95–111.
67. Pearl, J., 2000. *Causality: Models, Reasoning and Inference*. Cambridge University Press, Cambridge.
68. Pearl, J., 2010. Physical and Metaphysical Counterfactuals, in: Technical Report R-359.
69. Presson, P., Benassi, V., 1996. Illusion of control: A meta-analytic review. *J. Soc. Behav. Personal.* 11, 493–510.
70. Rosenzweig, P., 2013. What Makes Strategic Decisions Different. *Harv. Bus. Rev.* 89–93.
71. Roy, A., 1952. Safety first and the holding of assets. *Econom. J. Econom. Soc.* 431–449.
72. Seo, M.-G., Goldfarb, B., Barrett, L., 2010. Affect and the framing effect within individuals over time: Risk taking in a dynamic investment simulation. *Acad. Manage. J.* 53, 411–431.
73. Shachter, R., 1986. Evaluating Influence Diagrams. *Oper. Res.* 34, 871–882.
74. Shachter, R., 2012. Effective Decision Analysis: A Cause to Believe In. *Decis. Anal. Today* 31, 29–33.
75. Shachter, R., Heckerman, D., 2010. Pearl Causality and Value of Control, in: R. Dechter, H. Geffner and J. Y. Halpern (Eds.), *Heuristics, Probability, and Causality: A Tribute to Judea Pearl*. College Publications, pp. 431–447.
76. Skyrms, B., 1982. Causal Decision Theory. *J. Philos.* 79, 695–711.
77. Strickland, L., Lewicki, R., Katz, A., 1966. Temporal orientation and perceived control as determinants of risk-taking. *J. Exp. Soc. Psychol.* 2, 143–51.
78. Teller, P., 1973. Conditionalization and observation. *Synthese* 26, 218–258.
79. Tsoukiàs, A., 2008. From decision theory to decision aiding methodology. *Eur. J. Oper. Res.* 187, 138–161.
80. Tversky, A., Kahneman, D., 1986. Rational Choice and the Framing of Decisions. *J. Bus., The Behavioral Foundations of Economic Theory* 59, S251–S278.

-
81. Tversky, A., Kahneman, D., 1992. Advances in prospect theory: cumulative representation of uncertainty. *J. Risk Uncertain.* 5, 297.
 82. Ulkumen, G., Erner, C., Fox, C., Tannenbaum, D., Walters, D., 2014. Perceived Nature of Market Uncertainty Predicts Investment Behaviors. Presented at the Behavioural Decision Research in Management, London.
 83. Von Mises, L., 1949. *Human Action: A Treatise on Economics*, 4th ed. Fox & Wilkes, San Fransisco.
 84. Von Mises, R., 1957. *Probability, Statistics and Truth*, 2nd ed. Macmillan, New York.
 85. Von Winterfeldt, D., Kavet, R., Peck, S., Mohan, M., Hazen, G., 2012. The Value of Environmental Information without Control of Subsequent Decisions. *Risk Anal.* 32, 2113–2132.
 86. Winkler, R., 1996. Uncertainty in probabilistic risk assessment. *Reliab. Eng. Syst. Saf.* 54, 127–132.
 87. Wiseman, R., Gomez-Mejia, L., 1998. A Behavioral Agency Model of Managerial Risk Taking. *Acad. Manage. Rev.* 23, 133–153.
 88. Wu, G., Gonzalez, R., 1996. Curvature of the Probability Weighting Function. *Manag. Sci.* 42, 1676–1690.
 89. Young, D., Goodie, A., Hall, D., 2011. Modeling the Impact of Control on the Attractiveness of Risk in a Prospect Theory Framework. *J. Behav. Decis. Mak.* 24, 47–70.

PRELUDE TO CHAPTER 2

This chapter is in the area of Behavioural Decision Making. Attitudes to risk are a widely studied topic in Behavioural Decision Making and there has been ample work on examining and explaining contextual variations in risk attitudes. Closer to the topic of this thesis, there has been some empirical work on studying how ‘control’ over uncertainty affects risk taking. Control, however, has been conceptualized in many distinct ways (e.g. illusory control, source preference, probability alterability).

In this chapter, control is conceptualized as the opportunity to modify probabilities of an uncertain event, before playing a gamble, and will be manipulated as a fixed shift of probability mass from the worse outcome to the better outcome. A novel randomized controlled study is presented to explore two behavioural research questions related to control: 1) its effect on risk taking; and 2) behavioural perceptions of how control is valued.

The endeavour of the study is to expand prevailing understanding of how and why control affects risk taking, and the relationship between observations in a laboratory setting and accounts of managerial risk taking from the field, where uncertainty is often viewed as being controllable. Understanding the behavioural perceptions of how control is valued can be pertinent to analytic methods which measure the value of exerting control.

Some basic concepts of the methods used to measure attitudes to risk and the basics of corresponding theories of decision making, that are relevant for this chapter, can be found in the Annexure.

Chapter 2

The Moderating Effect of 'Probability Control' on Risk Taking

Shweta Agarwal*

ABSTRACT

When faced with risky choices, decision makers can often exert some control on the risks they experience, either by modifying the probabilities or the outcomes. Previous research on the topic of 'control' suggests that people are willing to accept more risk in decision situations where probabilities are 'controllable' (e.g. depend on skill) than when they depend on pure chance. A novel study is designed to investigate whether 'control', defined as the opportunity to alter the probabilities of a gamble, before playing it, can affect risk taking and tested for three types of gambles — gains-only, losses-only and mixed gambles. While for decisions without control, the choice patterns are found to be consistent with the fourfold pattern of risk attitudes predicted by Cumulative Prospect Theory (CPT), this pattern is 'neutralized' in tasks where subjects could modify probabilities. Control was found to increase risk taking when risk-averse behaviour persists but reduce risk taking otherwise. These results are explained using previous findings about the relationship between affect and CPT as well as cognitive accounts of risk-seeking behaviour for losses, thus providing a deeper insight into *why* control affects risk acceptance. Consistent with the certainty effect, there is evidence that control is valued more when it guarantees gains or makes sure losses probable. The relation between this research and agency-based models of managerial attitudes to risk are also discussed.

Key-words: decisions under risk, control, risk taking, Cumulative Prospect Theory

*Management Science Group, Department of Management, London School of Economics and Political Science, London, UK. E-mail: s.agarwal@lse.ac.uk.

2.1 Introduction

When making decisions in the real world, individuals often have an opportunity to exert some control over the uncertainties they encounter (Goodie, 2003; Young et al., 2011). For example, in contrast to the weather on a particular day or the outcome of a sporting game (watched on television), which are ‘uncontrollable’ events, activities which depend on skill, such as how well one performs in an exam or whether a doctor will perform a medical surgery successfully are ‘controllable’ events (Brandstätter and Schwarzenberger, 2001; Rosenzweig, 2014). Similarly, in the corporate world, in addition to managing risks by controlling the outcomes (e.g. purchasing insurance), managers also endeavour to modify risks by influencing the probabilities of the uncertain events (e.g. operational risk management, marketing or product launch decisions) (March and Shapira, 1987; Courtney, 2001; Goodie and Young, 2007; Rosenzweig, 2013, pp. 23–44). Understanding attitudes to risk when uncertainties are controllable, labelled here as the ‘decision with control’ (DWC) paradigm, is important because managers view risk taking as being fundamentally different from ‘gambling’ and often associate good management with being able to affect the odds of uncertain outcomes favourably (March and Shapira, 1987; Rosenzweig, 2014, pp. 23–44). The behavioural effect of ‘control’ on risk taking is also particularly relevant to models of agency theory and corporate governance that seek to explain managerial risk taking (Lefebvre and Vieider, 2014; Wiseman and Gomez-Mejia, 1998).

The majority of empirical research on decision making under risk, however, examines choice behaviour using gambles where the probabilities of outcomes are based on pure chance events such as coin flips, drawing balls from an urn or spinning a roulette wheel. In these tasks, subjects cannot alter the probabilities of outcomes before playing the gambles, labelled here as the ‘decision without control’ (DWOC) paradigm. The stability of attitudes to risk in the DWOC paradigm has been studied extensively from a number of perspectives and there is empirical support that risk-seeking or risk-averse behaviour depends on objective contingencies of gambles such as decision frame and the magnitude of probabilities (Tversky and Kahneman, 1992; Tversky and Wakker, 1995). A widely accepted theory of choice for the DWOC paradigm is Cumulative Prospect Theory (CPT) (Tversky and Kahneman, 1992) which predicts a *fourfold pattern* of risk attitudes — risk-aversion for moderate-to-large probabilities of gains and small probability of losses; risk-seeking for moderate-to-large probabilities of losses and small probabilities of gain. The psychological implications of this theory, namely, loss aversion, reflection effect and overweighting of small probabilities can account for many decisions in the real world, such as coexistence of lottery purchases and insurance (Tversky and Kahneman, 1992), over-betting on long-shot horses (Camerer, 2004), disposition effects (Camerer, 2004), investor decisions (Olsen,

1997) and incidence of fraudulent financial reporting (Fung, 2014). It has been suggested that the extent to which our understanding of attitudes to risk in the DWOC paradigm reflects decision making in situations where an agent can influence the risks, such as strategic decisions faced by managers, is questionable (Rosenzweig, 2013). This paper, therefore, studies the effect of control on risk taking and, in particular, on the fourfold pattern of risk attitudes.

Over the years, various conceptualizations of control have been studied. Broadly, these can be classified into three types: i) effect of ‘illusionary control’ (introduced by Langer (1975), see Presson and Benassi (1996) for a review), where control is manipulated by introducing skill related but probability independent cues in tasks;¹ ii) effect of ‘control’ as a distinct source of uncertainty (Li, 2011), where control is manipulated as the opportunity to pick a random number to bet on; and iii) the effect of ‘probability alterability’ (Goodie, 2003; Young et al., 2011) manipulated by comparing bet acceptance on skill (subjective probability of answering a general knowledge question correctly) as opposed to equiprobable chance events. A common finding of all these studies is that perception of control² increases risk acceptance and this observed effect of control is not mediated by competence when probabilities depend on skill. Thus, although interlinked, inherent controllability of probabilities and competence in task (competence hypothesis (Heath and Tversky, 1991)) influence risk taking independently. Various dependent measures such as choice of bet (Brandstätter and Schwarzenberger, 2001; Goodie, 2003), coefficient of the probability weighting function (Young et al., 2011), investment amount (Li, 2011) have been used to capture the effect of control on risk taking.

In this paper, ‘control’ is conceptualized as the opportunity to modify probabilities of a gamble before playing them and closely corresponds to Goodie’s (2003) conception of probability alterability. The present study extends the prevailing insights on the effect of control on risk taking in three important ways. Firstly, in contrast to previous studies on control, where probability was only *perceived* to be controllable and subjects did not have opportunity to actually alter the probabilities, in this study subjects were allowed to actually modify the probability of the events before playing the gamble. Thus, control was an objective characteristic of the task and not subject to judgement. Secondly, since emotional responses and, consequently, attitudes to risk have been found to be sensitive to decision frame (Bryant and Dunford, 2008; Seo et al., 2010), the effect of control is examined

¹ Illusionary control was manipulated in several ways such as appearance of skill (compete against an attractive or awkward confederate in a chance task), degree of choice (choose or accept a lottery ticket), familiarity with a chance task (opportunity to try the task) and confidence (time spent contemplating if a lottery will be won).

² Perceptions of controllability and the corresponding effects on risk taking can depend on the relevance of skill in the task (Chau and Phillips, 1995; Horswill and McKenna, 1999).

separately for all domains of payoffs — gains, losses and mixed gambles. Arguably, control can interact with the interdependence of affect, risk taking and decision frame, which can be relevant to any general explanation that can be offered about the effect of control on risk taking. However, previous studies on control have mostly focussed on gains-only gambles (Young et al., 2011; Goodie and Young, 2007; Brandstätter and Schwarzenberger, 2001). Thirdly, in this study, the effect of control is studied for the entire probability scale. In previous studies for which the concept of control closely corresponds to the one used in this paper (Goodie, 2003; Goodie and Young, 2007; Young et al., 2011), the effect of control has been studied only in situations where the probability of success is at least 0.5, because of how control was manipulated (as confidence in the answer to a general knowledge question). Considering the full probability scale, allows us to better relate the observed effects of control to the variations in risk attitude captured by descriptive theories of risk such as Cumulative Prospect Theory (Tversky and Kahneman, 1992).

The findings from this study underscore the crucial mediating effect of payoffs and prior probabilities of success on any general conclusions that can be made on the effect of control on risk taking. In relation to the fourfold pattern of risk attitudes, findings from the present study suggest that control increases risk taking when risk-averse behaviour persists (i.e. for large probability of gains and small probabilities of losses), whereas the opposite is true when risk-seeking behaviour is prevalent (i.e. control reduces risk taking for small probability of gains and large probabilities of losses). Explanations for this moderating effect of control based on the interaction of control with affect, cognitive accounts of risk taking and regulatory focus are offered.

This paper also examines preferences for exerting control at particular probability levels of ‘success’ and this phenomenon is formally described here as *perceived value of control (PVoC)* — the difference in expected utility of a risky option before and after control is exerted based on the CPT framework. PVoC is an extension of the ‘Expected Value of Control’ concept from Decision Analysis, where the benefit of exerting control is assessed by measuring the standard difference in utility (without probability weighting) (Matheson and Matheson, 2005; Matheson, 1990). According to CPT, however, the psychological value of probability changes that make a possibility an impossibility, or an impossibility a possibility, is greater than intermediate changes to probability. For a fixed change in expected utility, if PVoC is found to depend on prior probability of success, then it can inform analysis based on ‘Expected Value of Control’. In particular, the resulting recommendations may need to be sensitive to whether exerting control eliminates uncertainty or not (and not just the magnitude of utility change) in order to better tie individuals’ preferences for exerting control with theoretical solutions. An area where valuing control is especially useful is Operational Risk management (Agarwal and

Montibeller, 2014) where a number of uncertainties can be modified but due to limited resources and positive cost of mitigating risks, control can be exerted only selectively. The findings from this study reveal that, in line with CPT, there is a strong preference for exerting control at the boundaries of the probability scale, i.e. when exerting control can guarantee a positive payoff or make a sure (or nearly sure) loss probable.


This paper is organized as follows. Section 2 describes the experimental set up. Section 3 and 4 discuss the analysis, findings and interpretations for the effect of control on risk taking and perceived value of control, respectively. In section 5, some auxiliary analyses corresponding to demographic and experimental factors are presented. In the concluding section, the relation of this research to managerial attitudes to risk is discussed and some directions for future research are proposed. The Appendix discusses the estimation of probability weighting functions for the data in this study and contains some other accompanying tables.

2.2 Design of experiments

335 graduate students (approximately 170 female) from London volunteered to participate in the study. Subjects were recruited online through an email invitation which contained a link to the experiment. The study was administered online and subjects were randomly assigned to one of three experiments, corresponding to three distinct payoff domains — gains-only, mixed, losses-only. Subjects were allowed to take the experiment any time they wanted and were told that the experiment would be closed as soon as the target number of respondents is reached. The three experiments employed a within-subject test of the effect of control on risk taking. Each experiment consisted of two conditions — Decision with control (DWC) and Decision without control (DWOC) — which were presented separately as two parts of the experiment.

In each condition, subjects chose between two gambles. In the DWC condition, the two gambles presented were of the type $A_{DWC} = (£x, p_A; £y, 1 - p_A)$ vs. $B_{DWC} = (£x, p_B; £y, 1 - p_B)$, $x > y$, and subjects were told that they would play *both* gambles. They were asked to choose the gamble in which they prefer to improve the probability of the better outcome ($£x$) from p_* to $p'_* = p_* + .2$ before an opportunity to play *both* gambles. Playing the two binary gambles, A_{DWC} and B_{DWC} , is theoretically equivalent to playing a ‘compound’ gamble with three outcomes of the type $(£2x, p; £(x + y), 1 - p - q; £2y, 1 - q)$, where $p = p_A * p_B$, and $q = (1 - p_A) * (1 - p_B)$. This is illustrated in Figure 2-1.

Figure 2-1: The equivalence between playing two binary gambles and playing one corresponding three outcome ‘compound’ gamble.



In the DWOC condition, subjects were presented with two gambles A_{DWOC} and B_{DWOC} which corresponded to the ‘compound’ gambles obtained by *modifying* the probability in A_{DWC} and B_{DWC} , respectively, and subjects simply choose which gamble they prefer to play. Thus, if $p'_A = p_A + 0.2$ and $p'_B = p_B + 0.2$, the gambles in DWOC condition were:

$$\begin{aligned} A_{DWOC} &= (£2x, p'_A p_B; £(y+x), (1-p'_A)p_B + (1-p_B)p'_A; £2y, (1-p'_A)(1-p_B)) \\ B_{DWOC} &= (£2x, p'_B p_A; £(y+x), (1-p'_B)p_A + (1-p_A)p'_B; £2y, (1-p'_B)(1-p_A)) \end{aligned} \quad (2-1)$$

The study consisted of 15 tests of the effect of control on risk taking. The tasks designed for the DWC condition were such that the payoffs were same for all 15 questions (gains-only: $x = 1500$, $y = 0$; losses-only: $x = 0$, $y = -1500$; mixed $x = 1500$; $y = -1500$) but the probabilities, p_A and p_B , were varied to be all pairwise combinations of 0%, 5% 20%, 50%, 75%, 80%. Table 2-1 gives the set of 15 choice tasks for the DWC condition. The tasks for the DWOC condition were derived from the DWC tasks using equation (2-1) (see Table 2-12 in Appendix for a complete list).

Table 2-1: List of 15 choice tasks for the DWC condition. Choice tasks for the DWOC condition were derived from these gambles using equation (2-1).

Task	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
p_A	.80	.80	.80	.80	.80	.75	.75	.75	.75	.50	.50	.50	.20	.20	.05
p_B	.75	.50	.20	.05	0	.50	.20	.05	0	.20	.05	0	.05	0	0

The choice of payoffs (£1500) was based on a rationale similar to the one adopted in other studies (Stott, 2006), i.e. to correspond to some real world financial activity in subjects’ lives. £1500 reflects the approximate monthly expenditure of students in London. The probability levels were chosen to be sufficiently distinct and to cover the entire scale (without exceeding a total of 15 tasks per condition), with a focus on capturing decision making behaviour at the extreme ends of the probability scale, i.e. when for one of the choice options the best outcome is impossible before control (tasks 5, 9, 12, 14, 15) or can be guaranteed after exerting control (tasks 1-5). The choice of probabilities was also guided by the desire to obtain integer numbers without needing to round them when probabilities for the DWOC condition are derived, since the probabilities were converted to number of balls

of a particular colour in an urn. The gambles were presented as draws from a bag with 100 different coloured balls with a different payoff attached to each colour. A sample of the questions presented in the DWC and DWOC conditions for task 10 in the gains-only experiment is shown in Figure 2-2.

Figure 2-2: Stimuli for the DWC and the DWOC conditions for task 10 in the gains-only experiment. The design for the other payoff domains was identical except for the winning amounts.

Decision with Control (DWC)

Imagine you are in the following situation.

You are presented with **two** bags, A and B, each of which has 100 balls. The bags contain **red** and **blue** balls.

Suppose you can draw a ball, once, from **each** bag and the **total pay-off is calculated after you draw a ball from each of them**.

For each bag, if you draw a **blue** ball you win £1500, else nothing. The number of **blue** and **red** balls in each bag is shown below:

Bag A	Bag B
50 Blue balls (£1500)	20 Blue balls (£1500)
50 Red balls (£0)	80 Red balls (£0)

Suppose that before you draw a ball from **each** bag, you can replace 20 **red** balls (£0) with 20 **blue** balls (£1500) in **any one** of the bags.

Please indicate in which bag you would prefer to replace 20 red balls with blue balls.

Replace in Bag A
 Replace in Bag B

Decision without Control (DWOC)

Imagine you are in the following situation.

You are presented with **two** bags, A and B, each of which has 100 balls. The bags contain **red**, **blue** and **green** balls.

For each bag, if you draw a **green** ball you win £3000, **blue** ball you win £1500, else nothing.

The number of **red**, **blue** and **green** balls in each bag is shown below:

Bag A	Bag B
14 Green balls (£3000)	20 Green balls (£3000)
62 Blue balls (£1500)	50 Blue balls (£1500)
24 Red balls (£0)	30 Red balls (£0)

Suppose you can draw **one** ball, only once, from **any one** of the bags.

Please indicate from which bag you would prefer to draw a ball.

Draw from Bag A
 Draw from Bag B

The order in which the subjects took the DWC and the DWOC conditions, the order of 15 questions within each condition and appearance of the gambles on the left or right of the screen was counterbalanced. Measures to enhance quality of data were taken by including two practice questions in each condition, randomly repeating two questions in each condition and presenting two questions in the DWOC condition where one gamble clearly dominated

the other gamble. Subjects were asked to rate the clarity of the questions after the practice question and the ease with which they could answer the questions at the end of each part of the experiment on a 6 point Likert scale. Demographic information such as gender, age, financial support for education and nationality was also collected. At the end of the experiment, subjects were asked to provide a brief description of how they made the decisions during the experiment. Subjects answered a total of 40 gambles questions, 6 rating questions, 10 (optional) demographic questions and 2 (optional) free text questions. The estimated time for completing the experiment was 20 minutes but subjects were allowed to take as much time as they wanted. All choices were hypothetical but subjects were incentivized by telling them that three participants would be selected at random to play one of their answered questions for real money and would be given 10% of their winnings in the selected question (for the losses-only gambles experiment, 10% of the losses was deducted from an initial endowment of £300 and for the mixed gambles experiment the subject was paid average of £300 and 10% of the winnings). In all three experiments the maximum amount that could be won was £300 and the minimum was £0.

2.3 Effect of control on risk taking

To study the effect of control on risk taking and fourfold pattern of risk attitudes, the data was analysed from two perspectives: i) overall effect of control on risk taking for each domain; and ii) effect of control on fourfold pattern of risk attitudes. In addition, the coefficients of the probability weighting and expected utility functions were also estimated (Appendix 2-1).

2.3.1 Estimation and hypotheses

Risk taking and the effect of control was measured as follows. For each pair of gambles in the two conditions, B_{DWC} and B_{DWO} are the riskier options in the sense that within the standard expected utility paradigm, selection of B_* maximizes expected utility for a risk-seeking individual. For single domain gambles this corresponds to a convex utility function, i.e. the risk-aversion coefficients in the (power) utility function, ρ_G, ρ_L , are greater than 1. For mixed gambles this corresponds to a convex utility function in both domains or an S or inverse-S shaped function where the convexity is greater than the concavity.³ Intuitively, preference for taking risk means that a person is unwilling to forego probability mass on the best outcome for a lower probability of the worst outcome and therefore prefers to choose or create gambles that maximize the probability of the best outcome rather than minimize the

³ In the case of S shaped function this means, for given loss aversion λ , $\rho_G > \rho_L + \ln(\lambda)/\ln(3000)$ if $\rho_G, \rho_L > 1$. In the case of inverse-S shaped function this means $\rho_G < \rho_L + \ln(\lambda)/\ln(3000)$ if $\rho_G, \rho_L < 1$.

chance of the worst outcome. In terms of exerting control, this means that the person who is willing to take risk prefers to exert control on the inferior gamble which has a higher prior probability of the worse outcome. If gamble B_* is selected more often than the less risky gamble A_* , then it can be concluded that on average subjects prefer to take risk.

Two measures were used to test if control affects risk taking: 1) relative proportions of riskier choices in the DWC and DWOC conditions across all subjects and tasks; and 2) relative proportions of riskier choices in the DWC and DWOC conditions in tasks when a switch in risk preference is observed between the two conditions. Based on previous studies, one would expect the proportion of risky choices to be higher in the DWC condition, especially for gains-only gambles. The following hypothesis tests if control increases risk taking in the DWC condition:

Hypothesis 1: The proportion of riskier choices is higher in the DWC condition (control increases risk taking).

In the context of fourfold pattern of risk attitudes predicted by CPT, the effect of control at various levels of probabilities of gains (or losses) was examined. In the experiments presented in this paper, the gambles in both conditions contain three possible outcomes (e.g. £0, £1500, £3000 for gains-only gambles) and therefore it is difficult to identify a single probability of gain (or loss) in an unambiguous manner. The following approach was adopted in this paper. For a given choice task, the ‘base’ gamble in the DWC condition (Table 2-1) was used to define the overall probability of gain (or loss). This overall probability was determined by taking the average of the gain or loss probabilities in the two gambles A_{DWC} and B_{DWC} . For example, in the gains domain, for a choice task $A_{DWC} = (£1500, p_A; £0, 1 - p_A)$ vs. $B_{DWC} = (£1500, p_B; £0, 1 - p_B)$, the probability of gain is taken to be $(p_A + p_B)/2$. Table 2-2 shows the ordering of the 15 tasks in decreasing order of probability of gain for this approach.⁴

Table 2-2: Ordering of tasks (as numbered in Table 2-1) based on overall probability of gains.

Method (average probability): $P(\text{gain}) = (p_A + p_B)/2$; $P(\text{loss}) = 1 - (p_A + p_B)/2$															
Task	1	2	6	3	7	4	5	8	9	10	11	12	13	15	14
P(gain)	.78	.65	.63	.50	.48	.43	.40	.40	.38	.35	.28	.25	.13	.10	.03

⁴ An alternative approach would be to use the gambles in the DWOC condition as reference and to calculate the cumulative probability of the non-zero outcomes $(1 - (1 - p_A)(1 - p_B))$. It can be verified that although this approach yields different magnitudes of probabilities of gain (or loss), the ordering of tasks in terms of probability of gain is almost identical to the approach displayed in Table 2-2 (except for tasks 4 and 7) and both approaches reasonably cover the entire probability scale. Since it is the ordering of tasks that is important for the analysis to follow, the results are presented for only the approach displayed in Table 2-2.

For the DWOC condition, the following hypothesis is postulated to test if the observed choices are consistent with the fourfold pattern of risk attitudes.

Hypothesis 2: For tasks 14 and 15 in gains-only gambles, the proportion of riskier choices is greater than 0.5 and for all other tasks the proportion of riskier choices is smaller than 0.5. For tasks 1 and 2 in losses-only gambles the proportion of riskier choices is smaller than 0.5 and for all other tasks the proportion of riskier choices is greater than 0.5.

For the DWC condition, if control increases risk taking regardless of payoff domain and probability levels, then we expect the proportion of riskier choices in the DWC condition to be higher for all tasks. The corresponding hypothesis is:


Hypothesis 3: The proportion of riskier choices in the DWC condition is higher than in the DWOC condition for all tasks.

2.3.2 Results

The median completion time for each payoff domain was around 18 minutes for all the payoff domains. 31 (17 female) of 335 respondents were excluded from the analysis because they violated dominance in at least one of the two questions (23 subjects), indicated extremely low clarity after the practice questions (6 subjects) or made inconsistent choices in all four questions that were repeated (2 subjects) (break down across payoff domains in shown in Table 2-13). The final sample size consisted of 102 (50 female) respondents for gains-only gambles, 99 (53 female) respondents for losses-only gambles and 103 (50 female) respondents for the mixed gambles. The choices also appeared to be reasonably consistent in the questions that were repeated: in the DWOC condition consistency rates were close to 75% whereas in the DWC condition consistency rates were over 80% (see Table 2-14). These consistency percentages are marginally higher than values reported in other studies, which are around 70% (Brooks et al., 2013; Brooks and Zank, 2005; Weber and Kirsner, 1997; Wakker et al., 1994; Camerer, 1989; Starmer and Sugden, 1989).

Figure 2-3 shows histograms of the overall effect of control on risk taking for gains-only, losses-only and mixed gambles for the DWC and DWOC condition. Comparing the risky choices in the DWC and DWOC conditions, as hypothesized for gains-only gambles, control increases risk taking (Hypothesis 1 is not rejected ($p < 0.01$; McNemar)). However, for losses-only and mixed gambles control was found to decrease risk taking (Hypothesis 1 is rejected and the proportion of riskier choices are significantly different ($p < 0.01$; McNemar)).

Figure 2-3: Mean proportion of riskier choices in the decision with control (DWC) and the decision without control (DWOC) conditions for all domains.



The proportions of riskier choices in each choice task are shown in Table 2-3 and plotted in Figure 2-4. In the DWOC condition, the choice pattern is found to be consistent with the fourfold pattern of risk attitudes. For gains-only gambles, risk-seeking behaviour is observed for tasks 14 and 15, where the probability of gain is small (Hypothesis 2 is not rejected for 1 of 2 tasks) and risk-aversion is observed for the remaining tasks where the probability of gain is not low (Hypothesis 2 is not rejected for 11 of the remaining 13 tasks). Conversely, for losses-only gambles, risk-aversion is observed for small probability of losses (Hypothesis 2 is not rejected for 2 of 2 tasks) and risk-seeking behaviour is observed for the remaining tasks (Hypothesis 2 is not rejected for 11 of 13 tasks). Weaker but similar results, i.e. increase in risk taking as probability of gain decreases, is observed for mixed gambles.

Table 2-3: Proportion of riskier choices for each choice task. Statistically significant differences in the DWC and the DWOC conditions are indicated with *.

Task	P(gain/no loss)	Gains		Losses		Mixed	
		DWOC	DWC	DWOC	DWC	DWOC	DWC
1	0.78	0.24++	0.33++	0.34++	0.33++	0.23++	0.32++
2	0.65	0.19++	0.32++(**)	0.34++	0.42	0.20++	0.37+ (**)
6	0.63	0.25++	0.39+ (*)	0.53	0.52	0.31++	0.44
3	0.50	0.12++	0.26++(**)	0.47	0.49	0.21++	0.40+ (**)
7	0.48	0.23++	0.39+ (**)	0.75++	0.55 (**)	0.43	0.46
4	0.43	0.16++	0.33++(**)	0.7++	0.44 (**)	0.22++	0.38+ (*)
5	0.40	0.18++	0.32++ (*)	0.86++	0.61+(**)	0.57	0.50
8	0.40	0.15++	0.35++(**)	0.72++	0.53 (**)	0.55	0.47
9	0.38	0.25++	0.36++	0.81++	0.66++(*)	0.75++	0.58 (**)
10	0.35	0.30++	0.40	0.76++	0.58 (**)	0.50	0.50
11	0.28	0.26++	0.37+	0.77++	0.49 (**)	0.57	0.50
12	0.25	0.45	0.44	0.88++	0.63+(**)	0.83++	0.56 (**)
13	0.13	0.45	0.41	0.70++	0.55 (*)	0.61+	0.44 (*)
14	0.10	0.60	0.39+ (**)	0.81++	0.57 (**)	0.83++	0.50 (**)
15	0.03	0.66++	0.44 (**)	0.77++	0.63+ (*)	0.82++	0.56 (**)
Average		0.30++	0.37++(**)	0.68++	0.53+(**)	0.51	0.46++(**)


Proportion different from 0.5: ++ $p < 0.01$; + $p < 0.05$;

DWC and DWOC proportions different (McNemar's test): ** $p < 0.01$; * $p < 0.05$

In 9 (gains-only), 11 (losses-only), 8 (mixed) out of 15 tests the proportion of riskier choices were significantly different between the DWC and DWOC condition. However, it is

not the case that the proportion of riskier choices in the DWC condition is *always higher* than in DWOC condition. Roughly, Hypothesis 3 is rejected for small probabilities of gains and large probabilities of losses but not rejected for high probabilities of non-zero gains. The proportion of riskier choices is significantly *higher* in 7 gains-only tasks ($0.4 \leq P(\text{gain})$) but *lower* in 2 gains-only tasks ($P(\text{gain}) \leq 0.1$); significantly *lower* in 11 losses-only tasks ($0.5 \leq P(\text{loss})$); significantly *lower* in 5 mixed tasks (roughly $0.6 \leq P(\text{loss})$) but significantly *higher* for 3 mixed tasks (some tasks where $0.4 \leq P(\text{gain})$).

Figure 2-4: Proportion of riskier choices at various probability levels. Shaded points indicate proportions that are significantly different from 0.5; triangles indicate proportion in the DWC condition is significantly different from proportion in the DWOC condition.



As illustrated in Figure 2-4, the findings indicate that control has a moderating effect on attitudes to risk: when individuals are inclined to avert risk in the absence of control (small probability of loss or large probability of gain), control reduces risk-aversion. However, if

individuals are inclined to take risk in the absence of control (small probability of gain and large probability of loss), then control *decreases* risk taking.

2.3.3 Discussion

The main objective of this study was to test the effect of control on risk taking in relation to the fourfold pattern of risk attitudes predicted by Cumulative Prospect Theory (CPT). The findings from the DWOC condition provide additional support for CPT as a descriptive theory of choice under risk. The predictive strength of CPT is also found to be stronger for single domain gambles, compared to mixed domain gambles, reflecting conclusions of other studies (Brooks et al., 2013; Wu and Markle, 2008). Control was found to affect risk taking but the direction of the effect is mediated by decision frame (e.g. gains or losses) as well as the magnitude of probabilities. Overall, control has a moderating effect on risk taking. Three explanations are offered for this finding. The first explanation, also discussed by extant research on control, is based on the influence of control on the emotional and affective factors that underlie risk taking tendencies (Camerer, 1992; Isen and Patrick, 1983; Lerner and Keltner, 2001; Loewenstein et al., 2001; Rottenstreich and Hsee, 2001; Seo et al., 2010; Trepel et al., 2005).⁵ The second explanation explores the interaction of control with cognitive accounts of risk-seeking in the domain of losses (e.g. Sitkin and Pablo, 1992; Thaler and Johnson, 1990). The third explanation is based on regulatory focus theory (Higgins, 1996b, 1997, 1998) which has been acknowledged to bear comparison with the role of control (Goodie, 2003). These explanations, in the corresponding sequence, are discussed next.

An explanation for observed effects of control on risk taking based on the interaction of control with emotional factors that can affect risk taking behaviour has been discussed frequently by extant research on control. For example, there is empirical evidence that when situations are perceived to be controllable, individuals are also overly optimistic (Weinstein, 1980) which can lead to increased attractiveness of controllable options. The enhanced confidence and optimism in the presence of control can explain the positive relationship between perceived controllability and risk taking for gains-only prospects (see Young et al., 2011). For gains-only gambles, the results from this study support the general finding that when probabilities can be controlled, risk taking is higher. Decision with control has not been studied previously for losses, and the present study suggests that individuals take *less*

⁵ For example, probability distortion at the boundary (impossibility to possibility, possibility to certainty) is heightened for affect laden outcomes (Camerer, 1992; Rottenstreich and Hsee, 2001), anger induces risk taking whereas fear reduces risk taking (Lerner and Keltner, 2001) and risk taking or aversive biases are reduced if the decision maker experiences pleasant feelings (Seo et al., 2010).

risk when they can control losses. The effect of control for mixed gambles is similar to losses, possibly due to the symmetric design of the gambles in this study. While at first the findings for losses-only and mixed gambles may appear counterintuitive, one explanation can be offered by extending the explanations offered for gains-only gambles, that control mitigates the unpleasant feelings induced in a risky situation. Studies on the interaction of affect with decision frames, Seo et al. (2010) and Isen & Patrick (1983) have found that positive affect (for example, average self-reported scores of ‘happy’, ‘satisfied’, ‘enthusiastic’ and ‘relaxed’ in Seo et al’s (2010) study) attenuates the relationship between decision frames and risk taking. The findings of the present study mirror these effects, which is consistent with the explanation that control induces positive affect.

A cognitive explanation, which fits the data is based on ‘perceived risk’ (Sitkin and Pablo, 1992) or ‘risk bearing’ (Wiseman and Gomez-Mejia, 1998). Sitkin and Pablo’s (1992) notion of perceived risk, defined as threats to wealth, can explain CPT’s prediction of risk aversion for gains and risk seeking for losses (Wiseman and Gomez-Mejia, 1998). In the gains domain, since award of wealth is at risk, perceived risk is higher, resulting in conservative choices. Conversely, in the losses domain since the wealth is effectively lost, perceived risk is lower which in turn increases risk taking. In agency-based theories the related notion of risk bearing is offered as an explanation for conservative behaviour of executives when firm forecasts are positive vs. willingness to take strategic risks when the forecasts are unsatisfactory (Wiseman and Gomez-Mejia, 1998). It is possible that control mitigates perceived risk or risk bearing, which in turn reduces risk-aversion in the domain of gains and risk-seeking in the domain of losses, as observed in this study. An alternative explanation for risk taking in the domain of losses is related to aversion to losses: individuals prefer to minimize the chance of any loss rather than chance of large losses, which leads them to make riskier choices (Thaler and Johnson, 1990; Wiseman and Gomez-Mejia, 1998). Aversion to losses was reflected in the explanations offered by subjects, who reported making decisions so as to avoid sure losses, where possible. Although not directly measured by this study, it is possible that control reduces aversion to losses and with it decreases risk taking in the domain of losses.

An explanation for the data based on *regulatory focus* (Higgins, 1996b, 1997, 1998) is as follows. According to regulatory focus theory, decision makers either adopt a *prevention focus* (invoked by security needs and targeted towards avoiding pain) or *promotion focus* (invoked by self-actualization needs or aspiration mindedness and targeted towards obtaining pleasure). The emotions experienced (cheerfulness or dejection in promotion focus vs. agitation or quiescence in the prevention focus) (Higgins, 1996a, 1997) as well as the strategic inclination (precautionary in the prevention focus vs. eagerness to achieve gains in promotion focus) (Crowe and Higgins, 1997) are distinct in the two regulatory orientations.

There is evidence that these orientations and the related emotions interact with the framing of risky decisions (i.e. as gains vs. losses) (Bryant and Dunford, 2008; Idson et al., 2000) as well as probability distortion (Kluger et al., 2004) and therefore can explain or mediate the risk taking behaviour predicted by CPT. In particular, in relation to commission risks (gains or losses that can result from acting), a promotion focus (characteristic of gains-only gambles) will promote risk taking and prevention focus (characteristic of losses-only gambles) will lead to risk avoidance (Bryant and Dunford, 2008; Crowe and Higgins, 1997). In the context of control, one explanation is that when exerting control in gains-only prospect, an agent acts to promote gains and the promotion focus prevails. Conversely, when exerting control in losses-only prospects, an agent acts to prevent losses and prevention focus prevails. It is possible that control causes an individual to feel more responsible for the outcomes that ensue and thus makes the notion of commission more salient. As a result, relative to risk attitudes that prevail in the absence of control, when probabilities can be controlled more risk is accepted for gains-only gambles and risk is avoided for losses-only gambles.

2.4 Perceived value of control

The risk averting choices for gains-only gambles in the DWC condition, suggest that when no losses are involved, given a choice, subjects prefer to exert control so as to guarantee some gain over exerting control to make a large gain more likely. This inference is supported by the self-reported explanation for decisions provided by a majority of the subjects (elaborated in section 2.5.3 below). For losses-only gambles, neutral attitude to risk suggests that subjects are insensitive to probability levels when exerting control. However in four tasks (5, 9, 12, 14), where one of the gambles presented a sure loss, control was exerted to eliminate sure loss (proportion significantly higher than 0.5), suggesting an aversion for sure losses. This aversion for sure losses is replicated in mixed gambles. Additionally, there is indication that subjects prefer to guarantee no loss on at least one gamble (i.e. in tasks 1, 2, 3, 4 and 5 exert control on gambles for which the probability of ‘no loss’ before control is 0.8). Using an approach similar to Wu and Gonzalez (1996), a thorough analysis of preference for exerting control at various levels of probability is presented in this section.

2.4.1 Estimation

Consider tasks 2 and 6 in Table 2-1 which are reproduced in Table 2-4 below. In both tasks, gamble $B_{DWC} = (£1500, 0.5; £0, 0.5)$ is the same but gamble A_{DWC} is different. If preference for exerting control is driven by attitudes to risk only, then we should not find any significant difference between proportion of times A_{DWC} is selected in the two choice tasks (since in both tasks they are the less risky gambles). However, if the proportion of times

control is exerted varies, then arguably, the *perceived value of control* when $p_A = 0.8$ (task 2) is different from the perceived value of control when $p_A = 0.75$ (task 6).

Notice that preference for exerting control when $p_A = 0.8$ vs. when $p_A = 0.75$ is also captured by considering the proportion of times A_{DWC} is selected in tasks 3 and 7 ($B_{DWC} = (£1500,0.2; £0,0.8)$ is same) or 4 and 8 ($B_{DWC} = (£1500,0.05; £0,0.95)$ is same) or 5 and 9 ($B_{DWC} = (£1500,0; £0,1)$ is same). These pairs of tasks are reproduced in Table 2-4. The *overall* perceived value for exerting control at $p_A = 0.8$ vs. $p_A = 0.75$ was therefore examined by comparing the proportion of times A_{DWC} is selected in the group of tasks $A_{0.8} = \{2,3,4,5\}$ vs. in the ‘matched’ group of tasks $A_{0.75} = \{6,7,8,9\}$. The matching of task groups $A_{0.8}$ and $A_{0.75}$ for $p_A = 0.8$ and $p_A = 0.75$, respectively, is illustrated in Table 2-4.

Table 2-4: Matched groups of tasks, {2, 3, 4, 5} and {6, 7, 8, 9}, for a comparison between $p_A = 0.8$ and $p_A = 0.75$. Both 0.8 and 0.75 always feature on the less risky gamble, A_{DWC} and the other gamble, B_{DWC} , is same. The task numbers correspond to the numbers in Table 2-1.

Gambles in group $A_{0.8}$	A_{DWC}	p_A	Task 2 .80	Task 3 .80	Task 4 .80	Task 5 .80
	B_{DWC}	p_B	.50	.20	.05	0
Gambles in group $A_{0.75}$	B_{DWC}	p_B	.50	.20	.05	0
	A_{DWC}	p_A	Task 6 .75	Task 7 .75	Task 8 .75	Task 9 .75

Similarly, matched task groups were constructed for pairwise comparisons between other probability levels. Table 2-5 shows the grouping of tasks for 20 comparisons between the six probability levels, so that the probabilities that were compared always featured on either the less risky gamble (groups A_*) or on the more risky gamble (groups B_*). The shaded cell depicts the matched group of tasks for the comparison described in Table 2-4.

Table 2-5: ‘Matched’ grouping of tasks for pairwise comparison between all probability levels of the proportion of times control is exerted.

P(Gain)	0	0.05	0.2	0.5	0.75	0.8
0.05	$B_{0.05}=\{4,8,11,13\}$		$A_{0.05}=\{15\}$	$A_{0.05}=\{15\}$	$A_{0.05}=\{15\}$	$A_{0.05}=\{15\}$
	$B_0=\{5,9,12,14\}$		$A_2=\{14\}$	$A_5=\{12\}$	$A_{75}=\{5\}$	$A_8=\{9\}$
0.2	$B_2=\{3,7,10\}$	$B_2=\{3,7,10\}$		$A_2=\{13,14\}$	$A_2=\{13,14\}$	$A_2=\{13,14\}$
	$B_0=\{5,9,12\}$	$B_{05}=\{4,8,11\}$		$A_5=\{11,12\}$	$A_{75}=\{4,5\}$	$A_8=\{8,9\}$
0.5	$B_5=\{2,6\}$	$B_5=\{2,6\}$	$B_5=\{2,6\}$		$A_5=\{10,11,12\}$	$A_5=\{10,11,12\}$
	$B_0=\{5,9\}$	$B_{05}=\{4,8\}$	$B_2=\{3,7\}$		$A_{75}=\{3,4,5\}$	$A_8=\{7,8,9\}$
0.75	$B_{75}=\{1\}$	$B_{75}=\{1\}$	$B_{75}=\{1\}$	$B_{75}=\{1\}$		$A_{75}=\{6,7,8,9\}$
	$B_0=\{5\}$	$B_{05}=\{4\}$	$B_2=\{3\}$	$B_5=\{2\}$		$A_8=\{2,3,4,5\}$

The number of tasks used in each pair of matched task groups is different because the proportion of times a given probability level features on a less risky gamble, A_{DWC} , is different from the proportion of times it features on the more risky gamble, B_{DWC} . Using the data in the DWC condition, perceived value of control was examined by comparing the

proportion of times A_{DWC} (B_{DWC}) is selected for each pair of task groups $A_*(B_*)$ in Table 2-5.

2.4.2 Results

Table 2-6 shows the various pairwise comparisons of the proportion of times control was exerted at various probability levels after controlling for the relative riskiness of the gambles for which the probabilities are compared. For example, the shaded cell in Table 2-6a describes the comparison discussed in Table 2-4, in which 62% times control was exerted when $p_A = 0.75$ whereas 68% times control was exerted when $p_A = 0.8$.

Table 2-6: Proportion of times control is exerted at various probability levels when the relative risk at the level is same. The pair in each cell compares the proportion of times control is exerted on the row probability vs. on the column probability for the task groups in Table 2-5. Statistically significant differences are marked with *.

6a. Gains-only gambles						
P(gain)	0	0.05	0.2	0.5	0.75	0.8
0.05	0.37, 0.38		0.56, 0.61	0.56, 0.56	0.56, 0.64	0.56, 0.68*
0.2	0.35, 0.38	0.35, 0.35		0.59, 0.59	0.59, 0.64	0.59, 0.67*
0.5	0.36, 0.34	0.36, 0.34	0.36, 0.33		0.59, 0.62	0.59, 0.69**
0.75	0.33, 0.32	0.33, 0.33	0.33, 0.26	0.33, 0.32		0.62, 0.68**

6b. Losses-only gambles						
P(loss)	1	0.95	0.8	0.5	0.25	0.2
0.95	0.50, 0.61**		0.37, 0.43	0.37, 0.37	0.37, 0.34	0.37, 0.39
0.8	0.53, 0.62**	0.53, 0.49*		0.44, 0.43	0.44, 0.40	0.44, 0.47
0.5	0.56, 0.63**	0.56, 0.48*	0.56, 0.52		0.43, 0.42	0.43, 0.48
0.25	0.33, 0.61**	0.33, 0.44*	0.33, 0.49**	0.33, 0.61**		0.44, 0.50**

6c. Mixed gambles						
P(loss)	1	0.95	0.8	0.5	0.25	0.2
0.95	0.44, 0.54**		0.44, 0.50	0.44, 0.44	0.44, 0.42	0.44, 0.50
0.8	0.45, 0.54**	0.45, 0.45		0.53, 0.47	0.53, 0.48	0.53, 0.56
0.5	0.47, 0.54*	0.47, 0.42	0.47, 0.42		0.48, 0.50	0.48, 0.57**
0.25	0.32, 0.50**	0.32, 0.38	0.32, 0.40	0.32, 0.50**		0.51, 0.59**

** p value < 0.01; * p value < 0.05 (McNemar)

For gains-only gambles, the proportion of times control is exerted at various probability levels is not significantly different except when the prior probability of gain is 0.8, where it is significantly higher (last column of Table 2-6a). For losses-only gambles, the relationship between probability levels and perceived value of control is less clear cut. Perceived value of control is significantly higher when it makes a sure loss probable (first column of Table 2-6b) or when it is exerted on moderate-to-large probability of loss (last row of triangle of Table 2-6b). Subjects were agnostic between reducing large probabilities of losses and guaranteeing no loss (last column of Table 2-6b). For mixed gambles the

aversion for sure losses is replicated (first column of Table 2-6c). There is also indication that in situations when the probability of gain is moderately high (above 0.5) subjects sought to guarantee no loss (bottom right cells of Table 2-6c).

2.4.3 Discussion

The analysis for perceived value of control reveals that there is a strong preference to exert control on the boundaries of the probability scale and individuals aspire to eliminate sure losses or ensure some gain. This is compatible with the well-known ‘certainty effect’ (Tversky and Kahneman, 1986). The desire for certainty is again directly related to negative emotions such as ‘worry’ that decision under risk invokes and therefore the emotional impact is greater for modifications of extreme probabilities compared to intermediate probabilities (Loewenstein et al., 2001).

The present study supports findings of Payne et al. (1999), that individuals prefer to exert control (manipulated in their study by increasing the value of any one outcome by a fixed amount) such that it improves the probability of some gain and reduces the probability of any loss. In their study and other studies (Fennema and Wakker, 1997), control has been applied as an alternative methodological tool for validating CPT. Fennema and Wakker (1997) found that, in line with CPT, subjects prefer to alter probabilities of extreme outcomes, especially for flat tailed gambles. The present study confirms that observed preferences for exerting control can be captured by CPT’s modelling of diminishing sensitivity with respect to changes in probability (Fennema and Wakker, 1997). In relation, the analysis of the data in the DWC condition can be considered a different methodological test of CPT.

In terms of probability weighting functions, the findings suggest that for gains-only gambles the weighting function is close to linear throughout the probability range with some convexity at the right end of the scale (shape predicted by Segal (1987)). For losses-only gambles, the probability weighting function takes the popular inverse-S shape — convex at the upper end of the scale ($P(\text{loss}) > 0.8$) and concave at the lower end of the scale ($P(\text{loss}) < 0.5$). For mixed gambles the probability weighting function is similar to that of losses-only gambles, with a more pronounced boundary effect. Furthermore, there is indication that boundary effects are more salient drivers of the non-linearity of probability weighting function for ‘controllable’ probabilities especially when losses are a potential outcome of the risky decision (e.g. Tversky and Wakker, 1995). This differs from discussions about the shape of the probability weighting function which were obtained in experiments where subjects did not exert control (e.g. Wu and Gonzalez, 1996). An estimation of probability weighting functions is offered in Appendix 2-1.

The crucial desire to bring about certainty (or uncertainty in the case of losses), exhibited in this study suggests, that when evaluating the value of control in Decision Analysis, the calculations should be sensitive to whether control eliminates uncertainty or simply alters probabilities. Furthermore if CPT-based models are used to perform such analysis, different probability weighting functions may be needed to evaluate decisions where probabilities are not controllable vs. where probabilities are controllable.

2.5 Other analysis and findings

The demographic effects on risk taking, potential effects of the experimental manipulation and the qualitative self-reported explanations about how subjects chose were also analysed. The findings are presented in this section.

2.5.1 Demographic effects

At the end of the experiment, subjects were given the option to provide some demographic information about themselves. Four demographic variables were analysed — gender (male/female), age (18-22, 23-27, 28+), nationality (emerging market country or developed country) and sponsorship for current programme of study (not sponsored or sponsored) — using a stepwise logistic regression ($p=0.05$) with only first order effects. Table 2-7 reports the variables which were significant predictors of risk taking and changes in risk taking as well as the corresponding odds ratio. The proportions of riskier choices for each demographic group are shown in Appendix 2-2 (Table 2-15).

Broadly, previous research has found that risk taking reduces with age (Byrnes et al., 1999; Vroom and Pahl, 1971) and males are more likely to take risk (Byrnes et al., 1999; Charness and Gneezy, 2012). These trends are roughly confirmed by the present study but mediated by the payoff domain and presence of control. When non-zero gains are a possible outcome, the odds of taking risk was approximately 30-50% higher for males, compared to females (no difference was observed for losses-only gambles). For mixed gambles, increase in risk taking in the DWC condition was significantly higher for males. Compared to students in the 18-22 years group, older students (23+ years group) made more conservative choices for gains-only gambles (odds of taking risk are 50% lower), regardless of whether they can exert control, but conversely, for losses-only gambles when they can exert control, the odds of them taking risk is twice as much.

Table 2-7: Odds ratios from stepwise regression on demographic variables (for risk taking in the DWC and the DWOC conditions and increase in risk taking in the DWC condition).

Baseline Variable	Variable	Gains			Losses			Mixed		
		DWOC	DWC	Higher risk in DWC	DWOC	DWC	Higher risk in DWC	DWOC	DWC	Higher risk in DWC
Male	Female	0.76	0.67					0.48	0.70	
18-22	23-27 yrs	0.54	0.52		3.58	1.97				0.72
18-22	28+ yrs	0.50	0.68		7.19	1.03				1.38
EM Nat.	Dev Nat.		1.81	1.72				2.13	1.81	
Unspn.	Spons.	1.43	2.61	1.48	0.55	0.70		0.79		
	n	1395	1395	1395	1440	1440	1440	1380	1380	1380
	LL	-843	-853	-1274	-908	-922	-1365	-947	-906	-1333
	R-square	0.02	0.06	0.01	0.00	0.07	0.02	0.01	0.05	0.02

Only statistically significant odds ratios are reported.

No difference in risk taking behaviour was observed between students from developed countries and emerging market countries in the DWOC condition, but the odds of risk taking for students from developed countries is nearly twice that of students from emerging market countries when they can exert control and the gambles contain positive payoffs. Fully sponsored students took more risk than students who were partially or not sponsored for gains-only gambles, especially in the DWC condition.

2.5.2 Experimental factors

It is possible that the design of the experiment and layout of questions affected the responses. Although, on average, subjects found the instructions in the practice questions reasonably clear (Likert rating >4.5 on 1-6 scale), the reported clarity for the DWOC condition was significantly higher (>5). The cognitive difficulty of making the decisions (asked at the end of each condition) was more or less same for losses-only and mixed gambles but for gains-only gambles subjects found decisions in the DWOC condition more difficult. The mean scores are reported in Appendix 2-2 (Table 2-16). A regression was performed to examine any effect of the order in which the two conditions appeared and the side of the screen on which the less risky gamble (gamble A) appeared. The side of the screen on which gamble appeared did not affect risk taking. However, whether the DWC condition was experienced first or second affected overall risk taking, which highlights an important limitation of within subject experimental designs. For single domain gambles, in both conditions, the proportion of riskier choices was relatively higher if the DWC condition was experienced first (except for DWOC, losses-only gambles). For mixed gambles, more risk was taken in the condition that was experienced first. The odds ratio and proportions of riskier choices is reported in Appendix 2-2 (Table 2-17 and Table 2-18, respectively).

2.5.3 Qualitative responses

The qualitative self-reported explanations provided by subjects were also analysed. Only a small proportion of participants noticed that the expected value was the same (20% when gains were present and 9% for losses-only gambles). A few subjects appear to have miscalculated expected value and did not realize they were equal. Loosely, five other strategies used by subjects were identified for all the payoff domains — 1) maximize the chance of the best payoff; 2) minimize the chance of the worst payoff; 3) trade off probabilities on extreme payoffs (DWOC); 4) Equalize risk (DWC); and 5) focus on achieving probabilities above and below a threshold (DWC). For gains-only gambles, majority of the subjects attempted to seek some gain (strategy 2). For losses-only gambles, both minimizing probability of any loss (strategy 1) and probability of worst loss (strategy 2) were equally popular but subjects also consciously avoided sure losses (especially in the DWC condition). For mixed gambles, minimizing losses (strategy 2) was most popular, followed by strategies 1 and 3.

2.6 Conclusion

Most laboratory and field studies that have sought to validate Cumulative Prospect Theory (CPT) are based on a paradigm where individuals cannot exert any control on the probabilities. In many real world decisions, however, individuals can exert some control on probabilities of events. Scholars have noted that risk acceptance and attitudes to risk as described by CPT are sensitive to situational and emotional factors such as affect (e.g. Isen and Patrick, 1983; Rottenstreich and Hsee, 2001; Seo et al., 2010), anxiety (Thompson, 1981), time pressure (Kocher et al., 2013) and perceived ‘controllability’ of the probabilities of a gamble (Brandstätter and Schwarzenberger, 2001; Goodie, 2003; Langer, 1975; Weber and Milliman, 1997). The aim of this paper was to further the understanding of how one of the situational factors — control — affects the ‘risk biases’ predicted by CPT, by considering all payoff domains and different probabilities of success. In particular the generality of previous findings (focused mostly on gains-only gambles), that control increases risk taking, was investigated. Preferences for exerting control at various probability levels, formalized here as perceived value of control, were also examined.

The findings from this study cast a new perspective on the effect of control on risk taking. In particular, it is found that control does not always increase taking — control increases risk taking in favourable situations but reduces risk taking in less favourable situations. Thus, in the context of CPT, control has a moderating effect on the fourfold pattern of risk attitudes and this effect is robust across all payoff domains. This result closely corresponds to the effect that positive affect has on the fourfold pattern of risk attitudes (Isen

and Patrick, 1983; Seo et al., 2010). An account of the observed effects was constructed by combining previous explanations regarding the effect of control on emotions — that control reduces the negative feelings experienced in risky situations (Thompson, 1981) or enhances optimism (Weinstein, 1980; Young et al., 2011) — with findings on the relation between affect and risk taking (Isen and Patrick, 1983; Seo et al., 2010). A cognitive explanation based on the mitigating effect of control on risk bearing, proposed as a determinant of risk taking (Sitkin and Pablo, 1992; Wiseman and Gomez-Mejia, 1998), was also offered. Additionally, an alternative interpretation within the context of regulatory focus (Crowe and Higgins, 1997; Idson et al., 2000) was explored; control possibly makes commission risks more salient, thus enhancing risk taking for positive outcomes but reducing it for negative outcomes. With regards to how control is valued and its relation to prior probability of success it was found that, in line with the certainty effect (Tversky and Kahneman, 1986), control was preferred more when it guaranteed (eliminated) probable gains (losses) or made sure losses probable.

The findings from this study can also contribute to the agency-based literature that has attempted to integrate risk taking in the corporate world with the behavioural insights offered by Prospect Theory (Lefebvre and Vieider, 2014; Wiseman and Gomez-Mejia, 1998). In particular they can help understand the complex contextual variations in managerial risk taking. For example, although managers believe that risk taking is essential for success, risk is avoided if it poses a significant threat to the current position of the firm (March & Shapira 1987). Alternatively, by viewing stock options as a compensation instrument that endows executives with greater control over their earnings, the findings from this study are consistent with findings that executive compensation through stock options increases risk taking (Anderson and Fraser, 2000; Lefebvre and Vieider, 2014; Wiseman and Gomez-Mejia, 1998) but post regulation in the 1990's, the relation between managerial stock holding and banks (Anderson and Fraser, 2000) was found to be negative. One oversimplified explanation is that the threat of penalties for taking unwarranted risks enhance the element of losses in the risky decisions that are considered by managers and, as found in this study, risk taking is lower despite having control. The connections proposed between this study and managerial risk taking can be informative for designing executive compensation contracts and open some interesting avenues for future research.

This study has some limitations inherent to the design of the experiments. Firstly, the gambles in the DWOC condition were presented as a multiple outcome gambles whereas those in the DWC condition were presented as binary gambles. The number of outcomes can affect how individuals reason about the gambles and subsequent patterns of choices of studies (Brooks et al., 2013; Payne, 2005). More experiments, with varied designs that test of the effect of control, are needed to confirm the robustness of the findings presented in this

study. Secondly, in this study, control was manipulated as an unbiased, objective contingency of the gambles and the level of control was fixed. In real world situations control is subject to perception and depending on the context individuals can feel different levels of control. A potential area for further research is to check if the moderating effect of control found in this study depends on perceived levels of control in any systematic way. Finally, the explanation offered for the findings in this study were based on the impact control has on the emotions that underlie risk taking, its interaction with regulatory focus and cognitive arguments related to the notion of risk bearing. While an attempt was made to extract cognitive explanations provided by subjects, the emotional effects or impact on regulatory focus was not directly captured. The validity of the explanations offered in support of the findings presented in this paper, therefore, also merits further empirical investigation.

Acknowledgements

This research is sponsored by the AXA Research Fund. The author would like to thank Gilberto Montibeller and Alec Morton for insightful discussions and feedback on this paper. The author would also like to thank Greg Davies, Emily Haisley, Peter Brooks and Barbara Fasolo for helpful comments on the experimental design; and Ian Durbach for feedback on the manuscript. Statistical analysis of the data and its interpretation is the responsibility of the author only.

2.7 Appendices for Chapter 2

Appendix 2-1 Parameter estimation

While measuring the effect of control by comparing the proportion of risky choices with and without control is fairly unambiguous, capturing the effect of control on risk acceptance parametrically, within the Cumulative Prospect Theory (CPT) framework, poses some challenges. In particular, this effect can be captured in two ways: as a less (more) concave (convex) value function or an elevated probability weighting function. Previous studies have suggested capturing the effect of control on risk acceptance as an elevated probability weighting function (Gonzalez and Wu, 1999; Young et al., 2011). Compared to gains-only bets on chance (uncontrollable) probabilities, Young et al. (2011) found significantly elevated probability weighting function for (gains-only) bets with knowledge relevant (controllable) probabilities but no significant difference in the risk-aversion coefficient or probability distortion. The effect of control for mixed gambles and losses-only gambles was not investigated in their study. Here, the parametric effect of control within the CPT framework is formally tested by fitting probability weighting functions and utility functions for the DWC and DWOC condition at pooled participant level.

Cumulative Prospect Theory and Model Estimation

In the CPT framework, if p_1, p_2, \dots, p_n denote the probabilities attached to the payoffs x_1, x_2, \dots, x_n , $x_i < x_{i+1}$ and $w(p_i)$ and $U(x_i)$ represent the nonlinear transformations, the expected utility, $EU(X)$, of the prospect $X = \{x_1, p_1; x_2, p_2; \dots; x_n, p_n\}$ is given by:

$$EU(X) = \sum_{i=1}^n W(p_i)U(x_i) \tag{2-2}$$

where if $x_k = 0$, the decision weights, $W(p_i)$, for gains and losses are given by:

$$W(p_i) = \begin{cases} w\left(\sum_{j=i}^n p_j\right) - w\left(\sum_{j=i+1}^n p_j\right); & k < i \leq n \\ \left(1 - w\left(1 - \sum_{j=1}^i p_j\right)\right) - \left(1 - w\left(1 - \sum_{j=1}^{i-1} p_j\right)\right); & 1 \leq i < k \\ 0; & i = k \end{cases} \tag{2-3}$$

Wakker (2010, pp. 219–222) provides a detailed explanation of the relationship between the weighting functions for losses and gains. The power function is a popular choice for the utility function (Wakker, 2010, p. 256) whereas the Tversky and Kahneman (1992), the Linear-in-Log-odds function (Goldstein and Einhorn, 1987; Gonzalez and Wu, 1999) and the Prelec (Prelec, 1998) functions are frequently used for probability weighting functions (see Stott (2006) for a comprehensive list). Studies that have looked at the best fitting probability weighting functions (Cavagnaro et al., 2013; Stott, 2006), however, have found that the Prelec function and the Linear-in-Log odds function provide better empirical fits. These functions, along with the psychological interpretation of the parameters, are shown in Table 2-8.

Table 2-8: The utility function and probability weighting functions that are commonly assumed for Cumulative Prospect Theory.

Utility function	Probability weighting functions
Power: $U(x) = \begin{cases} x^{\rho_G} & x \geq 0 \\ -\lambda x^{\rho_L} & x < 0, \lambda > 0 \end{cases}$	Prelec (Pr): $w_{Pr}(p) = \exp(-\beta(-\ln(p))^\alpha)$ Linear-in-Log-odds (GE): $w_{GE}(p) = \frac{\beta p^\alpha}{\beta p^\alpha + (1-p)^\alpha}$
ρ: Attitude to risk (risk neutral when $\rho = 1$, risk averse when $\rho < 1$, risk-seeking when $\rho > 1$) λ: Loss aversion (loss-averse for $\lambda > 1$, gain-seeking for $\lambda < 1$).	α: Sensitivity to probabilities (complete insensitivity when $\alpha = 0$, diminishing for intermediate probabilities when $\alpha < 1$, diminishing for extreme probabilities s when $\alpha > 1$) β: Optimism (neutral when $\beta = 1$, optimistic $\beta < 1$, pessimistic $\beta > 1$)

Since the data in the present study is obtained from choices between pairs of gambles, a stochastic choice analysis at pooled participant level, akin to the analysis in studies by Wu and Markle (2008), Stott (2006), Camerar and Ho (1994) and Wu and Gonzalez (1996), was performed. In this analysis, it is assumed that choices are generated by a stochastic process, where each choice depends on the deterministic difference in expected utility of the gambles (evaluated based on the subject’s true underlying preference functions) *plus* some noise. If

the noise terms are assumed to be independent and identically distributed and follow a double exponential (Gumbel Type II extreme value) distribution, which is a popular choice (Sokol-Hessner et al., 2013; Wu and Markle, 2008; Stott, 2006; Camerer and Ho, 1994; Carbone and Hey, 1994), the probability of selecting gamble A over B ($P(A > B)$) is given by the following Logit choice function:

$$P(A > B) = \frac{1}{1 + \exp(-\mu(EU(A) - EU(B)))} \quad (2-4)$$

where μ measures the sensitivity of the choices to utility differences (Wu and Markle, 2008). For a given model for $U(\cdot)$ (choice of functional forms for utility and probability weighting), the best fitting parameters $\{\theta_i^*\}$ are estimated by maximizing the Log-likelihood of generating the observed proportions of choices of each gamble or equivalently the Logarithm of the joint probability of the choices, $\prod_{t,n} P(A > B|\{\theta_i\})$, observed for all n participants and t tasks. If $\{\theta_i\}$ is the set of parameters assumed in the expected utility model $EU(\cdot)$, t is the index for the 15 tasks and N_{tA} and N_{tB} denote the number of times gamble A and B were selected in task t , the objective function to be maximized is:

$$\max_{\mu, \{x_i\}} \sum_{t=1}^{15} N_{tA} \ln(P(A > B|\{\theta_i\})) + N_{tB} \ln(1 - P(A > B|\{\theta_i\})) \quad (2-5)$$

A more detailed discussion of this approach, including review of studies where similar analysis is performed, can be found in Stott (2006). The maximum Log-likelihood provides a measure of the fit of the model assumed for $EU(\cdot)$ as follows: if the model does no better than chance, assigning $P(A \geq B) = P(A \leq B) = 0.5$ for every task and participant, the Log-likelihood for the $15 * n$ observations would be $15 * n * \ln(0.5)$, yielding in the present study, maximum Log-likelihood values of -1060.5, -1029.32 and -1070.9 for the gains-only, losses-only and mixed gambles, respectively. Conversely, a model which generates a perfect prediction would have a maximum Log-likelihood value of $15 * n * \ln(1) = 0$.

Since previous studies have reported that the best-fitting model (one which maximizes the Log-likelihood) can depend on the specification of the choice of the functional forms assumed in the CPT framework (Cavagnaro et al., 2013; Stott, 2006), three different models for calculating $EU(\cdot)$ (by varying choice of probability weighting functions) were tested — 1) No probability weighting; 2) Prelec (Pr) function; and 3) Lin-in-Log-odds (GE) function. In addition, for the DWC condition, based on the perceived value of control (equation (2-6)), two other methods were used to calculate the expected utility for different choices of probability weighting functions — 1) Prelec function (Pr); and 2) Lin-in-Log-odds function (GE). The expected utility in equation (2-4) is replaced by PVoC which, for binary gambles $\{x_1, p_1; x_2, p_2\}; x_1 < x_2$, is given by:

$$PVoC(X) = (W(p_1 - 0.2) - W(p_1))U(x_1) + (W(p_2 + 0.2) - W(p_2))U(x_2) \quad (2-6)$$

Given the design of the present study (single domains, zero valued payoff), the objective function was simplified where possible, by dropping some parameters.⁶ The non-linear optimization was performed in Matlab; a procedure for minimizing the negative of equation (2-5) using the `fmincon` function was written, where all variables except μ were constrained to be positive.

The estimation was carried out similar to Wu and Markle (2008), where μ, α, β were taken to be free parameters and the utility function $(\rho_G, \rho_L, \lambda)$ was fixed.⁷ The probability weighting function for mixed gambles was assumed to be the same for losses and gains. For the CPT models, the parameter estimation was repeated for 5 different values of ρ_G and ρ_L (0.25, 0.5...,1.25) and λ (0.5, 1,...,2.5) and the set of solutions which maximized equation (2-5) across all these fixed values of ρ, λ is reported.⁸ Since the solution was also found to be sensitive to the initial values supplied in the optimization, for each model the optimization was repeated for 5 different initial values⁹ for μ, α, β (0.25, 0.5...,1.25) and the set of parameters which maximized equation (2-5) across all the iterations were selected. Adjusting the Log-likelihood (LL) for the number of free parameters (k) in a model, the Akaike Information Criterion ($AIC = -2*LL+2k$) was used to determine which model provided the best fit.

Results and Discussion

The ordering of models based on increasing AIC is shown in Table 2-9. The CPT framework with the GE probability weighting function was found to be the best fitting model in most cases.

⁶ In the PVoC model, the differential, $PVoC(A) - PVoC(B)$, for gains-only gambles is given by $(W(p_2^A + .2) - W(p_2^A)) - (W(p_2^B + .2) - W(p_2^B))U(1500)$. $U(1500)$ can be absorbed into the coefficient μ in equation (2-4) and equation (2-5) is maximized for various combinations of μ and ρ_G . Thus, when fitting the model for PVoC, the utility function can be ignored. The same is true for losses-only gambles. Similarly, for the CPT framework when estimating the parameters for the losses-only domain, the loss aversion coefficient can be ignored.

⁷ In a previous attempt at estimating the parameters, both the utility $(\rho_G, \rho_L, \lambda)$ and probability weighting parameters α, β as well as the sensitivity parameter μ were assumed to be free parameters. Two observations are made about the solutions obtained in this attempt. First for single domain gambles, the values of ρ_G, ρ_L that were estimated were very small (less than 0.1). One reason for the extreme concavity (convexity for losses) might be the spacing of the payoffs in this experiment. It is likely that £1500 and £3000 may not have been perceived to be very different. Given that the utility function also captures the perceived psychological worth of a given outcome, the corresponding utility function would be relatively flat between the £1500 and £3000. Second, for mixed gambles, the optimal utility function and μ was not unique, but the probability weighting function was unique. This may be because the number of free parameters is larger. Given these difficulties with the estimation, the restrictions used by Wu and Markle (2008) were introduced.

⁸ For mixed gambles, when fitting the CPT framework, the optimal solution for the utility function was again found to be non-unique in the DWOC condition and the value of β as well as the utility function was found to be non-unique for the DWC condition. For the DWOC condition, the utility function selected is the one for which the value of β best approximate the values obtained in the first trial. For the DWC condition the same utility function was selected.

⁹ Chosen to manage the computational challenges of estimating the parameters.

Table 2-9: Ordering of models based on AIC (lower AIC implies better explanatory power of a model). ‘Pr’ stands for the Prelec probability weighting function and ‘GE’ for the Linear-in-Log-odds probability weighting function.

Gains-only		Losses-only		Mixed	
DWOC	DWC	DWOC	DWC	DWOC	DWC
CPT-Pr 1716.6	CPT-GE 2020.1	CPT-GE 1720.3	CPT-GE 2028.0	CPT-GE 1826.3	PVoC-GE 2110.1
CPT-GE 1727.0	CPT-Pr 2023.3	CPT-Pr 1720.7	CPT-Pr 2029.2	CPT-Pr 1826.8	PVoC-Pr 2115.8
	PVoC-GE 2026.1		PVoC-GE 2029.5		CPT-GE 2116.0
	PVoC-Pr 2026.6		PVoC-Pr 2029.5		CPT-Pr 2116.0
EU 1795.1	EU 2030.5	EU 1849.8	EU 2056.0	EU 2149.8	EU 2146.3

The parameter estimates for the CPT framework with GE probability weighting function is shown in Table 2-10. The estimates for the Prelec probability weighting function are qualitatively similar.¹⁰ Comparing the Log-likelihood (LL) for the DWC and DWOC conditions in Table 2-10, it can be seen that the CPT fit is superior for the DWOC condition.

Table 2-10: Parameter estimates for the CPT and the PVoC models (with the GE probability weighting function).

	Gains-only			Losses-only			Mixed		
	DWOC	DWC	DWC	DWOC	DWC	DWC	DWOC	DWC	DWC
	CPT	CPT	PVoC	CPT	CPT	PVoC	CPT	CPT	PVoC
α	0.72	0.89	0.91	0.62	0.18	0.00	0.31	0.00	0.41
β	1.98	2.24	0.33	0.50	0.40	2.00	1.53	0.82	1.18
ρ_G	0.25	0.25	-				0.25	0.25	0.25
ρ_L				0.25	0.25	-	0.25	0.25	0.25
λ				-	-	-	1.50	1.50	1.00
μ	4.35	3.38	2.34	3.59	0.45	0.79	0.59	0.08	0.13
LL	-860.48	-1007.07	-1010.06	-857.16	-1011.02	-1011.76	-910.17	-1055.01	-1052.07
LR	400.06	106.90	100.90	344.29	98.99	35.13	321.49	31.81	37.68


This is not surprising as the CPT model has been originally proposed for choices people make in the absence of control. The PVoC model which reflects preference for exerting control at particular probability levels does not do better, except for mixed gambles where it is marginally better. Consistent with previous findings (Young et al., 2011), it is found that changes in risk taking attitudes in the presence of control are captured not by the curvature of the utility function but by the elevation of the probability weighting function.

¹⁰ Except for the gains-DWC and mixed-DWC conditions, where the utility functions differed from the utility functions in the GE probability weighting function fit (less risk-aversion for gains-DWC and gain seeking for mixed-DWC) which affected the elevation of the probability weighting function (less elevated).

The parameter estimates for the utility and probability weighting function have been noted to vary across various studies (see Stott, 2006, p. 116 for gains-only tasks). The probability sensitivity parameter, α , for gains-only gambles in the DWOC condition obtained in this study ($\alpha = 0.72$) somewhat resemble the ones reported in Brandstatter et al. (2002) ($\alpha = 0.77, \beta = 0.88$), Wu and Gonzalez (1996) ($\alpha = 0.68, \beta = 0.84$) and Young et al's (2011) within-subject experiment ($\alpha = 0.78, \beta = 0.97$), but the value of $\beta = 1.98$ in this study is comparatively much higher. (This is probably due to the low value of $\rho_G = 0.25$ in this study compared to the value of $\rho_G > 0.85$ assumed in the other studies).¹¹ In the within-subject design Young et al. (2011) found $\alpha = 0.77$ and $\beta = 1$ for the knowledge condition, which was their manipulation for controllable probabilities.

Figure 2-5 plots the probability weighting functions for the estimates in Table 2-10. Similar to the within subject test of Young et al. (2011), for gains-only gambles, compared to the DWOC condition the elevation of probability weighting function is significantly higher for the DWC condition ($t = 4.2, p < 0.0001$). Young et al. (2011) found no significant difference in the sensitivity parameter between conditions, but in this study the sensitivity to probability was found to be higher for the DWC condition (gains-only gambles) (α is closer to 1, $t = 13.8, p < 0.0001$). The opposite is true when losses are present (for losses-only gambles $\alpha: t = 17.8, p < 0.0001; \beta: t = 4.9, p < 0.0001$ and for mixed gambles $\alpha: t = 8.1, p < 0.0001; \beta: t = 16.2, p < 0.0001$). The finding in Wu & Markle's (2008) study, that compared to single domain gambles, sensitivity to probabilities captured by α is diminished for mixed domain gambles, is also replicated in this study for both DWOC ($t = 10.0, p < 0.0001$) and DWC conditions ($t = 5.6, p < 0.0001$).

Figure 2-5: Comparison of the probability weighting functions estimated for the DWOC and the DWC conditions for all three domains of gambles (based on parameter estimations in Table 2-10).



As one would expect, the probability weighting functions obtained for perceived value of control in the DWC condition, compared to the CPT probability weighting functions, have a

¹¹ 0.88 in Brandstatter et al. (2002) and 0.853 in Young et al. (2011). The value for Wu and Gonzalez (1996) was not reported.

different elevation (lower for gains-only gambles ($t = 19.79, p < 0.0001$), higher for losses-only gambles ($t = 14.79, p < 0.0001$), and higher for mixed gambles ($t = 6.79, p < 0.0001$)). This is because the effect of attitude to risk (risk-aversion for gains and risk-seeking for losses) that is normally captured by the utility function is transferred to the parameter β , yielding a lower weighting function for gains-only gambles and a higher probability weighting function for losses-only and mixed gambles.

Analysis of Choice Patterns

The data collected in the experiments is not sufficient to estimate parameters at the individual level. However, ranking the tasks from high to low probability of the best outcome (e.g. Table 2-2), the choice patterns (Brooks et al., 2013; Harless and Camerer, 1994) of each individual was studied. Tasks 5, 9, 12 in Table 2-2 were eliminated for this part of the analysis so that the patterns could be compared to CPT predictions (for various choices of parameters).

Six different choice patterns were identified — 1) All risk-averse choices (A...); 2) All risk-seeking choices (B...); 3) Risk-averse for large probabilities of the better outcome and switch to risk-seeking choices at some threshold probability level (A...B...); 4) Risk-seeking for large probabilities of the better outcome and switch to risk-averse choices at some threshold probability level (B...A...); 5) Risk-seeking for extreme probabilities of the better outcome and risk-seeking otherwise (B...A...B...); 6) Risk-averse for extreme probabilities of the better outcome and risk-averse otherwise (A...B...A...). The first four patterns are consistent with CPT predictions (for some parameters of the probability weighting and utility functions, depending on the threshold probability level for switching attitude to risk). The proportion of choices for each pattern are shown in Table 2-11.

Table 2-11: Choice patterns examined and the number of choices in each pattern.

Choice Pattern	Gains (n=102)		Losses (n=99)		Mixed (n=103)	
	DWC	DWOC	DWC	DWOC	DWC	DWOC
A...	27	17	19	8	17	5
B...	8	1	15	2	7	3
A...B...	15	22	25	9	16	23
B...A...	12	2	16	3	14	4
B...A...B...	5	4	2	7	6	9
A...B...A...	8	7	7	13	7	10
Total classified	75	53	84	42	67	54
Consistent with CPT	62	42	75	22	54	35

Nearly half of the observations could be classified into one of the six patterns. The relationship between prior probabilities and risk taking was found to be more consistent in the DWC condition than in the DWOC condition — more than half of the participants' choices can be explained using a CPT framework in the DWC condition whereas for the DWOC condition it was less than 40%.

Appendix 2-2 Accompanying tables

Table 2-12: Pairs of gambles for the 15 choice tasks. x and y denote the payoffs for the three payoff domains. For gains-only gambles, $x = \text{£}1500$; $y = \text{£}0$; for losses-only gambles, $x = \text{£}0$; $y = -\text{£}1500$; and for mixed gambles, $x = \text{£}1500$; $y = -\text{£}1500$.

Task	DWC A		DWC B		DWOC A			DWOC B		
	$P(y)$	$P(x)$	$P(y)$	$P(x)$	$P(2y)$	$P(y+x)$	$P(2x)$	$P(2y)$	$P(y+x)$	$P(2x)$
1	0.2	0.8	0.25	0.75	0	0.25	0.75	0.01	0.23	0.76
2	0.2	0.8	0.5	0.5	0	0.5	0.5	0.06	0.38	0.56
3	0.2	0.8	0.8	0.2	0	0.8	0.2	0.12	0.56	0.32
4	0.2	0.8	0.95	0.05	0	0.95	0.05	0.15	0.65	0.2
5	0.2	0.8	1	0	0	1	0	0.16	0.68	0.16
6	0.25	0.75	0.5	0.5	0.03	0.5	0.47	0.08	0.4	0.52
7	0.25	0.75	0.8	0.2	0.04	0.77	0.19	0.15	0.55	0.3
8	0.25	0.75	0.95	0.05	0.05	0.9	0.05	0.19	0.62	0.19
9	0.25	0.75	1	0	0.05	0.95	0	0.2	0.65	0.15
10	0.5	0.5	0.8	0.2	0.24	0.62	0.14	0.3	0.5	0.2
11	0.5	0.5	0.95	0.05	0.28	0.68	0.04	0.37	0.5	0.13
12	0.5	0.5	1	0	0.3	0.7	0	0.4	0.5	0.1
13	0.8	0.2	0.95	0.05	0.57	0.41	0.02	0.6	0.35	0.05
14	0.8	0.2	1	0	0.6	0.4	0	0.64	0.32	0.04
15	0.95	0.05	1	0	0.75	0.25	0	0.76	0.23	0.01

Table 2-13: Number of respondents excluded.

	Gains	Losses	Mixed
Total number of responses	112	114	109
Deleted: Low clarity in practice questions	1	2	3
Time taken low + reversals	1	0	1
Violate dominance	8	13	2
Sample Size	102	99	103

Table 2-14: Inconsistency rates for each of the two randomly selected questions that were repeated (for both conditions).

	Gains			Losses			Mixed		
	Q1 only	Q2 only	Both	Q1 only	Q2 only	Both	Q1 only	Q2 only	Both
DWOC	26%	21%	8%	24%	24%	7%	21%	15%	5%
DWC	14%	12%	2%	10%	22%	8%	19%	19%	7%

Table 2-15: Proportion of riskier choices in each condition by demographics.

	Gains				Losses				Mixed			
	DWOC	DWC	P value	n	DWOC	DWC	P value	n	DWOC	DWC	P value	n
Male	0.33	0.42	0.00	49	0.66	0.52	0.00	45	0.54	0.56	0.55	50
Female	0.28	0.31	0.08	50	0.69	0.55	0.00	53	0.48	0.36	0.00	50
18-22 yrs	0.43	0.46	0.49	16	0.62	0.24	0.00	15	0.49	0.45	0.37	21
23-27 yrs	0.28	0.33	0.01	63	0.68	0.54	0.00	64	0.54	0.43	0.00	57
28+ yrs	0.26	0.39	0.00	20	0.72	0.72	1.00	20	0.46	0.54	0.02	22
EM Nation	0.33	0.29	0.13	44	0.67	0.46	0.00	36	0.51	0.37	0.00	37
Dev. Nation	0.29	0.42	0.00	50	0.68	0.57	0.00	63	0.51	0.56	0.06	56
Part/No Spons.	0.26	0.26	0.94	36	0.70	0.62	0.00	38	0.51	0.37	0.11	37
Spons.	0.33	0.43	0.00	61	0.66	0.46	0.00	59	0.51	0.56	0.01	61

Table 2-16: Difficulty (end of each section) and Clarity (end of practice section) of questions.

Difficulty	Gains		Losses		Mixed	
	DWOC	DWC	DWOC	DWC	DWOC	DWC
Extremely Difficult (1)	0	0	2	1	0	1
Fairly Difficult (2)	1	9	4	1	1	1
Somewhat Difficult (3)	9	12	14	12	13	13
Somewhat Easy (4)	24	18	15	21	32	35
Fairly Easy (5)	49	46	48	43	41	40
Extremely Easy (6)	19	17	16	21	16	13
Average	4.75	4.49	4.53	4.69	4.56	4.47
p value (2 tailed ttest)	0.01		0.20		0.25	

Clarity	Gains		Losses		Mixed	
	DWOC	DWC	DWOC	DWC	DWOC	DWC
Extremely Unclear (1)	1	3	1	2		
Fairly Unclear (2)	2	6	1	1	1	1
Somewhat Unclear (3)	0	10	3	14	4	9
Somewhat Clear (4)	15	15	11	17	12	15
Fairly Clear (5)	40	41	48	41	47	52
Extremely Clear (6)	44	27	35	24	39	26
Average	5.2	4.6	5.1	4.7	5.2	4.9
	0.00		0.01		0.00	

Table 2-17: Effect of the order in which the two conditions appeared and the side on which gamble A appeared (logistic regression).

DWOC	Gains			Losses			Mixed		
	Odds Ratio	Std. Err.	p value	Odds Ratio	Std. Err.	p value	Odds Ratio	Std. Err.	p value
Order (DWOC first)	0.69	0.17	0.00	0.88	0.13	0.27	1.35	0.08	0.00
Left right (A on left)	0.92	0.10	0.44	0.97	0.11	0.82	1.18	0.12	0.11
DWC									
Order (DWC first)	1.44	0.15	0.00	1.65	0.18	0.00	1.72	0.18	0.00
Left right (A on left)	0.93	0.10	0.48	1.03	0.11	0.80	1.15	0.12	0.18

$R^2 \leq 0.01$ for all domains and conditions

Table 2-18: Proportion of riskier choices corresponding to the order in which the two conditions appeared and the side on which gamble A appeared.

DWC	Gains			Losses			Mixed		
	% riskier choices	Std. Err.	p val Chi ² (1)	% riskier choices	Std. Err.	p val Chi ² (1)	% riskier choices	Std. Err.	p val (Chi2 1)
DWOC 1 st	26%	0.02	0.00	67%	0.02	0.26	55%	0.02	0.00
DWOC 2 nd	33%	0.02		70%	0.02		47%	0.02	
A on right	31%	0.02	0.53	68%	0.02	0.80	49%	0.02	0.11
A on left	29%	0.02		68%	0.02		53%	0.02	
DWC									
DWC 1 st	41%	0.02	0.00	60%	0.02	0.00	53%	0.02	0.00
DWC 2 nd	32%	0.02		48%	0.02		40%	0.02	
A on right	38%	0.02	0.51	53%	0.02	0.91	45%	0.02	0.17
A on left	36%	0.02		53%	0.02		48%	0.02	

2.8 References for Chapter 2

1. Agarwal, S., Montibeller, G., 2014. A Quantitative Method for Measuring the Value of Operational Risk Mitigations.
2. Anderson, R., Fraser, D., 2000. Corporate control, bank risk taking, and the health of the banking industry. *J. Bank. Finance* 24, 1383–1398.
3. Brandstätter, E., Kühberger, A., Schneider, F., 2002. A cognitive-emotional account of the shape of the probability weighting function. *J. Behav. Decis. Mak.* 15, 79–100.
4. Brandstätter, E., Schwarzenberger, H., 2001. Beyond the gambling paradigm: Internal controllability in decision-making. *Psychol. Rep.* 89, 259–266.
5. Brooks, P., Peters, S., Zank, H., 2013. Risk behavior for gain, loss, and mixed prospects. *Theory Decis.* 1–30.
6. Brooks, P., Zank, H., 2005. Loss averse behavior. *J. Risk Uncertain.* 31, 301–325.
7. Bryant, P., Dunford, R., 2008. The Influence of Regulatory Focus on Risky Decision-Making. *Appl. Psychol.* 57, 335–359.
8. Byrnes, J., Miller, D., Schafer, W., 1999. Gender differences in risk taking: A meta-analysis. *Psychol. Bull.* 125, 367.
9. Camerer, C., 1989. An experimental test of several generalized utility theories. *J. Risk Uncertain.* 2, 61–104.
10. Camerer, C., 1992. Recent tests of generalizations of expected utility theory, in: *Utility Theories: Measurements and Applications*. Springer, pp. 207–251.
11. Camerer, C., 2004. Prospect theory in the wild: Evidence from the field. Colin F Camerer George Loewenstein Matthew Rabin Eds *Adv. Behav. Econ.* 148–161.
12. Camerer, C., Ho, T.-H., 1994. Violations of the betweenness axiom and nonlinearity in probability. *J. Risk Uncertain.* 8, 167–196.

13. Carbone, E., Hey, J., 1994. Estimation of expected utility and non-expected utility preference functionals using complete ranking data, in: *Models and Experiments in Risk and Rationality*. Springer, pp. 119–139.
14. Cavagnaro, D., Pitt, M., Gonzalez, R., Myung, J., 2013. Discriminating among probability weighting functions using adaptive design optimization. *J. Risk Uncertain.* 47, 255–289.
15. Charness, G., Gneezy, U., 2012. Strong evidence for gender differences in risk taking. *J. Econ. Behav. Organ.* 83, 50–58.
16. Chau, A., Phillips, J., 1995. Effects of perceived control upon wagering and attributions in computer blackjack. *J. Gen. Psychol.* 122, 253.
17. Courtney, H., 2001. *20/20 Foresight: Crafting Strategy in an Uncertain World*. Harvard Business School Press, Boston, MA.
18. Crowe, E., Higgins, E.T., 1997. Regulatory Focus and Strategic Inclinations: Promotion and Prevention in Decision-Making. *Organ. Behav. Hum. Decis. Process.* 69, 117–132.
19. Fennema, H., Wakker, P., 1997. Original and cumulative prospect theory: a discussion of empirical differences. *J. Behav. Decis. Mak.* 10, 53–64.
20. Fung, M., 2014. Cumulative Prospect Theory and Managerial Incentives for Fraudulent Financial Reporting. *Contemp. Account. Res.*
21. Goldstein, W., Einhorn, H., 1987. Expression theory and the preference reversal phenomena. *Psychol. Rev.* 94, 236.
22. Gonzalez, R., Wu, G., 1999. On the Shape of the Probability Weighting Function. *Cognit. Psychol.* 38, 129–166.
23. Goodie, A., 2003. The effects of control on betting: Paradoxical betting on items of high confidence with low value. *J. Exp. Psychol. Learn. Mem. Cogn.* 29, 598–610.
24. Goodie, A., Young, D., 2007. The skill element in decision making under uncertainty: Control or competence? *Judgm. Decis. Mak.* 2, 189–203.
25. Harless, D., Camerer, C., 1994. The Predictive Utility of Generalized Expected Utility Theories. *Econometrica* 62, 1251–1289.
26. Heath, C., Tversky, A., 1991. Preference and belief: Ambiguity and competence in choice under uncertainty. *J. Risk Uncertain.* 4, 5–28.
27. Higgins, E.T., 1996a. Ideals, oughts, and regulatory focus. *Psychol. Action Link. Cogn. Motiv. Behav.* 91–114.
28. Higgins, E.T., 1996b. Emotional Experiences: The Pains and Pleasures of Distinct Regulatory Systems. *Emot. Interdiscip. Perspect.* 203.
29. Higgins, E.T., 1997. Beyond pleasure and pain. *Am. Psychol.* 52, 1280.
30. Higgins, E.T., 1998. Promotion and prevention: Regulatory focus as a motivational principle. *Adv. Exp. Soc. Psychol.* 30, 1–46.
31. Horswill, M., McKenna, F., 1999. The Effect of Perceived Control on Risk Taking. *J. Appl. Soc. Psychol.* 29, 377–391.

32. Idson, L.C., Liberman, N., Higgins, E.T., 2000. Distinguishing Gains from Nonlosses and Losses from Nongains: A Regulatory Focus Perspective on Hedonic Intensity. *J. Exp. Soc. Psychol.* 36, 252–274.
33. Isen, A., Patrick, R., 1983. The effect of positive feelings on risk taking: When the chips are down. *Organ. Behav. Hum. Perform.* 31, 194–202.
34. Kluger, A., Stephan, E., Ganzach, Y., HersHKovitz, M., 2004. The effect of regulatory focus on the shape of probability-weighting function: Evidence from a cross-modality matching method. *Organ. Behav. Hum. Decis. Process.* 95, 20–39.
35. Kocher, M., Pahlke, J., Trautmann, S., 2013. Tempus Fugit: Time Pressure in Risky Decisions. *Manag. Sci.*
36. Langer, E., 1975. The illusion of control. *J. Pers. Soc. Psychol.* 32, 311–328.
37. Lefebvre, M., Vieider, F., 2014. Risk taking of executives under different incentive contracts: Experimental evidence. *J. Econ. Behav. Organ.* 97, 27–36.
38. Lerner, J., Keltner, D., 2001. Fear, anger, and risk. *J. Pers. Soc. Psychol.* 81, 146.
39. Li, K., 2011. Preference towards control in risk taking: Control, no control, or randomize? *J. Risk Uncertain.* 43, 39–63.
40. Loewenstein, G., Weber, E., Hsee, C., Welch, N., 2001. Risk as feelings. *Psychol. Bull.* 127, 267–286.
41. March, J., Shapira, Z., 1987. Managerial perspectives on risk and risk taking. *Manag. Sci.* 33, 1404–1418.
42. Matheson, D., Matheson, J., 2005. Describing and Valuing Interventions That Observe or Control Decision Situations. *Decis. Anal.* 2, 165–81.
43. Matheson, J., 1990. Using influence diagrams to value information and control, in: R.M.Oliver & J.Q. Smith, *Influence Diagrams, Belief Nets and Decision Analysis*. John Wiley & Sons Ltd, pp. 25–48.
44. Olsen, R., 1997. Prospect theory as an explanation of risky choice by professional investors: Some evidence. *Rev. Financ. Econ.* 6, 225–232.
45. Payne, J., 2005. It is Whether You Win or Lose: The Importance of the Overall Probabilities of Winning or Losing in Risky Choice. *J. Risk Uncertain.* 30, 5–19.
46. Payne, J., Bettman, J., Schkade, D., 1999. Measuring Constructed Preferences: Towards a Building Code. *J. Risk Uncertain.* 19, 243–270–270.
47. Prelec, D., 1998. The Probability Weighting Function. *Econometrica* 66, 497–527.
48. Presson, P., Benassi, V., 1996. Illusion of control: A meta-analytic review. *J. Soc. Behav. Personal.* 11, 493–510.
49. Rosenzweig, P., 2013. What Makes Strategic Decisions Different. *Harv. Bus. Rev.* 89–93.
50. Rosenzweig, P., 2014. *Left Brain, Right Stuff: How Leaders Make Winning Decisions*. Public Affairs, USA.
51. Rottenstreich, Y., Hsee, C., 2001. Money, kisses, and electric shocks: On the affective psychology of risk. *Psychol. Sci.* 12, 185–190.

52. Segal, U., 1987. Some remarks on Quiggin's anticipated utility. *J. Econ. Behav. Organ.* 8, 145–154.
53. Seo, M.-G., Goldfarb, B., Barrett, L., 2010. Affect and the framing effect within individuals over time: Risk taking in a dynamic investment simulation. *Acad. Manage. J.* 53, 411–431.
54. Sitkin, S., Pablo, A., 1992. Reconceptualizing the determinants of risk behavior. *Acad. Manage. Rev.* 17, 9–38.
55. Sokol-Hessner, P., Camerer, C., Phelps, E., 2013. Emotion regulation reduces loss aversion and decreases amygdala responses to losses. *Soc. Cogn. Affect. Neurosci.* 8, 341–350.
56. Starmer, C., Sugden, R., 1989. Violations of the independence axiom in common ratio problems: An experimental test of some competing hypotheses. *Ann. Oper. Res.* 19, 79–102.
57. Stott, H., 2006. Cumulative prospect theory's functional menagerie. *J. Risk Uncertain.* 32, 101–130.
58. Thaler, R., Johnson, E., 1990. Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Manag. Sci.* 36, 643–660.
59. Thompson, S., 1981. Will it hurt less if I can control it? A complex answer to a simple question. *Psychol. Bull.* 90, 89–101.
60. Trepel, C., Fox, C., Poldrack, R., 2005. Prospect theory on the brain? Toward a cognitive neuroscience of decision under risk. *Cogn. Brain Res.* 23, 34–50.
61. Tversky, A., Kahneman, D., 1986. Rational Choice and the Framing of Decisions. *J. Bus., The Behavioral Foundations of Economic Theory* 59, S251–S278.
62. Tversky, A., Kahneman, D., 1992. Advances in prospect theory: cumulative representation of uncertainty. *J. Risk Uncertain.* 5, 297.
63. Tversky, A., Wakker, P., 1995. Risk Attitudes and Decision Weights. *Econometrica* 63, 1255–1280.
64. Vroom, V., Pahl, B., 1971. Relationship between age and risk taking among managers. *J. Appl. Psychol.* 55, 399.
65. Wakker, P., 2010. *Prospect theory: For risk and ambiguity*. Cambridge University Press Cambridge.
66. Wakker, P., Erev, I., Weber, E., 1994. Comonotonic independence: The critical test between classical and rank-dependent utility theories. *J. Risk Uncertain.* 9, 195–230.
67. Weber, E., Kirsner, B., 1997. Reasons for rank-dependent utility evaluation. *J. Risk Uncertain.* 14, 41–61.
68. Weber, E., Milliman, R., 1997. Perceived Risk Attitudes: Relating Risk Perception to Risky Choice. *Manag. Sci.* 43, 123–144.
69. Weinstein, N., 1980. Unrealistic optimism about future life events. *J. Pers. Soc. Psychol.* 39, 806.

-
70. Wiseman, R., Gomez-Mejia, L., 1998. A Behavioral Agency Model of Managerial Risk Taking. *Acad. Manage. Rev.* 23, 133–153.
 71. Wu, G., Gonzalez, R., 1996. Curvature of the Probability Weighting Function. *Manag. Sci.* 42, 1676–1690.
 72. Wu, G., Markle, A., 2008. An empirical test of gain-loss separability in prospect theory. *Manag. Sci.* 54, 1322–1335.
 73. Young, D., Goodie, A., Hall, D., 2011. Modeling the Impact of Control on the Attractiveness of Risk in a Prospect Theory Framework. *J. Behav. Decis. Mak.* 24, 47–70.

2.9 Supplementary material for Chapter 2

S 2-1: Other tables and figures


Table 2-19: Standard error and p values for Table 2-3.

Task	Gains					Losses					Mixed				
	Std err	p value (2 tail) Bin. test	Std err	p value (2 tail) Bin. test	p value McNemar test	Std err	p value (2 tail) Bin. test	Std err	p value (2 tail) Bin. test	p value McNemar test	Std err	p value (2 tail) Bin. test	Std err	p value (2 tail) Bin. test	p value McNemar test
1	0.042	0.00	0.047	0.00	0.08	0.048	0.00	0.048	0.00	0.88	0.042	0.00	0.046	0.00	0.12
2	0.039	0.00	0.047	0.00	0.01	0.048	0.00	0.050	0.16	0.23	0.040	0.00	0.048	0.01	0.01
6	0.043	0.00	0.049	0.04	0.01	0.050	0.62	0.050	0.84	0.88	0.046	0.00	0.049	0.24	0.06
3	0.032	0.00	0.044	0.00	0.00	0.050	0.69	0.051	1.00	0.77	0.041	0.00	0.048	0.05	0.00
7	0.042	0.00	0.049	0.04	0.00	0.044	0.00	0.050	0.42	0.00	0.049	0.17	0.049	0.43	0.67
4	0.036	0.00	0.047	0.00	0.00	0.046	0.00	0.050	0.31	0.00	0.041	0.00	0.048	0.02	0.01
5	0.038	0.00	0.047	0.00	0.02	0.035	0.00	0.049	0.04	0.00	0.049	0.17	0.050	1.00	0.29
8	0.035	0.00	0.048	0.00	0.00	0.045	0.00	0.050	0.62	0.01	0.049	0.32	0.049	0.55	0.18
9	0.043	0.00	0.048	0.01	0.11	0.040	0.00	0.048	0.00	0.02	0.043	0.00	0.049	0.11	0.01
10	0.046	0.00	0.049	0.06	0.14	0.043	0.00	0.050	0.13	0.01	0.050	1.00	0.050	1.00	1.00
11	0.044	0.00	0.048	0.01	0.11	0.043	0.00	0.051	1.00	0.00	0.049	0.17	0.050	1.00	0.28
12	0.050	0.37	0.049	0.28	0.89	0.033	0.00	0.049	0.02	0.00	0.037	0.00	0.049	0.24	0.00
13	0.050	0.37	0.049	0.09	0.56	0.046	0.00	0.050	0.42	0.03	0.048	0.03	0.049	0.24	0.02
14	0.049	0.06	0.049	0.04	0.00	0.040	0.00	0.050	0.23	0.00	0.037	0.00	0.050	1.00	0.00
15	0.047	0.00	0.049	0.28	0.00	0.043	0.00	0.049	0.02	0.03	0.038	0.00	0.049	0.24	0.00

Table 2-20: Gender effects for risk taking for each choice task.

Task	Gains						Losses						Mixed					
	DWOC			DWC			DWOC			DWC			DWOC			DWC		
	Male 49	Fem. 50	Sig. diff.	Male 49	Fem. 50	Sig. diff.	Male 45	Fem. 53	Sig. diff.	Male 45	Fem. 53	Sig. diff.	Male 50	Fem. 50	Sig. diff.	Male 50	Fem. 50	Sig. diff.
1	0.29	0.20	0.32	0.39	0.28	0.26	0.40	0.30	0.31	0.33	0.34	0.95	0.38	0.08	0.00	0.42	0.20	0.02
2	0.27	0.12	0.07	0.45	0.20	0.01	0.38	0.30	0.43	0.47	0.40	0.48	0.32	0.10	0.01	0.42	0.30	0.21
6	0.29	0.20	0.32	0.37	0.40	0.74	0.58	0.47	0.30	0.47	0.57	0.33	0.34	0.28	0.52	0.54	0.30	0.02
3	0.14	0.08	0.32	0.27	0.24	0.77	0.40	0.55	0.15	0.56	0.45	0.31	0.32	0.12	0.02	0.50	0.26	0.01
7	0.24	0.22	0.77	0.41	0.36	0.62	0.71	0.77	0.48	0.49	0.60	0.26	0.46	0.40	0.55	0.60	0.30	0.00
4	0.18	0.12	0.38	0.35	0.32	0.78	0.67	0.72	0.59	0.40	0.49	0.37	0.24	0.20	0.63	0.52	0.24	0.00
5	0.20	0.16	0.57	0.37	0.26	0.25	0.78	0.92	0.04	0.58	0.64	0.52	0.58	0.56	0.84	0.64	0.36	0.01
8	0.18	0.12	0.38	0.37	0.32	0.62	0.62	0.79	0.06	0.60	0.47	0.21	0.58	0.52	0.55	0.58	0.36	0.03
9	0.31	0.22	0.33	0.35	0.36	0.89	0.76	0.85	0.22	0.71	0.62	0.36	0.78	0.70	0.36	0.68	0.48	0.04
10	0.35	0.26	0.35	0.45	0.34	0.27	0.73	0.77	0.65	0.60	0.57	0.74	0.54	0.44	0.32	0.50	0.46	0.69
11	0.33	0.22	0.24	0.39	0.36	0.78	0.82	0.72	0.22	0.44	0.55	0.31	0.62	0.54	0.42	0.52	0.46	0.55
12	0.43	0.50	0.48	0.51	0.38	0.19	0.89	0.87	0.75	0.62	0.64	0.84	0.80	0.88	0.28	0.68	0.46	0.03
13	0.49	0.44	0.62	0.51	0.32	0.06	0.69	0.70	0.92	0.49	0.60	0.26	0.52	0.72	0.04	0.54	0.32	0.03
14	0.57	0.64	0.49	0.51	0.26	0.01	0.78	0.83	0.52	0.51	0.62	0.27	0.82	0.86	0.59	0.58	0.40	0.07
15	0.69	0.64	0.57	0.57	0.32	0.01	0.73	0.79	0.49	0.60	0.66	0.54	0.84	0.78	0.45	0.64	0.48	0.11
Mean	0.33	0.28	0.02	0.42	0.31	0.00	0.66	0.69	0.25	0.52	0.55	0.36	0.54	0.48	0.01	0.56	0.36	0.00

Figure 2-6: Predicted vs. observed proportion of riskier choices for each choice task for the CPT and the PVoC models in Table 2-10.



S 2-2: Experiment invitation e-mail

We would like to invite you to do a short online experiment (25 mins).

The link to the experiment is below and you can take it anytime you like, ideally, within the next week.

https://lse.qualtrics.com/SE/?SID=SV_3fUmWxmIKnyOqz3 (now defunct)

Compensation for your participation:

From all those who complete this experiment, **three respondents** will be selected at random, to play one of their answered questions for real money and each selected respondent can **earn up to £300**

(Participation in this experiment is entirely **voluntary** and your decision to not participate will have no bearing on your grades or access to LSE services).

Please note that this experiment is open only for a **limited time** and will be closed once we receive around 300 complete responses.

(Once the experiment is closed, three respondents will be selected at random and invited **via email** to play one of their answered questions for real money).

All data gathered in this experiment will be used in anonymised form for the research.

Your participation in this experiment is very valuable to our research and by participating you will help advance scientific research on decision under risk & uncertainty.

Thank you in advance for your support.

S 2-3: Sample of experiment (screenshots for one instance of gains-only gambles).

Welcome to the London School of Economics Management Science Experiment.

Thank you for volunteering to take part in this study on decision making under risk.

In return for your participation, you have a chance to **earn up to £300!**

Overview of the Experiment

This experiment consists of **two parts**. In each part there are 2 practice questions followed by around 17 questions which will be used in the research.

In all questions you will be asked to decide between two risky options. **There are no right or wrong answers in this experiment.**

This experiment should take you approximately **25 minutes** to complete. Please try to complete the experiment in one sitting and do it sincerely.

Once you are satisfied with your answer to a question, click 'Next' to **submit** it. Please note that once you submit your answer you **cannot go back** and change it. Please do not use the back browser button as this **will disrupt** the experiment.

Click Next to continue

0% 100%

Next

Terms and Conditions of this Experiment

Compensation

From all the participants who **complete** the experiment, **three** respondents will be **randomly** selected to play **one** of their answered questions (also selected at random) for **real money** and **each of them** can earn up to £300.

If you are selected to play one of your answered questions for real money, you will be paid 10% of the actual amount you win in the question. For example, if you win £1500 in the question, you will be paid £150.

- The maximum amount you can earn is £300 (if you win £3000 when you play the question).
- The minimum amount you can earn is £0 (if you win £0 when you play the question).
- You will not lose any of your own money.

The payment will be made in pounds.

Data Protection and Confidentiality

Your answers are **completely anonymous** and will only be used for academic research purposes. You can choose to drop out of the experiment at any point and your responses will be discarded.

NOTE: If you **want to be considered in the draw** to play one of your answered questions for real money, please **give your email address** at the end of the experiment when prompted. Any contact details you give will be kept **strictly confidential** and will not be used for any other purpose except to contact you if you are selected for the draw.

If you have any questions, please email **Shweta Agarwal** at s.agarwal@lse.ac.uk.

Before you begin the experiment, please give your consent to participate in this research and then click Next to continue.

Declaration of agreement for scientific use of responses in this experiment

By checking the box below, I agree that my responses in this experiment can be used for research reasons in the course of Shweta Agarwal's PhD at the London School of Economics and Political Science, UK.

Shweta Agarwal has confirmed that my anonymity is protected and I understand that I am entitled to revoke this agreement at any time.

According to LSE's data protection rules, Shweta Agarwal has affirmed appropriate analysis and saving of my responses in this experiment.

I have read and understand the terms of this experiment and agree to participate sincerely in this experiment.

0% 100%

Next

Welcome to Part 1 of 2

Click 'Next' to see the Practice Questions.

0% 100%

Next

Practice Question 1

Imagine you are in the following situation.

You are presented with **two** bags, A and B, each of which has 100 balls. The bags contain **red** and **blue** balls.

Suppose you can draw a ball, once, from **each** bag and the **total pay-off is calculated after you draw a ball from each of them**.

For each bag, if you draw a **blue** ball you win £1500, else nothing. The number of **blue** and **red** balls in each bag is shown below:

Bag A	Bag B
40 Blue balls (£1500)	70 Blue balls (£1500)
60 Red balls (£0)	30 Red balls (£0)

Suppose that before you draw a ball from **each** bag, you can replace 20 **red** balls (£0) with 20 **blue** balls (£1500) in **any one** of the bags.

Please indicate in which bag you would prefer to replace 20 red balls with blue balls.

Replace in Bag A



Replace in Bag B



0% 100%

Next

Practice Question 2

Imagine you are in the following situation.

You are presented with **two** bags, A and B, each of which has 100 balls. The bags contain **red** and **blue** balls.

Suppose you can draw a ball from **each** bag and the **total pay-off is calculated after you draw a ball from each of them**.

For each bag, if you draw a **blue** ball you win £1500, else nothing. The number of **blue** and **red** balls in each bag is shown below:

Bag A	Bag B
30 Blue balls (£1500)	10 Blue balls (£1500)
70 Red balls (£0)	90 Red balls (£0)

Suppose that before you draw a ball from **each** bag, you can replace 20 **red** balls (£0) with 20 **blue** balls (£1500) in **any one** of the bags.

Please indicate in which bag you would prefer to replace 20 red balls with blue balls.

Replace in Bag A



Replace in Bag B



0% 100%

Next

We hope you feel familiar with the nature of the questions.

On the scale below, please indicate how clear you feel the questions were.

	Extremely Unclear	Fairly Unclear	Somewhat Unclear	Somewhat Clear	Fairly Clear	Extremely Clear
I feel the questions were	<input type="radio"/>					

0%  100%

Next

We will now present you with 17 questions which will be used in the research.

The instructions for all the questions are same as in the Practice Questions.

Please try to answer as honestly as you can.

Remember, there are **no right or wrong answers** and you may be selected to **play any one of your answered questions for real money.**

Click next to continue.

0%  100%

Next

The instructions below are same as the instructions in the Practice Questions.

Imagine you are in the following situation.

You are presented with **two** bags, A and B, each of which has 100 balls. The bags contain **red** and **blue** balls.

Suppose you can draw a ball, once, from **each** bag and the **total pay-off is calculated after you draw a ball from each of them.**

For each bag, if you draw a **blue** ball you win £1500, else nothing. The number of **blue** and **red** balls in each bag is shown below:

Bag A	Bag B
50 Blue balls (£1500)	5 Blue balls (£1500)
50 Red balls (£0)	95 Red balls (£0)

Suppose that before you draw a ball from **each** bag, you can replace 20 **red** balls (£0) with 20 **blue** balls (£1500) in **any one** of the bags.

Please indicate in which bag you would prefer to replace 20 red balls with blue balls.

Replace in Bag A Replace in Bag B

0%  100%

Next

The instructions below are same as the instructions in the Practice Questions

Imagine you are in the following situation.

You are presented with two bags, A and B, each of which has 100 balls. The bags contain red and blue balls.

Suppose you can draw a ball, once, from each bag and the total pay-off is calculated after you draw a ball from each of them.

For each bag, if you draw a blue ball you win £1500, else nothing. The number of blue and red balls in each bag is shown below:

Bag A	Bag B
20 Blue balls (£1500)	50 Blue balls (£1500)
80 Red balls (£0)	50 Red balls (£0)

Suppose that before you draw a ball from each bag, you can replace 20 red balls (£0) with 20 blue balls (£1500) in any one of the bags.

Please indicate in which bag you would prefer to replace 20 red balls with blue balls.

Replace in Bag A Replace in Bag B

0% 100%

[Next](#)

The instructions below are same as the instructions in the Practice Questions

Imagine you are in the following situation.

You are presented with two bags, A and B, each of which has 100 balls. The bags contain red and blue balls.

Suppose you can draw a ball, once, from each bag and the total pay-off is calculated after you draw a ball from each of them.

For each bag, if you draw a blue ball you win £1500, else nothing. The number of blue and red balls in each bag is shown below:

Bag A	Bag B
0 Blue balls (£1500)	75 Blue balls (£1500)
100 Red balls (£0)	25 Red balls (£0)

Suppose that before you draw a ball from each bag, you can replace 20 red balls (£0) with 20 blue balls (£1500) in any one of the bags.

Please indicate in which bag you would prefer to replace 20 red balls with blue balls.

Replace in Bag A Replace in Bag B

0% 100%

[Next](#)

The instructions below are same as the instructions in the Practice Questions

Imagine you are in the following situation.

You are presented with two bags, A and B, each of which has 100 balls. The bags contain red and blue balls.

Suppose you can draw a ball, once, from each bag and the total pay-off is calculated after you draw a ball from each of them.

For each bag, if you draw a blue ball you win £1500, else nothing. The number of blue and red balls in each bag is shown below:

Bag A	Bag B
5 Blue balls (£1500)	80 Blue balls (£1500)
95 Red balls (£0)	20 Red balls (£0)

Suppose that before you draw a ball from each bag, you can replace 20 red balls (£0) with 20 blue balls (£1500) in any one of the bags.

Please indicate in which bag you would prefer to replace 20 red balls with blue balls.

Replace in Bag A Replace in Bag B

0% 100%

[Next](#)

The instructions below are same as the instructions in the Practice Questions

Imagine you are in the following situation.

You are presented with two bags, A and B, each of which has 100 balls. The bags contain red and blue balls.

Suppose you can draw a ball, once, from each bag and the total pay-off is calculated after you draw a ball from each of them.

For each bag, if you draw a blue ball you win £1500, else nothing. The number of blue and red balls in each bag is shown below:

Bag A	Bag B
80 Blue balls (£1500)	0 Blue balls (£1500)
20 Red balls (£0)	100 Red balls (£0)

Suppose that before you draw a ball from each bag, you can replace 20 red balls (£0) with 20 blue balls (£1500) in any one of the bags.

Please indicate in which bag you would prefer to replace 20 red balls with blue balls.

Replace in Bag A Replace in Bag B

0% 100%

[Next](#)

The instructions below are same as the instructions in the Practice Questions

Imagine you are in the following situation.

You are presented with two bags, A and B, each of which has 100 balls. The bags contain red and blue balls.

Suppose you can draw a ball, once, from each bag and the total pay-off is calculated after you draw a ball from each of them.

For each bag, if you draw a blue ball you win £1500, else nothing. The number of blue and red balls in each bag is shown below:

Bag A	Bag B
80 Blue balls (£1500)	50 Blue balls (£1500)
20 Red balls (£0)	50 Red balls (£0)

Suppose that before you draw a ball from each bag, you can replace 20 red balls (£0) with 20 blue balls (£1500) in any one of the bags.

Please indicate in which bag you would prefer to replace 20 red balls with blue balls.

Replace in Bag A Replace in Bag B

0% 100%

[Next](#)

The instructions below are same as the instructions in the Practice Questions

Imagine you are in the following situation.

You are presented with two bags, A and B, each of which has 100 balls. The bags contain red and blue balls.

Suppose you can draw a ball, once, from each bag and the total pay-off is calculated after you draw a ball from each of them.

For each bag, if you draw a blue ball you win £1500, else nothing. The number of blue and red balls in each bag is shown below:

Bag A	Bag B
0 Blue balls (£1500)	50 Blue balls (£1500)
100 Red balls (£0)	50 Red balls (£0)

Suppose that before you draw a ball from each bag, you can replace 20 red balls (£0) with 20 blue balls (£1500) in any one of the bags.

Please indicate in which bag you would prefer to replace 20 red balls with blue balls.

Replace in Bag A Replace in Bag B

0% 100%

[Next](#)

Congratulations! You have finished Part 1 of the experiment successfully.

On the scale below, please indicate how easily you were able to answer the questions.

	Very Difficult	Fairly Difficult	Somewhat Difficult	Somewhat Easy	Fairly Easy	Very Easy
I found the questions	<input type="radio"/>					

0% 100%

[Next](#)

Welcome to Part 2 of 2

Click 'Next' to see the Practice Questions.

0% 100%

[Next](#)

Practice Question 1

Imagine you are in the following situation.

You are presented with **two** bags, A and B, each of which has 100 balls. The bags contain **red**, **blue** and **green** balls.

For each bag, if you draw a **green** ball you win £3000, **blue** ball you win £1500, else nothing.

The number of **red**, **blue** and **green** balls in each bag is shown below:

Bag A	Bag B
36 Green balls (£3000)	42 Green balls (£3000)
58 Blue balls (£1500)	46 Blue balls (£1500)
6 Red balls (£0)	12 Red balls (£0)

Suppose you can draw **one** ball, only once, from any one of the bags.

Please indicate from which bag you would prefer to draw a ball.

Draw from Bag A



Draw from Bag B



0%  100%

Next

Practice Question 2

Imagine you are in the following situation.

You are presented with **two** bags, A and B, each of which has 100 balls. The bags contain **red**, **blue** and **green** balls.

For each bag, if you draw a **green** ball you win £3000, **blue** ball you win £1500, else nothing.

The number of **red**, **blue** and **green** balls in each bag is shown below:

Bag A	Bag B
9 Green balls (£3000)	5 Green balls (£3000)
42 Blue balls (£1500)	50 Blue balls (£1500)
49 Red balls (£0)	45 Red balls (£0)

Suppose you can draw **one** ball, only once, from any one of the bags.

Please indicate from which bag you would prefer to draw a ball.

Draw from Bag A



Draw from Bag B



0%  100%

Next

We hope you feel familiar with the nature of the questions.

On the scale below, please indicate how clear you feel the questions were.

	Extremely Unclear	Fairly Unclear	Somewhat Unclear	Somewhat Clear	Fairly Clear	Extremely Clear
I feel the questions were	<input type="radio"/>					

0%  100%

Next

We will now present you with 19 questions which will be used in the research.

The instructions for all the questions are same as in the Practice Questions.

Please try to answer as honestly as you can.

Remember, there are **no right or wrong answers** and you may be selected to **play any one of your answered questions for real money.**

Click next to continue.

0%  100%

Next

The instructions below are same as the instructions in the Practice Questions.

Imagine you are in the following situation.

You are presented with **two** bags, A and B, each of which has 100 balls. The bags contain **red**, **blue** and **green** balls.

For each bag, if you draw a **green** ball you win £3000, **blue** ball you win £1500, else nothing.

The number of **red**, **blue** and **green** balls in each bag is shown below:

Bag A	Bag B
0 Green balls (£3000)	4 Green balls (£3000)
40 Blue balls (£1500)	32 Blue balls (£1500)
60 Red balls (£0)	64 Red balls (£0)

Suppose you can draw **one** ball, only once, from **any one** of the bags.

Please indicate from which bag you would prefer to draw a ball.

Draw from Bag A Draw from Bag B

0%  100%

Next

The instructions below are same as the instructions in the Practice Questions.

Imagine you are in the following situation.

You are presented with two bags, A and B, each of which has 100 balls. The bags contain red, blue and green balls.

For each bag, if you draw a green ball you win £3000, blue ball you win £1500, else nothing.

The number of red, blue and green balls in each bag is shown below:

Bag A	Bag B
5 Green balls (£3000)	20 Green balls (£3000)
95 Blue balls (£1500)	65 Blue balls (£1500)
0 Red balls (£0)	15 Red balls (£0)

Suppose you can draw one ball, only once, from any one of the bags.

Please indicate from which bag you would prefer to draw a ball.

Draw from Bag A Draw from Bag B

0% 100%

[Next](#)

The instructions below are same as the instructions in the Practice Questions.

Imagine you are in the following situation.

You are presented with two bags, A and B, each of which has 100 balls. The bags contain red, blue and green balls.

For each bag, if you draw a green ball you win £3000, blue ball you win £1500, else nothing.

The number of red, blue and green balls in each bag is shown below:

Bag A	Bag B
32 Green balls (£3000)	20 Green balls (£3000)
66 Blue balls (£1500)	80 Blue balls (£1500)
12 Red balls (£0)	0 Red balls (£0)

Suppose you can draw one ball, only once, from any one of the bags.

Please indicate from which bag you would prefer to draw a ball.

Draw from Bag A Draw from Bag B

0% 100%

[Next](#)

The instructions below are same as the instructions in the Practice Questions.

Imagine you are in the following situation.

You are presented with two bags, A and B, each of which has 100 balls. The bags contain red, blue and green balls.

For each bag, if you draw a green ball you win £3000, blue ball you win £1500, else nothing.

The number of red, blue and green balls in each bag is shown below:

Bag A	Bag B
5 Green balls (£3000)	2 Green balls (£3000)
35 Blue balls (£1500)	41 Blue balls (£1500)
60 Red balls (£0)	57 Red balls (£0)

Suppose you can draw one ball, only once, from any one of the bags.

Please indicate from which bag you would prefer to draw a ball.

Draw from Bag A Draw from Bag B

0% 100%

[Next](#)

The instructions below are same as the instructions in the Practice Questions.

Imagine you are in the following situation.

You are presented with two bags, A and B, each of which has 100 balls. The bags contain red, blue and green balls.

For each bag, if you draw a green ball you win £3000, blue ball you win £1500, else nothing.

The number of red, blue and green balls in each bag is shown below:

Bag A	Bag B
19 Green balls (£3000)	5 Green balls (£3000)
62 Blue balls (£1500)	90 Blue balls (£1500)
19 Red balls (£0)	5 Red balls (£0)

Suppose you can draw one ball, only once, from any one of the bags.

Please indicate from which bag you would prefer to draw a ball.

Draw from Bag A Draw from Bag B

0% 100%

[Next](#)

The instructions below are same as the instructions in the Practice Questions.

Imagine you are in the following situation.

You are presented with two bags, A and B, each of which has 100 balls. The bags contain red, blue and green balls.

For each bag, if you draw a green ball you win £3000, blue ball you win £1500, else nothing.

The number of red, blue and green balls in each bag is shown below:

Bag A	Bag B
50 Green balls (£3000)	56 Green balls (£3000)
50 Blue balls (£1500)	38 Blue balls (£1500)
0 Red balls (£0)	6 Red balls (£0)

Suppose you can draw one ball, only once, from any one of the bags.

Please indicate from which bag you would prefer to draw a ball.

Draw from Bag A Draw from Bag B

0% 100%

[Next](#)

The instructions below are same as the instructions in the Practice Questions.

Imagine you are in the following situation.

You are presented with two bags, A and B, each of which has 100 balls. The bags contain red, blue and green balls.

For each bag, if you draw a green ball you win £3000, blue ball you win £1500, else nothing.

The number of red, blue and green balls in each bag is shown below:

Bag A	Bag B
47 Green balls (£3000)	52 Green balls (£3000)
50 Blue balls (£1500)	40 Blue balls (£1500)
3 Red balls (£0)	8 Red balls (£0)

Suppose you can draw one ball, only once, from any one of the bags.

Please indicate from which bag you would prefer to draw a ball.

Draw from Bag A Draw from Bag B

0% 100%

[Next](#)

The instructions below are same as the instructions in the Practice Questions.

Imagine you are in the following situation.

You are presented with two bags, A and B, each of which has 100 balls. The bags contain red, blue and green balls.

For each bag, if you draw a green ball you win £3000, blue ball you win £1500, else nothing.

The number of red, blue and green balls in each bag is shown below:

Bag A	Bag B
75 Green balls (£3000)	76 Green balls (£3000)
25 Blue balls (£1500)	23 Blue balls (£1500)
0 Red balls (£0)	1 Red balls (£0)

Suppose you can draw one ball, only once, from any one of the bags.

Please indicate from which bag you would prefer to draw a ball.

Draw from Bag A Draw from Bag B

0% 100%

[Next](#)

The instructions below are same as the instructions in the Practice Questions.

Imagine you are in the following situation.

You are presented with two bags, A and B, each of which has 100 balls. The bags contain red, blue and green balls.

For each bag, if you draw a green ball you win £3000, blue ball you win £1500, else nothing.

The number of red, blue and green balls in each bag is shown below:

Bag A	Bag B
40 Green balls (£3000)	40 Green balls (£3000)
45 Blue balls (£1500)	50 Blue balls (£1500)
15 Red balls (£0)	10 Red balls (£0)

Suppose you can draw one ball, only once, from any one of the bags.

Please indicate from which bag you would prefer to draw a ball.

Draw from Bag A Draw from Bag B

0% 100%

[Next](#)

Congratulations! You have finished Part 2 of the experiment successfully.

On the scale below, please indicate how easily you were able to answer the questions.

	Very Difficult	Fairly Difficult	Somewhat Difficult	Somewhat Easy	Fairly Easy	Very Easy
I found the questions	●	●	●	●	●	●

0%  100%

[Next](#)

This question is optional but your response will be highly appreciated and helpful in this research.

Please tell us something about any decision rules you used in this experiment. If you used different decision rules in each part of this experiment, then please indicate this in your response.

0%  100%

[Next](#)

You are nearly done!

Before you finish this experiment, we would like to ask you a few (optional) questions about your self.

Click next to continue.

0%  100%

[Next](#)

What is the degree that you are currently studying?

- Undergraduate degree (3rd year)
- Postgraduate degree
- MPhil/PhD
- Other

What is the title of the program are you studying?

How are you funding your education? Select all that apply.

- I have full scholarship
- I have partial scholarship
- I am sponsored by my family
- I have taken a loan
- I am self funded
- I work part-time
- Other

Do you have any work experience?

0%  100%

[Next](#)

What is your gender?

- Male
- Female
- Other

What is your age?

- 18-22 years
- 23-27 years
- 28-31 years
- 32 years or above

What is your nationality?

0%  100%

[Next](#)

We assure you that all your responses have been recorded anonymously.

Please enter your email address below if you wish to be considered for the draw to play one of the questions for real money.

(Your email will be kept **strictly confidential** and only used to contact you if you are selected to play for real money).

0%  100%

Next

We welcome any feedback you have on this experiment. Please enter your comments in the box below.

0%  100%

Next

This experiment is now over.

Thank you for your participation.

We will contact you on the email address you provided, if you are selected to play one of the questions for real money.

If you have any questions about this experiment, please email **Shweta Agarwal** at s.agarwal@lse.ac.uk.

0%  100%

S 2-4: Raffle winner announcements

Dear Participants,

Thank you for participating in the recent Online Decision Science Experiment.

Congratulations!

You have been selected to play one of your answered questions for real money and can win upto £300 when you play one of the questions.

You will need to come to the London School of Economics on Friday, 1st March at 2.30 pm to play one of your answered questions for real money.

The 'lottery session' will last approximately 30 mins.

- First, you will draw a slip from a bag which will determine the question you will play.

- Then you will play the selected question by drawing from one (or two) bags that contain coloured balls.

The payment will be made by cheque (issued by LSE) and sent to you by post. You should receive the cheque within one month of playing the question.

Please reply to this email to confirm that you can attend the lottery session. If you are unable to make the time or cannot attend in person, then please let us know ASAP.

Please note that if we do not hear from you by Tuesday 26th February 2013, you will be disqualified from this opportunity to play one of your questions for real money and we will randomly select another participant.

We look forward to hearing from you.

S 2-5: Experiment closed e-announcement

Dear Participant,

Thank you for participating in the recent Online Decision Science Experiment. We have now closed the experiment and randomly selected three participants to play one of their questions for real money. (The random selection was done using a computer in the presence of a neutral witness who was completely unacquainted with the study).

Unfortunately, you have not been selected to play one of your answered questions for real money. We highly value your contribution in this study to help advance research on decision making under risk. We will notify you when the findings on this study are published. Thank you once again for your time.

S 2-6: Procedure and instruction for the raffle session

Instructions:

Thank you for participating in the experiment and attending the lottery session today.

Review of Experiment

In the experiment you saw two types of questions:

1. Bags with 2 colours of balls: You had to choose the bag in which you will replace 20 red balls (£0) with yellow balls (£1500) before drawing one ball, once, from each bag.

2. **Bags with 3 colours of balls:** You had to choose from which bag you will draw one ball, once. The payoff to you would be £0, £1500, £3000 depending on the colour of the ball you draw.

You were told you would be paid **£10% of your total payoff** if you are invited to play one of the questions for real money.

The table on the next page lists all the questions you saw, the choices you made and the final contents of the bags (from which you will draw a ball) based on your decision.

About the session

This session will be conducted in two parts:

1. One question will be randomly selected by the computer in your presence.
2. We will count the number of coloured m&ms (instead of balls) in the selected question and mix them in a bag. You will then draw one m&m from **each bag** in the selected question.

Payment

The payoffs for each colour m&m you draw and the corresponding amount that will be paid is as follows:

Colour of m&m drawn	Payoff shown in experiment	Amount that will be paid
Red	£ 0	£ 0
Yellow	£ 1500	£ 150
Green	£ 3000	£ 300

You will need to provide your **name** and **address** and a cheque for the total amount you earn will be sent to you by post.

Please let us know if you any questions. We will begin the session when you are ready.

Table 2-21: Choices of a sample participant selected for raffle (Draw/Modify from a bag)

Q. #	List of Questions (number of balls in each bag) and your Choices							Contents of the bags based on your choices from which you will draw a ball					
	Bag A =1			Bag B =2			Bag Selected	Bag A =1			Bag B =2		
	R	Y	G	R	Y	G		R	Y	G	R	Y	G
£0	£1500	£3000	£0	£1500	£3000	£0	£1500	£3000					
1	20	80		25	75		1	0	100		25	75	
2	20	80		80	20		1	0	100		80	20	
3	20	80		100	0		1	0	100		100	0	
4	25	75		80	20		1	5	95		80	20	
5	25	75		100	0		1	5	95		100	0	
6	50	50		95	5		1	30	70		95	5	
7	80	20		95	5		1	60	40		95	5	
8	95	5		100	0		2	95	5		80	20	
9	20	80		50	50		2	20	80		30	70	
10	20	80		95	5		1	0	100		95	5	
11	25	75		50	50		2	25	75		30	70	
12	25	75		95	5		1	5	95		95	5	
13	50	50		80	20		1	30	70		80	20	
14	50	50		100	0		1	30	70		100	0	
15	80	20		100	0		1	60	40		100	0	
16	0	25	75	1	23	76	1	0	25	75			
17	0	80	20	12	56	32	1	0	80	20			
18	0	100	0	16	68	16	1	0	100	0			
19	4	77	19	15	55	30	1	4	77	19			
20	5	95	0	20	65	15	2				20	65	15
21	28	68	4	37	50	13	1	28	68	4			
22	57	41	2	60	35	5	2				60	35	5
23	75	25	0	76	23	1	2				76	23	1
24	0	50	50	6	38	56	1	0	50	50			
25	0	95	5	15	65	20	1	0	95	5			
26	3	50	47	8	40	52	1	3	50	47			
27	5	90	5	19	62	19	1	5	90	5			
28	24	62	14	30	50	20	2				30	50	20
29	30	70	0	40	50	10	2				40	50	10
30	60	40	0	64	32	4	2				64	32	4
31	20	80		25	75		1	0	100		25	75	
32	20	80		95	5		1	0	100		95	5	
33	5	95	0	20	65	15	1	5	95	0			
34	57	41	2	60	35	5	2				60	35	5
35	50	40	10	60	30	10	1	50	40	10			
36	15	45	40	10	50	40	2				10	50	40

PRELUDE TO CHAPTER 3

In the previous chapter, a simplistic view of ‘control’ was adopted — as a fixed change in probability mass from one of two outcomes to another. In the real world, however, control is often subject to judgement and the controllable uncertain state may have multiple outcomes, which can make calculating the impact of exerting control burdensome. Therefore, in complex decision situations, analytic methods which capture beliefs about control quantitatively and compute the revised probability distribution can be helpful.

Shifting the focus away from examining control in a behavioural context, this chapter studies the topic of control from a modelling perspective, in the context of Decision Analysis. It explores the limitations of existing approaches in Decision Analysis for modelling uncertainty control and expands them to enable general beliefs about how actions can influence uncertainty to be modelled as a probability revision process. Links with probability revision rules in other disciplines are established and theoretical foundations for the procedure developed, akin to the coherence criterion that supports Bayes rule as the probability revision rule for information, are proposed and defended.

It is shown that the procedure developed can lower the judgemental burden of modelling the effect of interventions on probabilities, dramatically. The theoretical foundations proposed can potentially inform the suitability of any probability revision rule for interventions, while also giving such procedures a normative anchorage that is similar to Bayes rule. It is the hope that these developments can enhance the applicability of Decision Analysis in areas of risk management.

A brief introduction to Decision Analysis and relevant modelling tools can be found in the Annexure.

Chapter 3

Stochastic Interventions: A New Probability Revision Rule to Model the Effect of Actions on Uncertainties

Shweta Agarwal*, Gilberto Montibeller*, Alec Morton**

ABSTRACT

Managers frequently deal with risk by considering uncertainty as an element of the decision problem over which they can exert control, for example lobbyists trying to exert influence over regulators, or advertisers designing marketing campaigns to improve the chance of success of a new product. Although such interventions on uncertain events can be modelled using influence diagrams, the procedure for analysing them requires eliciting probability distributions corresponding to each action, which can be judgmentally very burdensome. In this paper, we address this challenge of analysing interventions in Decision Analysis. We present a tractable probability revision procedure for interventions, which is analogous to Bayes rule for information, but does not require probabilities to be assigned to actions. Our proposed procedure builds upon an existing method to model interventions in Decision Analysis (generic controller) and links this method to a class of linear probability revision rules (imaging), proposed in Philosophy. We also ground the notion of probability revision rules for interventions in theoretical foundations, similar to the coherence criterion that supports Bayes rule, by establishing two key properties. We expose the undesirable inconsistencies that can arise when these properties are violated and prove that our proposed probability revision procedure satisfies these properties. The benefit of our proposed approach to decision analysts, in terms of requiring fewer elicitation in complex influence diagrams, is demonstrated and potential real world applications are also suggested.

Key-words: decisions under uncertainty, uncertainty management, stochastic interventions, probability revision rules.

*Management Science Group, Department of Management, London School of Economics and Political Science, London, UK. E-mail: s.agarwal@lse.ac.uk; g.montibeller@lse.ac.uk.

**Department of Management Science, Strathclyde Business School, Glasgow, UK. E-mail: alec.morton@strath.ac.uk.

3.1 Introduction


Managers are often required to make decisions under uncertainty and manage risks. In terms of Decision Analysis frameworks, the analysis of such decisions entails producing assessments of uncertainty (probabilities of events), defining the impacts of uncertain events to the decision maker (consequences) and calculating the expected utility of each decision alternative. Risks can be managed by either controlling consequences (e.g. purchasing insurance) or, when possible, by influencing the probability of events (e.g. pre-empting competition by building new capacity (Courtney et al., 1997), increasing the chance of a product being successful through advertising, lobbying to influence legislations, or affecting stock prices through mergers and acquisitions (Sudarsanam, 2003)). Evidence suggests that the latter approach, of attempting to influence uncertainty, pervades most decisions about strategic risks in the corporate world. In particular, rather than accepting risks as inherent to a situation, managers and entrepreneurs often seek to ‘change the odds’ (March and Shapira, 1987), ‘shape the future’ (Courtney, 2001) or ‘make things happen’ (Rosenzweig, 2014). This paper focusses on analysing decision problems where actions can influence the probabilities of uncertain events.

The concept of influencing uncertainty or ‘control’ has implicitly been a part of the Decision Analysis repertoire since the eighties (Howard and Matheson, 1984) and over time this concept has seen some formal developments (Matheson 1990; Heckerman and Shachter 1994; Matheson and Matheson 2005; Shachter and Heckerman 2010). More recently, the notion of ‘control’ has received attention from a number of perspectives. In the context of causal thinking, it has been suggested that explicitly modelling the effect of actions on uncertainties can facilitate analysis of strategic decisions, such as understanding decision opportunities or developing new and robust strategies (Shachter, 2012). The relevance of control when analysing the value of information in real world decisions has also been explored (von Winterfeldt et al., 2012). In related disciplines, such as behavioural decision making, formally capturing the relation between event controllability and risk taking preferences is an ongoing area of research (Agarwal, 2014; Li, 2011; Young et al., 2011). In this paper, we will further develop the notion of ‘control’ in Decision Analysis and present a way to explicitly model the effect of actions on uncertain states.

Within a Decision Analysis framework, the dependence of probabilities on actions can be conveniently modelled using graphical tools such as influence diagrams (Howard and Matheson, 1984). The influence of actions on an uncertain state is represented by an arc from the decision node to the uncertainty node. Such actions can be regarded as interventions on the ‘controllable’ state node (Matheson and Matheson 2005) and typically one of the

constituent options of such decisions is a ‘do-nothing’ option. Figure 3-1 shows an influence diagram for a firm’s decision, whether or not to advertise (denoted by the decision node $A^D = \{a^D, \neg a^D\}$), to affect the chance of obtaining a higher market share (denoted by the state node $S^M = \{s_i^M\}$) than its competitor and the corresponding payoffs (denoted by the value node V).

Figure 3-1: Influence diagram for market share problem.



In Figure 3-1, if $P(s_i^M | a^D)$ represents the probability conditional on the decision to advertise and $U(s_i^M, a^D)$ is the corresponding monetary value or utility of the decision and state, the expected utility of the decision a^D is given by:

$$EU(a^D) = \sum_i P(s_i^M | a^D) * U(s_i^M, a^D) \quad (3-1)$$

Similarly, the expected utility of the decision $\neg a^D$ can be computed. Analysing managerial decisions to control the probability of events using equation (3-1), however, poses some serious difficulties. A normative problem with using equation (3-1) is that the corresponding expected utility calculation is not consistent with the commonly assumed axioms (e.g. sure-thing principle) of standard decision theories (Savage, 1954; von Neumann and Morgenstern, 1944) that underpin utility calculations as a guide to rational choice. Another problem with modelling interventions as probabilities conditional on decisions is that, typically, the conditional probabilities $P(s_i^M | a^D)$ and $P(s_i^M | \neg a^D)$ in equation (3-1) are elicited separately for each decision. From a decision analyst’s perspective, as we will show in this paper, this procedure can be judgmentally very burdensome when multiple interventions need to be analysed, especially when the state node that can be influenced also depends on other state nodes. While various solutions to the normative problem have been proposed (see Matheson 1990; Gibbard and Harper 1978, 22; Jeffrey 1965) (discussed in Supplementary Note S 3-1), the more pressing problem for a decision analyst — of a tractable procedure for modelling the effect of actions on uncertainties — merits further research and will be addressed in this paper.

Current methods for modelling interventions circumvent the problems associated with eliciting probabilities conditional on decisions, since they focus mostly on deterministic interventions, which eliminate the uncertainty of the controllable state node entirely and set the probability of one of its constituent outcomes to 1. Such interventions are described as ‘atomic interventions’ by Pearl (1994), or ‘set decisions’ by Heckerman & Shachter (1994), or ‘perfect control’ interventions by Matheson & Matheson (2005). Managerial interventions, however, can be stochastic, i.e. interventions that alter the probability distribution of the controllable state node. Matheson & Matheson (2005) describe such interventions as ‘imperfect control’ interventions: interventions which bring about a desired outcome of the controllable state (i.e. sets its probability to 1) with some positive probability and leave the probability distribution unaffected, otherwise. The post-intervention probability of the controllable state is then described by a linear mixture of its pre-intervention distribution and the distribution that assigns probability one to the desired outcome of the uncertain state.

In the first part of this paper, we will show how Matheson & Matheson’s (2005) procedure, called the generic controller, can be regarded as a probability revision procedure for interventions which is analogous to Bayes rule for information. In particular, if S is the controllable state node, the description of an intervention along with the pre-intervention distribution of S , $P(S|\neg a)$ (associated with the ‘do nothing’ option), describes the post-intervention distribution, $P(S|a)$. We will then develop a generalization of their procedure to allow more arbitrary movements of probability mass between states. We will also prove that our generalization of the generic controller coincides with imaging (Gärdenfors, 1982) — a probability revision rule proposed as an alternative to Bayes conditionalization, for assigning probabilities to conditional statements from their unconditional probabilities. Imaging has been advocated in subject areas of causality and causal decision theory as a suitable procedure for calculating probabilities conditional on actions (Joyce, 1999; Pearl, 1994) since, unlike Bayes conditionalization, imaging has the advantage that the probability revision is described even when the prior probability of the conditioning event is 0 or 1 (as is the case for actions). We demonstrate that compared to standard procedures for analysing interventions (e.g. eliciting post-intervention probabilities directly), the probability revision approach for analysing stochastic interventions is less information intensive and requires fewer elicitations when the controllable state depends on several other uncertain states.

In the second part of the paper, we establish theoretical foundations for probability revision rules for interventions. The foundational underpinning of Bayes rule, as a probability revision procedure for modelling the effect of information on subjective probabilities, is a coherence criterion: if a decision maker’s pre-information and post-

information beliefs about probabilities of a state do not satisfy Bayes rule then the decision maker becomes vulnerable to a ‘Dutch book’ (i.e. a series of bets on the occurrence of uncertain events, which the decision maker will willingly buy or sell, but guarantee a net loss to the decision maker whatever happens) (Teller, 1973) (argument presented in Supplementary Note S 3-2). We will extend this coherence criterion of Bayes rule to probability revision rules for interventions by establishing two key properties that probability revision rules for interventions should satisfy. We show that when these properties are satisfied, if the decision maker’s pre-intervention beliefs about the controllable state are coherent (satisfy Bayes rule), then, the post-intervention beliefs will also be coherent (i.e. immune to Dutch books). We also prove that our proposed probability revision procedure satisfies these key properties of probability revision rules for interventions.

This paper makes two main contributions. Our first contribution is a tractable and general probability revision procedure that generalizes an existing approach of modelling control (generic controller (Matheson and Matheson, 2005)) in Decision Analysis and links it to probability revision rules proposed in Philosophy (imaging (Lewis, 1976; Gardenfors, 1982)). We demonstrate its usefulness to decision analysts in terms of alleviating the judgmental burden of eliciting probabilities in ‘large’ influence diagrams that contain multiple interventions on controllable states that depend on other states. A second contribution is grounding the notion of probability revision rules for interventions in theoretical foundations, similar to the coherence criterion that supports Bayes rule. In particular, we state two key properties that probability revision rules for interventions, should satisfy and expose the undesirable inconsistencies that can arise when these properties are violated. We also prove that the method we are proposing fulfils these two properties.

This paper is organized as follows. In section 2, we provide a formal definition of an influence diagram and introduce the notations. In section 3, we review the generic controller of Matheson and Matheson (2005), develop a generalization of this procedure and offer a formal proof of the equivalence between our generalization and imaging. In section 4 we state and expose some properties that probability revision rules for interventions should satisfy and prove that the probability revision procedure proposed in section 3 satisfy these properties. In the concluding section we provide a summary and discuss directions for future research. The appendices contain a discussion on eliciting beliefs about interventions, some proofs and other technical discussions.


3.2 Definition of an influence diagram and notations used in this paper

In this section, we will provide a formal definition of an influence diagram (Howard and Matheson, 1984, 2005), its inputs and how the decisions in an influence diagram are evaluated. The section also introduces the notations that are used in the rest of the paper.

An influence diagram is a directed graph $G = (\mathcal{N}, \ell)$ with nodes $\mathcal{N} = \{N^i | i = 1, 2, \dots, |\mathcal{N}|\}$, where $|\mathcal{N}|$ represents the number of nodes in the graph, directed arcs $\ell = ((N^i, N^j) | i, j = 1, 2, \dots, |\mathcal{N}|; i \neq j)$ and no directed cycles. Following Shachter (1986), the nodes in an influence diagram can be partitioned into sets of *state* nodes $\mathcal{S} = \{S^j | j = 1, 2, \dots, |\mathcal{S}|\}$ where $S^j = \{s_i^j | i = 1, 2, \dots, |S^j|\}$, *decision* nodes $\mathcal{A} = \{A^j | j = 1, 2, \dots, |\mathcal{A}|\}$ where $A^j = \{a_i^j | i = 1, 2, \dots, |A^j|\}$ and a value node $V = \{v_i | i = 1, 2, \dots, |V|\}$. We will refer to s_i^j, a_i^j, v_i as the *atomic* outcomes of the corresponding nodes.

For any subset $\mathcal{A}' \subseteq \mathcal{A}$, the vector of joint decisions $\mathcal{A}'_i = (a_{i_1}^1, a_{i_2}^2, \dots, a_{i_{|\mathcal{A}'|}}^{|\mathcal{A}'|}) \in \prod_{\mathcal{A}'} A^j$ is defined to be a *policy* corresponding to \mathcal{A}' . Similarly, for any subset $\mathcal{S}' \subseteq \mathcal{S}$ the vector $\mathcal{S}'_i = (s_{i_1}^1, s_{i_2}^2, \dots, s_{i_{|\mathcal{S}'|}}^{|\mathcal{S}'|}) \in \prod_{\mathcal{S}'} S^j$ is defined to be an *event* corresponding to \mathcal{S}' . For a given node N^K , the sets of state nodes $\mathcal{S}(N^K)$ and decision nodes $\mathcal{A}(N^K)$ are called the *parent* state and decision nodes of N^K , respectively, if there exists a directed path from every $N^i \in (\mathcal{S}(N^K) \cup \mathcal{A}(N^K))$ to N^K . The states $\mathcal{S}(N^K) \supseteq \mathcal{S}^*(N^K) = \{S^j | (S^j, N^K) \in \ell\}$ and $\mathcal{A}(N^K) \supseteq \mathcal{A}^*(N^K) = \{A^j | (A^j, N^K) \in \ell\}$ denote the state and decision nodes, respectively, that are *direct predecessors* of N^K . Figure 3-2, which is an extended version of the market share example in Figure 3-1, illustrates the notations for a hypothetical influence diagram.

An influence diagram which contains a value node is called an oriented influence diagram. We will call a state node S^K *controllable* if there is an arc from a decision node into S^K , i.e. $\mathcal{A}^*(S^K) \neq \emptyset$, and a *conditional state* if there is an arc from another state node into S^K , i.e. $\mathcal{S}^*(S^K) \neq \emptyset$. The decisions $\mathcal{A}^*(S^K)$ of a controllable state will be called *interventions* on S^K . A^K is a *conditional decision* if a state node is its direct predecessor, i.e. $\mathcal{S}^*(A^K) \neq \emptyset$. An influence diagram is said to be in the *canonical form* if it has no controllable states, i.e. $\mathcal{A}^*(S^j) = \emptyset, \forall S^j$. In Figure 3-2, S^C and S^M are controllable and conditional states and A^T and A^D are the corresponding interventions.

Figure 3-2: Notations for components of a hypothetical influence diagram (extended version of Figure 1).

The value node, V , and state nodes, S^j , have real value mappings associated with them which describe the inputs required for an influence diagram: for the value node we have $U: \prod_{\mathcal{A}} A^j \times \prod_{\mathcal{S}} S^j \rightarrow \mathbb{R}$, which is the utility of a particular policy and event, $U(\mathcal{A}_i, \mathcal{S}_j)$, to the decision maker and for any state node, S^K , we have $P: S^K \times \prod_{\mathcal{S}^*(S^K)} S^j \times \prod_{\mathcal{A}_i^*(S^K)} A^j \rightarrow [0,1]$, which is the conditional probability $P_{\mathcal{A}_i^*(S^K)}(s_i^K | \mathcal{S}_i^*(S^K))$ of the atomic states $s_i^K \in S^K$, given $\mathcal{S}_i^*(S^K)$ occurs and under the hypothesis that policy $\mathcal{A}_i^*(S^K)$ is selected. Given the real value mappings for V and \mathcal{S} the influence diagram can be evaluated using expected utility calculations and the optimal policy $\mathcal{A}_{opt} = (a_{i_1}^1, a_{i_2}^2, \dots, a_{i_{|A^i|}}^{|A^i|}) \in \prod_{\mathcal{A}} A^j$ can be determined. Since expected utility calculations will not play a role in what is discussed in this paper, the procedure is not described here. Details of the procedure (similar to the algorithm proposed in Shachter (1986)) can be found in Supplementary Note S 3-3.

In the rest of the paper, vectors are distinguished from scalars using bold characters. In particular, $\mathbf{P}(S^A, S^B, \dots, S^H)$ is used to represent the joint probability distribution for $\{S^A, S^B, \dots, S^H\}$, $\mathbf{P}(S^K, \mathcal{S}_i^*(S^K))$ to represent sections of the joint probability distribution for $\{S^K \cup \mathcal{S}^*(S^K)\}$ and $P(s_j^K, \mathcal{S}_i^*(S^K))$ to represent the individual probabilities.

3.3 Modelling interventions as probability revision

In this section, we will discuss how interventions on a controllable state node S^K can be modelled as revision of ‘prior’ probabilities of S^K . In the market share example (Figure 3-1) this corresponds to describing a way to obtain P_{aD} from $P_{\neg aD}$. We will show that for a controllable state S^K , when $\mathcal{S}^*(S^K) \neq \emptyset$, a probability revision rule for modelling the effect of actions on probability distributions can significantly reduce the effort required to elicit the inputs for an influence diagram.

When the probability of a state node S^K depends on the outcome of another state node, $\mathcal{S}^*(S^K) = S^R = \{s_i^R\}$, the rule that is suitable for updating the probability of s_i^K , given an observation s_l^R , is Bayes conditionalization and we have $\forall s_i^K, P(s_i^K | s_l^R) = P(s_i^K, s_l^R) / \sum_i P(s_i^K, s_l^R) = P(s_i^K, s_l^R) / P(s_l^R)$. In the case of interventions, when the probability of S^K depends on actions, the revised probabilities of S^K cannot be calculated using Bayes conditionalization, $P(s_i^K | a) = P(s_i^K \& a) / P(a)$, since normally a decision maker does not assign probabilities to actions. We therefore need an alternative probability revision rule for interventions which does not involve assigning probabilities to actions.

In this section we will develop a probability revision rule for interventions. We will first review the generic controller, which is an existing approach for describing post-intervention probabilities as a probability revision of pre-intervention probabilities. We will then generalize this procedure and show that our generalization of the generic controller coincides with a class of linear probability revision rules called ‘imaging’. The probability revision rules discussed in this paper will be described as functions that quantitatively describe how the probabilities of the controllable state are redistributed as a result of action or observation. Formally, we define a probability revision function as follows:

Definition 3-1: Probability revision function


Suppose $\mathbf{P}(S^K) = \langle P(s_1^K), P(s_2^K), \dots, P(s_{|S^K|}^K) \rangle$ is a probability distribution for a controllable state S^K , where $|S^K|$ is the number of outcomes of S^K . Let $\Delta^{|S^K|}$ denote the $|S^K|$ -dimensional probability simplex. A probability revision function given an action or observation c , denoted by f_c , is a function $f_c: \Delta^{|S^K|} \rightarrow \Delta^{|S^K|}$.

We will only consider the probability revision of marginal distributions in this section; probability revision of joint distributions is discussed in the next section, where we present some properties that probability revision functions should satisfy.

3.3.1 The generic controller

In Decision Analysis, one way to model stochastic interventions is by using the concept of the generic controller (Matheson and Matheson 2005). In this approach, interventions are *atomic*, i.e. raise the probability of one of the atomic states $s_i^K \in S^K$ to 1. Let $s_t^K \in S^K$ be the desired atomic state, i.e. the atomic state for which the decision maker wishes to change the probability to 1. Let $\mathbf{e}_t^K = \langle e(s_i^K, s_1^K), e(s_i^K, s_2^K) \dots, e(s_i^K, s_{|S^K|}^K) \rangle$ be the probability function where $e(s_i^K, s_j^K) = 1$ for $i = j$ and 0 otherwise. The probability vector \mathbf{e}_t^K represents an atomic intervention which assigns probability value 1 to the desired atomic state s_t^K . An intervention is called “perfect” if it increases the probability of s_t^K to 1 and “imperfect” (or stochastic) if it increases the probability of s_t^K to some value less than 1. An imperfect intervention is modelled by introducing a new random variable (Q in Figure 3-3b) to the existing influence diagram (Figure 3-3a), with an arc pointing to the controllable state node (S^K in Figure 3-3). This variable Q , also known as the quality node, has two outcomes, perfect and useless, with probabilities q and $1 - q$ respectively, but is *hidden* in the influence diagram since the outcomes are not actually observed. In the revised influence diagram, the probabilities of the controllable state node, $P(s_i^K)$, are described by probabilities conditional on the quality node: $\forall s_i^K \in S^K, P(s_i^K | useless) = P(s_i^K)$ and $P(s_i^K | perfect) = e(s_t^K, s_i^K)$.

Figure 3-3: An influence diagram (a) modified to represent the generic controller (b).



For the generic controller, the post-intervention distribution, $\mathbf{P}_a(S^K) = f_a(\mathbf{P}(S^K))$, of the controllable state node S^K is given by: $\forall s_i^K \in S^K, P_a(s_i^K) = P(s_i^K | useless) * P(useless) + P(s_i^K | perfect) * P(perfect)$ or equivalently:

$$\forall s_i^K \in S^K, \quad P_a(s_i^K) = \begin{cases} P(s_i^K) * (1 - q) + 0 * q, & i \neq t \\ P(s_i^K) * (1 - q) + 1 * q, & i = t \end{cases} \quad (3-2)$$

Box 1 shows a numerical application of the generic controller to the market share example described in Figure 3-1.

Box 1: Example of the generic controller for the market share example in Figure 3-1

Suppose the market share node has three outcomes which correspond to no market share (s_o^M), partial market share (s_p^M) and full market share (s_f^M) and the pre-intervention distribution (firm does not advertise) is given by:


s_o^M	s_p^M	s_f^M
0.2	0.6	0.2

If a^D is modelled as an atomic intervention which increases the probability of full market share, s_f^M , and the quality of the intervention is $q = 0.5$, then the post-intervention probabilities, $P_{a^D}(s_i^M)$, (using equation (3-2)) are given by:

s_o^M	s_p^M	s_f^M
0.1	0.3	0.6

Viewed through the algebraic lens, the generic controller moves the pre-intervention probability vector to a different point on the standard probability simplex (convex hull of all atomic interventions e_j^K). The post-intervention probability vector is located on the line that connects the pre-intervention probability vector to the vertex which represents the perfect atomic intervention. Figure 3-4 shows the probability simplex for the market share example discussed in Box 1.


Figure 3-4: Probability simplex for three atomic states $\{s_o^M, s_p^M, s_f^M\}$ where s_f^M is the most desired state and the intervention has quality $q = 0.5$.



A key advantage of the generic controller, which is not explicitly discussed by its proponents, is that this approach clearly separates beliefs about uncertainty (probability

distribution of S^K) from views about intervention (modelled as the distribution of a hidden chance variable Q). The generic controller can also be easily applied to influence diagrams when there are multiple stochastic interventions on the controllable state node or when the controllable state node has other direct predecessors which are state nodes. For instance, consider a modified example of the market share example shown in Figure 3-1, where instead of one intervention, we have two interventions on the market share node — advertising (A^D) and decision to takeover an existing competitor (A^T) (Figure 3-5).

Figure 3-5: Influence diagram for market share problem with multiple interventions.



Suppose we model the interventions in Figure 3-5 as generic controllers and also suppose that both advertising, a^D , and decision to take over, a^T , increases the probability of full market share, s_f^M , with qualities q^D and q^T , respectively. Then we can apply the generic controller twice so that the post-intervention probability $\mathbf{P}_{a^T, a^D}(S^M) = (f_{a^T} \circ f_{a^D})(\mathbf{P}(S^M))$ is given by:

$$\forall s_i^M \in S^M,$$

$$P_{a^T, a^D}(s_i^M) = \begin{cases} P(s_i^M) * (1 - q^D)(1 - q^T), & i \neq f \\ P(s_i^M) * (1 - q^D)(1 - q^T) + q^D + q^T - q^D * q^T, & i = f \end{cases} \quad (3-3)$$

Notice in equation (3-3), a convenient property of the generic controller is that the calculation of the post-intervention probability is invariant to the order in which the probability revision function for a^D and a^T are applied, i.e. $(f_{a^T} \circ f_{a^D})(\mathbf{P}(S^M)) = (f_{a^D} \circ f_{a^T})(\mathbf{P}(S^M))$. This order invariance property stems from the fact that the generic controller is formulated as a Bayesian probability revision (by conditioning on a fictitious state node (e.g. Q in Figure 3-3)), and Bayes rule has the special property that iterated application of Bayes rule is invariant to the order in which the revisions are applied (demonstrated in Supplementary Note S 3-4).

We will now offer a different perspective of the generalized controller. Rewriting equation (3-2) for the generic controller, we have:


$$\forall s_i^K \in S^K, \quad P_a(s_i^K) = \begin{cases} P(s_i^K) * (1 - q) + 0 * \sum_{j \neq i} P(s_j^K), & i \neq t \\ P(s_i^K) + q * \sum_{j \neq i} P(s_j^K), & i = t \end{cases} \quad (3-4)$$

It can be seen that the post-intervention probability is a linear mixture of the pre-intervention probabilities, i.e. $P_a(s_i^K) = \sum_j P(s_j^K) \rho^a(s_i^K, s_j^K)$, where the mixing coefficients, $\rho_a(s_i^K, s_j^K)$, for $P_a(s_i^K)$ are given by:

$$\forall s_i^K \in S^K, \quad \rho^a(s_i^K, s_j^K) = \begin{cases} 1 - q, & i = j; i \neq t, \\ 0, & i \neq j; i \neq t \\ 1, & i = j; i = t \\ q, & i \neq j; i = t \end{cases} \quad (3-5)$$

From a decision theoretic perspective, $\rho^a(s_i^K, s_j^K)$ can be interpreted as encoding the decision maker's views about how an intervention redistributes the probabilities of a controllable state node. Consider the expression for the pre-intervention probability $P(s_i^K)$. From the definition of e_j^K , we can write $P(s_i^K) = \sum_j P(s_j^K) e(s_i^K, s_j^K)$, $\forall s_i^K \in S^K$. The value $e(s_i^K, s_j^K)$ can be interpreted as the proportion of probability mass of $P(s_j^K)$ which is concentrated on $P(s_i^K)$. In the expression for post-intervention probability, $P_a(s_i^K) = \sum_j P(s_j^K) \rho^a(s_i^K, s_j^K)$, $\rho^a(s_i^K, s_j^K)$ can be interpreted as the proportion of probability mass of s_j^K that is transferred to s_i^K , ($i \neq j$) as a result of the intervention. In the generic controller, equal proportions, q , of probability masses are transferred from the less desirable states to the desirable states (illustrated in Figure 3-6 for the market share example in Box 1).

Figure 3-6: Proportion of probability mass transfer illustrated for the generic controller applied to the market share example, where $q = 0.5$.



The motivation for re-interpreting the generic controller is that it helps us to establish a link between the generic controller and a class of probability revision functions called ‘imaging’, that has been proposed in Philosophy as an alternative to Bayes rule. In the next section, we

will generalize the generic controller and then, in section 3.3.3, expose how the generalization is related to imaging.


3.3.2 Generalizing the generic controller

We will first propose an extension of the generic controller which allows interventions that can bring about more than one desirable state to be represented. We will then propose a further generalization which also enables interventions that prevent undesirable states from occurring to be represented. Suppose instead of having only two outcomes for the hidden chance (quality) node Q , perfect and useless with probabilities q and $1 - q$, the outcomes of Q can be various perfect interventions corresponding to the desired atomic states $\{s_i^K\} \in S^{K'} \subset S^K$, each with quality q_i ($\sum q_i \leq 1; q_i \geq 0$) and a useless intervention with quality $(1 - \sum q_i)$. Then $\{s_i^K \in S^{K'}\} \cup \{useless\}$ will be the outcomes of the hidden chance node with probabilities $\{q_i\}_{s_i^K \in S^{K'}} \cup \{1 - \sum_{s_i^K \in S^{K'}} q_i\}$. For this extension, called the *extended generic controller (EGC)*, the post-intervention distribution, $\mathbf{P}_a(S^K) = f_a(\mathbf{P}(S^K))$, of the controllable state node is given by:

$$\forall s_i^K \in S^K, \quad P_a(s_i^K) = \begin{cases} \left(1 - \sum_{s_j^K \in S^{K'}} q_j\right) * P(s_i^K), & s_i^K \notin S^{K'} \\ \left(1 - \sum_{s_j^K \in S^{K'}} q_j\right) * P(s_i^K) + q_i * 1, & s_i^K \in S^{K'} \end{cases} \quad (3-6)$$

Consider the pre-intervention probability distribution for the market share example in Box 1. Suppose the firm believes that advertising has a positive effect on both full and partial market share, so $S^{M'} = \{s_p^M, s_f^M\}$ and the corresponding perfect interventions have qualities $q_p = 0.3$ and $q_f = 0.5$. Then the post-intervention probabilities (using equation (3-6)) are given by $P_{aD}(s_o^M) = 0.04$, $P_{aD}(s_p^M) = 0.42$ and $P_{aD}(s_f^M) = 0.54$. Graphically, this extension of the generic controller moves the pre-intervention probability towards the simplex which connects the states in $S^{K'}$ (Figure 3-7). In the special case where $S^{K'}$ contains only one atomic state, this simplex is a point and corresponds to the generic controller.

Figure 3-7: Probability simplex for three atomic states $\{s_o^M, s_p^M, s_f^M\}$ where the probabilities of both s_f^M and s_p^M are raised ($S^{M'} = \{s_p^M, s_f^M\}$) and the intervention has qualities $q_p = 0.3, q_f = 0.5$.




In terms of probability mass transfers, the $\rho^a(s_i^K, s_j^K)$ corresponding to the EGC are given by:

$$\forall s_i^K \in S^K, \quad \rho^a(s_i^K, s_j^K) = \begin{cases} 1 - \sum_{s_j^K \in S^{K'}} q_j, & i = j; i, s_i^K \notin S^{K'}, \\ 0, & i \neq j; s_i^K \notin S^{K'} \\ 1 - \sum_{s_j^K \in \{S^{K'} - s_i^K\}} q_j, & i = j; s_i^K \in S^{K'} \\ q_i, & i \neq j; s_i^K \in S^{K'} \end{cases} \quad (3-7)$$

For the market share example, where $S^{M'} = \{s_p^M, s_f^M\}$ and $q_p = 0.3, q_f = 0.5$, the coefficients $\rho^a(s_i^K, \cdot) = \langle \rho^a(s_i^K, s_o^M), \rho^a(s_i^K, s_p^M), \rho^a(s_i^K, s_f^M) \rangle$ are given by: $\rho^a(s_o^K, \cdot) = \langle 0.2, 0, 0 \rangle$, $\rho^a(s_p^K, \cdot) = \langle 0.3, 0.5, 0.3 \rangle$, $\rho^a(s_f^K, \cdot) = \langle 0.5, 0.5, 0.7 \rangle$. Figure 3-8 illustrates the corresponding probability mass transfers.

Figure 3-8: Proportion of probability mass transfer illustrated for the EGC applied to the market share example, where $S^{M'} = \{s_p^M, s_f^M\}$ and $q_p = 0.3, q_f = 0.5$.



We note that since the EGC is also formulated as a Bayesian probability revision (conditioning on the fictitious node Q), when multiple interventions are modelled as EGCs, the EGCs share the order invariance property of the generic controller. One of the limitations of the generic controller and EGC is that they can only be used to model interventions that raise the probability of atomic states and it is not immediately clear how they can be used to model interventions that are designed to prevent an undesirable state from occurring. However, if we do not restrict $\rho^a(s_i^K, s_j^K)$ to the form in equation (3-7) and allow $\rho^a(s_i^K, s_j^K)$ to take any value in $[0,1]$, with the requirement that $\sum_j \rho^a(s_i^K, s_j^K) = 1, \forall i$ (to conserve the sum of pre-intervention probabilities), then we have a further generalization of the generic controller which can be used to describe interventions that reduce the probability of an undesired atomic state. Formally, the probability revision function f_a and post-intervention distribution, $\mathbf{P}_a(S^K) = f_a(\mathbf{P}(S^K))$, for the generalized generic controller (GGC) is given by:

$$\begin{aligned} \forall s_i^K \in S^K, \quad P_a(s_i^K) &= \sum_j P(s_j^K) \rho^a(s_i^K, s_j^K); \rho^a(s_i^K, s_j^K) \\ &\in [0,1], \sum_i \rho^a(s_i^K, s_j^K) = 1 \end{aligned} \tag{3-8}$$

One approach for eliciting $\rho^a(s_i^K, s_j^K)$, in a manner similar to the ‘probability wheel’ (Shephard and Kirkwood, 1994), is described in Appendix 3-1. (In Supplementary Note S 3-5, the elicitation is illustrated for the market share example, through a conversation between a manager and a decision analyst).

We offer an algebraic interpretation for the GGC which is different from those in Figure 3-4 and Figure 3-7. Graphically, $\{\mathbf{e}_j^K\}$ can be interpreted as the vertices of a standard simplex and $\mathbf{P}(S^K)$ is a point on this simplex. The probability revision described by the generic controller defines a new simplex with respect to the vectors $\{\rho^a(s_1^K, \cdot), \rho^a(s_2^K, \cdot), \dots, \rho^a(s_n^K, \cdot)\}$ and $\mathbf{P}_a(S^K)$ is the projection of $\mathbf{P}(S^K)$ on the new simplex. Unlike the generic controller and EGC, the GGC is not order invariant and therefore it may not be possible to formulate the GGC as a Bayesian probability revision process. However, as we will show in section 3.4, the GGC shares some desirable properties of Bayes rule which are relevant when the controllable state node, S^K , has other direct predecessors which are state nodes, i.e. $\mathcal{S}^*(S^K) \neq \emptyset$. In the next subsection we will discuss the correspondence between the GGC and probability revision rules that are developed in Philosophy.

3.3.3 Generalized generic controller and imaging

Based on the semantics for belief revision formulated by Stalnaker (1968), Lewis (1976) proposed a probability revision rule called *imaging* (generalized by Gardenfors (1988, pp. 108–18, 1982)) to describe probabilities of conditional statements as revisions of probabilities of unconditional statements. In contrast to Bayes rule, which *rescales* the probabilities once information about some states becomes available, imaging describes probability revision in terms of *probability mass transfer* between states, based on the decision maker's qualitative views about how 'similar' the states are (Joyce, 1999, p. 198). For instance, suppose $\mathbf{P}(S^K)$ depends on the outcome of S^R . Given $s_l^R \in S^R$ is observed, while the revised probability for $s_i^K \in S^K$ based on Bayes rule is given by $P(s_i^K | s_l^R) = P(s_i^K, s_l^R) / P(s_l^R)$, the revised probability based on imaging, $P_{s_l^R}(s_i^K)$, does not depend on the prior probability $P(s_i^K, s_l^R)$ but, instead, is described by probability transfers from the joint events $(s_i^K, s_j^R) \in (S^K \times \{S^R - s_l^R\})$ to $(s_i^K, s_l^R) \in (S^K \times s_l^R)$, where the proportion of probability transfer, $\rho^{s_l^R}(s_i^K, s_j^R)$, from (s_i^K, s_j^R) to (s_i^K, s_l^R) depends on some notion of similarity between (s_i^K, s_j^R) and (s_i^K, s_l^R) . Gardenfors (1988, pp. 108–18, 1982) proved that the characterizing property for generalized imaging, which distinguishes it from Bayes rule, is a linearity condition defined as follows:

Definition 3-2: Linear probability revision function

Let $\mathbf{P}_\alpha(S^K) = f_\alpha(\mathbf{P}(S^K))$ be the post-intervention probabilities for a controllable state node $S^K = \{s_i^K\}$, for any probability function \mathbf{P} . A probability revision function f_α is linear if for some probability functions \mathbf{P}' , \mathbf{P}'' , $\alpha \in [0,1]$, $\mathbf{P}(S^K) = \alpha\mathbf{P}'(S^K) + (1 - \alpha)\mathbf{P}''(S^K)$, then $\mathbf{P}_\alpha(S^K) = \alpha\mathbf{P}'_\alpha(S^K) + (1 - \alpha)\mathbf{P}''_\alpha(S^K)$.

This definition is illustrated in Box 2.

It can be shown that the GGC in equation (3-9) is characterized by the *same* linearity property and thus coincides with imaging.

Theorem 3-1: Representation theorem for the generalized generic controller

A probability revision function f_α is linear if and only if $\mathbf{P}_\alpha(S^K) = f_\alpha(\mathbf{P}(S^K))$ is a generalized generic controller.

The proof for Theorem 3-1 can be found in Appendix 3-2. (In Supplementary Note S 3-6 the relationship between Bayes rule and the linearity property is discussed).

Box 2: Example to illustrate the definition of a linear probability revision function

Suppose there are three urns A, B, C which contain 100 balls, either coloured yellow or coloured black. The number of yellow coloured balls in each urn is as follows:

Urn A 20	Urn B 60	Urn C 40
-------------	-------------	-------------

Consider the following two alternatives:

(1)	(2)
Flip a fair coin and draw a ball from urn A if the coin shows heads else draw a ball from urn B	Draw a ball from urn C


Let $\{s_{ye}, s_{bl}\}$ denote the uncertain outcomes of the draw in both alternatives. The probability of drawing a yellow ball ($P(s_{ye}) = 0.4$) is the same in both alternatives. Now consider two different rules to *revise* the contents of each urn:

- 1) replace 50% of black balls with yellow balls;
- 2) replace $x\%$ of black balls with yellow balls where x is the number of yellow balls in the urn.

If the first rule is applied, the revised probability of drawing a yellow ball is the same (0.7) in both alternatives and therefore the corresponding probability revision function $f'(P(s_{ye})) = P(s_{ye}) + 0.5P(s_{bl})$, $f'(P(s_{bl})) = 0.5P(s_{bl})$ is linear. However if the second rule is applied, the revised probability of drawing a yellow ball is different in each alternative (0.60 in the first alternative and 0.64 in the second alternative) and therefore the corresponding probability revision function $f''(P(s_{ye})) = P(s_{ye})(1 + P(s_{bl}))$, $f''(P(s_{bl})) = P(s_{bl})(1 - P(s_{ye}))$ is not linear.

We will now demonstrate the advantage of modelling stochastic interventions as probability revision rules in influence diagrams. Suppose in Figure 3-1 the firm realizes that the probability of the market share is in fact conditional on whether or not a competitor advertises its product. To incorporate this information, the influence diagram is modified and a direct predecessor $S^C = \{s_d^C, s_{-d}^C\}$, with two scenarios (competitor advertises, competitor does not advertise) is added to the controllable state node S^M (Figure 3-9). The distribution of the market share node, S^M , is described by the conditional probability distributions $P(S^M | s_j^C)$.

Figure 3-9: Influence diagram in Figure 2 modified to include another direct predecessor for the market share node S^M .



Suppose the intervention on S^M is modelled as a probability revision and $\mathbf{P}_a(S^M) = f_a(\mathbf{P}(S^M))$. Then the post-intervention conditional probabilities $P_a(s_i^M | s_j^C)$ can be calculated by applying f_a to $\mathbf{P}(S^M | s_j^C)$ without requiring new elicitations. On the other hand, in the absence of a probability revision procedure, the probability distributions for $P_a(s_i^M | s_j^C)$ would need to be re-elicited. In general, with a probability revision function like the GGC, the judgmental burden of eliciting probability distributions is reduced considerably when the controllable node has several direct predecessors that are state nodes. In particular, for n atomic states of the controllable node, k options for the intervention and $m \geq 1$ atomic states of direct predecessors of the controllable state node, compared to direct elicitation of post-intervention conditional probabilities, which requires $(n - 1) * m * k$ elicitations, the GGC requires at most $(n - 1) * ((k - 1) * n + m)$ elicitations. For large influence diagrams, when the total number of atomic states of the direct predecessors of the controllable state node is greater than the number of atomic states of the controllable state node, the number of elicitations for the GGC will be much fewer than direct elicitation of post-intervention probabilities (illustrated in Table 3-1).

It is important to note that when the controllable state has direct predecessors that are state nodes, one cannot in general make inferences about these predecessor state nodes from the post-intervention distribution of the controllable state, i.e. the inference about $P_a(s_j^C | s_i^M)$ is not valid and need not be equal to $P(s_j^C | s_i^M)$ (illustrated in Supplementary Note S 3-7 for the market share example). In other words, all arcs into and out of the controllable state node are non-backtracking and cannot be reversed. This restriction is also assumed by Matheson and Matheson (2005).

Table 3-1: Numerical comparison of the number of elicitation required for each approach.

Atomic 'controllable' states n	Atomic decision options k	Total direct predecessor atomic states m	Direct elicitation $(n - 1) * k * m$	Generalized generic controller $(n - 1) * ((k - 1) * n + m)$
3	2	1	4	8
3	2	2	8	10
3	2	3	12	12
3	2	4	16	14
3	2	5	20	16
3	2	10	40	26
3	3	1	6	14
3	3	2	12	16
3	3	3	18	18
3	3	4	24	20
3	3	5	30	22
3	3	10	60	32

The generalized generic controller or, equivalently, linear probability revision functions are one of many ways of describing a probability revision function for stochastic interventions. In the next section we offer some arguments for why the GGC is a suitable probability revision procedure for modelling stochastic interventions.

3.4 Some foundations for probability revision rules for interventions

In this section we will state two properties that probability revision functions for interventions should satisfy and expose why these properties are reasonable requirements. We will show that the GGC satisfies both these requirements. Both the properties that we state are related to interventions that are performed on a controllable state S^K that also has a state node S^R as a direct predecessor (e.g. Figure 3-9), and thus correspond to probability revision of joint distributions $\mathbf{P}(S^K, S^R)$. In order to discuss these properties we therefore, first extend Definition 3-1 to describe a probability revision function for joint probability distributions (e.g. $f_a(\mathbf{P}(S^K, S^R))$).

Definition 3-3: Probability revision function for joint distributions

Suppose $S^K = \{s_i^K\}$ is a controllable state node, $|S^K|$ is the number of outcomes of S^K , $\mathcal{S}^*(S^K) = \{S^A, S^B, \dots, S^H\}$ are the direct predecessors of S^K , $\mathcal{S}_i^*(S^K) = (s_{A_i}^A, s_{B_i}^B, \dots, s_{H_i}^H) \in \prod_{\mathcal{S}^*(S^K)} S^j$ are the events corresponding to $\mathcal{S}^*(S^K)$, $|\mathcal{S}^*(S^K)|$ is the total number of events of $\mathcal{S}^*(S^K)$ and $P(s_{K_i}^K, s_{A_i}^A, \dots, s_{H_i}^H)$ are the joint probabilities of the events. For any $\mathcal{S}_i^*(S^K)$, denote the corresponding section of the joint probability distribution for $(S^K \cup \mathcal{S}_i^*(S^K))$ by the vector $\mathbf{P}(S^K, \mathcal{S}_i^*(S^K)) = \langle P(s_1^K, \mathcal{S}_i^*(S^K)), \dots, P(s_{|S^K|}^K, \mathcal{S}_i^*(S^K)) \rangle$. A probability revision function for S^K ,

given an action or observation c , is a vector-valued function $f_c = \langle f_c^1, f_c^2, \dots, f_c^{|S^*(S^K)|} \rangle: \Delta^{|S^K||S^*(S^K)|} \rightarrow \Delta^{|S^K||S^*(S^K)|}$, where the component functions $f_c^i: \mathbf{P}(S^K, S_i^*(S^K)) \rightarrow [0,1]^{|S^K|}$, for any $\mathbf{P}(\cdot)$.

Box 3 illustrates a hypothetical probability revision function for the market share example in Figure 3-9.

Box 3: Numerical Example of a probability revision function

In Figure 3-9, suppose, like in Box 1, there are three atomic states for the market share node ($S^M = \{s_o^M, s_p^M, s_f^M\}$) and that $P(s_f^M)$ is higher if the competitor does not advertise (s_{-d}^C). Consider the following joint distribution of (S^M, S^C) :

	s_{-d}^C	s_d^C
s_o^M	1/20	3/20
s_p^M	6/20	6/20
s_f^M	3/20	1/20

Let p_{ij} denote the probability for row i and column j of this table. An example of a probability revision function (that represents the effect of the firm advertising a^D) is $f_a = \langle f_a^{-d}, f_a^d \rangle: \Delta^6 \rightarrow \Delta^6$ where $f_a^i: \Delta^3 \rightarrow \Delta^3$ is given by:

$$f_a^{-d}(p_{11}, p_{21}, p_{31}) = \left(\frac{2(p_{11}+p_{21}+p_{31})}{9}, \frac{2(p_{11}+p_{21}+p_{31})}{9}, \frac{2(p_{11}+p_{21}+p_{31})}{9} \right)$$

$$f_a^d(p_{12}, p_{22}, p_{32}) = \left(\frac{1}{3}, \frac{(p_{12}+p_{22}+p_{32})}{3}, \frac{(p_{12}+p_{22}+p_{32})}{3} \right)$$

The post-intervention distribution is given by:

$$f_a^{-d} \left(\frac{1}{20}, \frac{6}{20}, \frac{3}{20} \right) = \left(\frac{1}{9}, \frac{1}{9}, \frac{1}{9} \right) \text{ and } f_a^d \left(\frac{3}{20}, \frac{6}{20}, \frac{1}{20} \right) = \left(\frac{2}{6}, \frac{1}{6}, \frac{1}{6} \right).$$

For a controllable state S^K , where $S^*(S^K) = S^R$, the probability revision function for the GGC and Bayes rule is defined as follows:

Generalized generic controller (for an intervention a on S^K):

Let $f_a = \langle f_a^1, f_a^2, \dots, f_a^{|S^R|} \rangle: \Delta^{|S^K||S^R|} \rightarrow \Delta^{|S^K||S^R|}$ be a probability revision function

for the GGC and $P_a(s_i^K, s_m^R)$ represent the corresponding revised probabilities.

Then,

$$\forall m, \quad P_a(s_i^K, s_m^R) = \sum_j P(s_j^K, s_m^R) \rho^a(s_i^K, s_j^K);$$

$$0 \leq \rho^a(s_i^K, s_j^K) \leq 1, \quad \sum_i \rho^a(s_i^K, s_j^K) = 1 \quad (3-9)$$

Bayes rule (for conditioning the distribution $P(S^K, S^R)$ on an observed state s_l^R):

Let $g_{s_l^R} = \langle g_{s_l^R}^1, g_{s_l^R}^2, \dots, g_{s_l^R}^{|S^R|} \rangle: \Delta^{|S^K||S^R|} \rightarrow \Delta^{|S^K||S^R|}$ be a probability revision function for Bayes rule and $P_{s_l^R}(s_i^K, s_m^R)$ represent the corresponding revised probabilities. Then,

$$\forall m, P_{s_l^R}(s_i^K, s_m^R) = \begin{cases} 0 & m \neq l \\ \frac{P(s_i^K, s_m^R)}{\sum_j P(s_j^K, s_m^R)} & m = l \end{cases} \quad (3-10)$$

For example, in the numerical example in Box 3, Bayes conditioning on s_d^C is described by $g_{s_d^C} = \langle g_{s_d^C}^{-d}, g_{s_d^C}^d \rangle: \Delta^6 \rightarrow \Delta^6$, $g_{s_d^C}^{-d}(\frac{1}{20}, \frac{6}{20}, \frac{3}{20}) = (0, 0, 0)$, $g_{s_d^C}^d(\frac{3}{20}, \frac{6}{20}, \frac{1}{20}) = (\frac{3}{10}, \frac{6}{10}, \frac{1}{10})$ and conditioning on s_{-d}^C by $g_{s_{-d}^C} = \langle g_{s_{-d}^C}^{-d}, g_{s_{-d}^C}^d \rangle: \Delta^6 \rightarrow \Delta^6$, $g_{s_{-d}^C}^{-d}(\frac{1}{20}, \frac{6}{20}, \frac{3}{20}) = (\frac{1}{10}, \frac{6}{10}, \frac{3}{10})$, $g_{s_{-d}^C}^d(\frac{3}{20}, \frac{6}{20}, \frac{1}{20}) = (0, 0, 0)$.

We will now state the properties that a probability revision function for joint distributions must fulfill when the controllable state node, S^K , has a direct predecessor, i.e. $\mathcal{S}^*(S^K) \neq \emptyset$, and expose the importance of these properties.

3.4.1 Basic properties for probability revision rules for interventions

Suppose S^K is a controllable state node, with a direct predecessor node $\mathcal{S}^*(S^K) = S^R$, $P(S^K, S^R)$ is any joint probability distribution for $\{S^K, S^R\} = \{(s_i^K, s_j^R) | i = 1, 2, \dots, |S^K|; j = 1, 2, \dots, |S^R|\}$ and $f_a: \Delta^{|S^K||S^R|} \rightarrow \Delta^{|S^K||S^R|}$ is a probability revision function for an intervention on S^K . Then, two basic properties that f_a should satisfy are:

Fixed-point at zero: $\forall j, f_a^j(\mathbf{0}) = \mathbf{0}$ where $\mathbf{0}$ represents the $|S^K|$ -dimensional zero vector

Bayes conditionalization preserving: $\forall j, f_a^j(P(S^K, s_j^R)) = f_a^j(P(S^K | s_j^R)) * P(s_j^R)$

The first property ensures that a probability revision function for an intervention does not generate counterintuitive probability distributions for S^K if it is applied after information is received about S^R . To see this, consider the probability revision function f_a defined in Box 3, which does not satisfy property (i) ($f_a^d(0, 0, 0) = (\frac{1}{3}, 0, 0)$). Let $g_{s_{-d}^C}(P(S^K, S^R))$ represents Bayes conditioning with respect to the observation that the competitor does not advertise (s_{-d}^C). Table 3-2 illustrates the various transformations of the probability distribution in Box 3 when f_a and $g_{s_{-d}^C}$ are applied.

Table 3-2: Transformations of the probability distribution in Box 3, where f_a is the probability revision function in Box 3 and $g_{s_{-d}^C}$ is Bayes conditioning once s_{-d}^C (competitor does not advertise) is observed.

	$f_a(P(\cdot))$		$(g_{s_{-d}^C} \circ f_a)(P(\cdot))$		$g_{s_{-d}^C}(P(\cdot))$		$(f_a \circ g_{s_{-d}^C})(P(\cdot))$	
	s_{-d}^C	s_d^C	s_{-d}^C	s_d^C	s_{-d}^C	s_d^C	s_{-d}^C	s_d^C
s_o^M	1/9	2/6	1/3	0	1/10	0	2/9	1/3
s_p^M	1/9	1/6	1/3	0	6/10	0	2/9	0
s_f^M	1/9	1/6	1/3	0	3/10	0	2/9	0

Notice in Table 3-2 that when the probability revision function f_a is applied *after* information is received that s_{-d}^C will occur (i.e. $(f_a \circ g_{s_{-d}^C})$), f_a assigns positive probability to s_d^C , a state that is known to *not* occur. Thus, to avoid generating such conflicting beliefs, probability revision functions for interventions must fulfil property (i). (A necessary and sufficient condition for property (i) is stated in Theorem 3-3 in Appendix 3-4).

The second property states that if the distribution of S^K is revised to incorporate the effect of an intervention, f_a , and information about $S^*(S^K) = S^R$ (e.g. s_j^R is observed), then the corresponding component function, f_a^j , when applied to its section of the joint distribution, $\mathbf{P}(S^K, s_j^R)$, and the distribution conditional on the observation s_j^R , $\mathbf{P}(S^K | s_j^R)$, should satisfy Bayes rule. This property implies a *coherence condition* for probability revision functions for interventions: if a probability revision function for interventions does not satisfy property (ii), then it is possible to construct a Dutch book against the decision maker.

To see this, consider the probability revision function f_a defined in Box 3, which does not satisfy property (ii) (for f_a^d , $f_a^d(\mathbf{P}(S^M, s_d^C)) = (\frac{2}{6}, \frac{1}{6}, \frac{1}{6})$, $f_a^d(\mathbf{P}(S^M | s_d^C)) * P(s_d^C) = f_a^d(\frac{3}{10}, \frac{6}{10}, \frac{1}{10}) * 0.5 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}) * 0.5 = (\frac{1}{6}, \frac{1}{6}, \frac{1}{6}) \neq f_a^d(\mathbf{P}(S^M, s_d^C))$). The following Dutch book can be constructed for this probability revision function:

Let $\mathbf{P}_a(S^M | s_d^C) = f_a^d(\mathbf{P}(S^M | s_d^C))$ represent the post-intervention probabilities if the competitor advertises and $\mathbf{P}_a(S^M, s_d^C) = f_a(\mathbf{P}(S^M, s_d^C))$ represents the post-intervention probabilities if the competitor's decision is unknown. From Box 3, the probability that the competitor will not advertise, $P(s_{-d}^C)$, is 0.5 and $\mathbf{P}_a(S^M, s_d^C) = (\frac{2}{6}, \frac{1}{6}, \frac{1}{6})$. Suppose the firm commits to advertising. Given the firm's post intervention beliefs, $(\mathbf{P}_a(S^M | s_d^C))$ and $\mathbf{P}_a(S^M, s_d^C)$, a bookie designs two bets, which are shown in Table 3-3.

Table 3-3: Bets designed by a bookie which show the payoff to the owner of the bet.

Bet A on whether or not Competitor advertises		Bet B on the joint events (Competitor's decision and market share outcome)	
s_{-a}^C	s_a^C		
\$ 33	\$ -	s_o^M	\$ - \$ 100
		s_p^M	\$ - \$ -
		s_f^M	\$ - \$ -

Assuming risk neutrality, the fair price for each bet, based on the firm's post-intervention beliefs is:

$$\text{Bet A: } \$ 33 * P(s_{-a}^C) = \$16.5$$

$$\text{Bet B: } \$ 100 * P_a(s_a^C, s_o^M) = \$33$$

The bookie sells both the bets to the firm at the maximum price the firm is willing to pay (the firm pays the fair price $\$33 + \$16.5 = \$49.5$) and the bookie will make payments to the firm depending on the outcome of the events. Now consider the potential payoffs to the firm. If the competitor does not advertise, the firm wins Bet A and the bookie pays the firm \$33 (payoff to the firm is $-\$49.5 + \$33 = -\$16.5$). If the competitor advertises, the bookie buys Bet B back from the firm (thus declaring it void) at the minimum price the firm is willing to accept, which is the fair price $\$ 100 * P_{a,s_a^C}(s_o^M, s_a^C) = \$ 100 * 0.33 = \33 (payoff to the firm is $-\$49.5 + \$33 = -\$16.5$). Thus, the firm incurs a loss of \$16.5 regardless of which state occurs.

A general design of the Dutch book for probability revision functions which do not satisfy property (ii) is presented in Appendix 3-3.

An alternative interpretation of the basic properties (i) and (ii) is that probability revision functions should commute with Bayes rule: when $\mathcal{S}^*(S^K) = \{S^R, A\}$, the revised distribution of S^K should be the same regardless of the order in which the probability revisions for the intervention (e.g. f_a) and information (e.g. $g_{s_a^R}$) are applied. (The equivalence between this re-interpretation and the basic properties is proved in Appendix 3-4). We will now prove that the GGC satisfies both the basic properties of probability revision functions for interventions.

3.4.2 GGC satisfies the properties of probability revision rules for interventions

In this subsection, first we prove that the basic properties of probability revision functions, stated in section 3.4.1, are satisfied by functions which preserve scalar multiplication. Using this result, we will prove that the GGC satisfies the basic properties.

Theorem 3-2: Functions which satisfy the basic properties of probability revision functions

Suppose S^K is a controllable state node with a direct predecessor $S^*(S^K) = S^R$, $\mathbf{P}(S^K, S^R)$ is any joint probability distribution for $\{S^K, S^R\} = \{(s_i^K, s_j^R) | i = 1, 2, \dots, |S^K|; j = 1, 2, \dots, |S^R|\}$, $f_a: \Delta^{|S^K||S^R|} \rightarrow \Delta^{|S^K||S^R|}$ is a probability revision function for an intervention on S^K and $g_{s_i^R}: \Delta^{|S^K||S^R|} \rightarrow \Delta^{|S^K||S^R|}$ represents Bayes conditioning with respect to $s_i^R \in S^R$. If, for any scalar $\alpha \in [0, 1]$, $f_a^j(\alpha \mathbf{P}(S^K, s_j^R)) = \alpha f_a^j(\mathbf{P}(S^K, s_j^R))$, $\forall j$ then $\forall j$, (i) $f_a^j(\mathbf{0}) = \mathbf{0}$ and (ii) $f_a^j(\mathbf{P}(S^K | s_j^R)) * P(s_j^R) = f_a^j(\mathbf{P}(S^K, s_j^R))$.

Proof:

Since $\forall j$, $f_a^j(\alpha \mathbf{P}(S^K, s_j^R)) = \alpha f_a^j(\mathbf{P}(S^K, s_j^R))$, by setting $\alpha = 0$, have $\forall j$, $f_a^j(\mathbf{0}) = \mathbf{0}$.

$$\forall j, f_a^j(\mathbf{P}(S^K | s_j^R)) = f_a^j\left(\frac{\mathbf{P}(S^K, s_j^R)}{\sum_i P(s_i^K, s_j^R)}\right) = \frac{f_a^j(\mathbf{P}(S^K, s_j^R))}{\sum_i P(s_i^K, s_j^R)} = \frac{f_a^j(\mathbf{P}(S^K, s_j^R))}{P(s_j^R)}$$

Therefore, $\forall j$, $f_a^j(\mathbf{P}(S^K | s_j^R)) * P(s_j^R) = f_a^j(\mathbf{P}(S^K, s_j^R))$.

Note that property (i) and (ii) can hold for functions that do not preserve scalar multiplication: for example, Bayes rule does not preserve scalar multiplication ($\forall s_i^R \in S^R$, $g_{s_i^R}(\alpha \mathbf{P}(s_i^K, s_j^R)) = g_{s_i^R}(\mathbf{P}(s_i^K, s_j^R))$), but it can be verified that Bayes rule satisfies property (i) and (ii) (since Bayes rule commute; see Supplementary Note S 3-4).

Corollary: The generalized generic controller satisfies the basic properties of probability revision functions for interventions.

Proof:

It is easy to verify that the GGC preserves scalar multiplication. Suppose S^K is a controllable state node, $S^*(S^K) = S^R$, $\mathbf{P}(S^K, S^R)$ is the joint probability distribution for $\{S^K, S^R\} = \{(s_i^K, s_j^R) | i = 1, 2, \dots, |S^K|; j = 1, 2, \dots, |S^R|\}$ and $f_a = \langle f_a^1, f_a^2, \dots, f_a^{|S^R|} \rangle: \Delta^{|S^K||S^R|} \rightarrow \Delta^{|S^K||S^R|}$ is a GGC. For any $\alpha \in [0, 1]$, we have,

$$\begin{aligned} \forall l, f_a^l(\alpha \mathbf{P}(S^K, s_l^R)) &= \left(\sum_j \alpha P(s_j^K, s_l^R) \rho^a(s_1^K, s_j^K), \dots, \sum_j \alpha P(s_j^K, s_l^R) \rho^a(s_{|S^K|}^K, s_j^K) \right) \\ &= \alpha f_a^l(\mathbf{P}(S^K, s_l^R)). \end{aligned}$$

Applying Theorem 3-2, we have that f_a satisfies the basic properties of probability revision functions.

3.5 Summary and directions for further research

Managers often deal with risks by considering uncertainty as an element of the decision problem over which they can exert control (Chelst and Bodily, 2000; Courtney, 2001; March and Shapira, 1987; Rosenzweig, 2013). Managerial actions to influence the probability of uncertain states can be modelled as *stochastic interventions* in influence diagrams (Howard and Matheson 1984; Howard and Matheson 2005; Matheson and Matheson 2005) with an arc from the decision node to a state node. However, within the Decision Analysis framework, evaluating a number of different stochastic interventions using standard procedures, such as eliciting post-intervention probabilities directly, can be judgmentally very burdensome. In this paper, we addressed this challenge associated with modelling stochastic interventions in Decision Analysis frameworks and developed a general probability revision method for interventions, which is analogous to Bayes rule (the probability revision rule for information), but does not need probabilities to be assigned to actions.

The method we proposed in this paper — the generalized generic controller (GGC) — builds upon the generic controller of Matheson & Matheson (2005). We first extended the generic controller (extended generic controller (EGC)) and offered an interpretation of the generic controller and the EGC in terms of ‘probability mass transfers’ between states. We also exposed their ‘Bayesian’ nature and that they share the order invariance property of Bayes rule. This property can be advantageous when modelling simultaneous interventions on the same state node. We then developed the GGC by further generalizing the generic controller to allow any arbitrary movements of probability mass between states, so that more general interventions, such as those that prevent an undesirable state from occurring, can be represented. We proved that the GGC is equivalent to ‘imaging’ (Gärdenfors, 1982; Lewis, 1976), a probability revision method proposed in Philosophy, thus establishing an interesting link between methods in Decision Analysis and other disciplines. It was noted that the GGC does not share the order invariance property of the generic controller and EGC.

The tractability of our proposed procedure for analysing stochastic interventions has several prescriptive advantages in terms of using Decision Analysis methods to analyse managerial influences on uncertainty. First, we demonstrated that compared to standard procedures for analysing interventions, the GGC is less information intensive and requires fewer elicitations when the controllable state node has several direct predecessor state nodes. This can reduce the effort needed when managers want to explore several potential interventions. Another important advantage of this procedure, which has not been recognized before, is that it separates beliefs about uncertainty from beliefs about the effect of

intervention on the probability distribution. This is particularly helpful if an analyst wants to introduce new predecessors to the controllable state node in the decision problem. With a probability revision function, as described in the paper, the post-intervention probabilities can systematically be obtained from the new conditional probabilities of the state node and the decision maker does not need to reassess all the probability distributions again.

A potential theoretical contribution of this paper is formally grounding the concept of probability revision rules for interventions, formulated as functions, in foundational principles which are akin to the coherence criterion and Dutch book argument that supports Bayes rule as the probability revision rule for information. We stated two basic properties for probability revision functions for interventions that are key when the controllable state depends on the outcome of other state nodes and Bayes rule is used to incorporate any information about these state nodes. The first property we stated, fixed-point at zero, ensures that when probability revision functions for interventions are applied together with Bayes rule, the post-intervention distribution does not assign positive probabilities to states that are known to not occur (i.e. have probability zero). The second property, Bayes conditionalization preserving, which requires that post-intervention beliefs should satisfy Bayes rule, extends the coherence arguments (such as immunity to Dutch books) that support Bayes rule (Teller, 1973), to probability revision functions for interventions. We proved that the GGC (or imaging) and, more generally, functions which preserve scalar multiplication satisfy these basic properties, thus providing additional support to previous suggestions to use imaging for modelling interventions (Joyce, 1999; Pearl, 1994).

One limitation of the approach proposed in this paper is that, in general, when modelling multiple interventions as GGCs, the post-intervention distribution is not invariant to the order in which the GGCs are applied. This limits the possible ways in which interventions can be formulated when multiple interventions are performed simultaneously on the same state node. We leave the formalization of a general condition under which the GGCs commute to future research. From a practitioner's perspective, the feasibility of using the GGC to model real world interventions needs to be tested. One area where the GGC can be potentially useful is in producing Operational Risk assessments, where many uncertainties pertaining to people, processes and systems are controllable and managers often have mitigations in place to modify the probability of loss events. The probability revision method, proposed here, provides a quantitative way to model the impact of mitigations on Operational Risk costs, explicitly. Agarwal and Montibeller (2014) explore the application of the GGC to measure Operational Risks in a real world case study.

We identify two theoretical directions for further research which could be useful. One of the motivations for generalizing the generic controller, beyond the EGC, was to establish

its correspondence with probability revision procedures in other disciplines. A direction for potentially useful research is to provide ecological validity for the GGC as a probability revision rule for interventions, by exposing the counterfactual reasoning that underpins imaging and unifying it with how managers contemplate interventions as mechanisms to bring about desired states. A second direction for further research relates to the lack of a normative foundation for influence diagrams in the non-canonical form. Potential solutions to this problem are based on eliminating any direct links between the decision node and state node by reformulating the state that can be influenced, without affecting the overall analysis of the decision problem (see Matheson 1990; Gibbard and Harper 1978, 22; Jeffrey 1965). We propose that another interesting approach would be to investigate if alternative versions of expected utility theory, with a revised set of axioms, such as Conditional Expected Utility (Bolker, 1967; Jeffrey, 1965; Luce and Krantz, 1971) or Causal Decision Theory (Armendt, 1986; Joyce, 1999; Lewis, 1981; Skyrms, 1982), can serve as normative foundations for influence diagrams that are not in the canonical form.

Acknowledgements

This research is sponsored by the AXA Research Fund. We would like to thank James Matheson for insightful discussions and feedback on the contents of this paper. Opinions and technical proofs are responsibilities of authors only.

3.6 Appendices for Chapter 3

Appendix 3-1: Eliciting the coefficients for the generalized generic controller using the influence wheel

Let the controllable state node be $S^K = \{s_1^K, s_2^K, \dots, s_{|S^K|}^K\}$. The coefficients, $\rho^\alpha(s_i^K, s_j^K)$, can be elicited as follows:

For each $[\rho^\alpha(s_1^K, s_j^K), \rho^\alpha(s_2^K, s_j^K), \dots, \rho^\alpha(s_{|S^K|}^K, s_j^K)]$:

Step 1: The decision maker is told that only s_j^K will occur. The only way the other s_i^K can occur is if he performs the intervention or spins an ‘influence wheel’.

Step 2 (elicit $\rho^\alpha(s_i^K, s_j^K)$ for $i \equiv j$): The decision maker is told if he spins the ‘influence wheel’ there is $x\%$ chance that s_j^K will occur and $(1 - x)\%$ chance one of the other states will occur. The value of x that makes the decision maker indifferent between spinning the wheel and performing the intervention is the value for $\rho^\alpha(s_j^K, s_j^K)$.

Step 3 (elicit $\rho^a(s_i^K, s_j^K)$ for $i \neq j$): The decision maker is told if he spins the ‘influence wheel’, there is $\rho^a(s_j^K, s_j^K)$ % chance that s_j^K will occur, x % chance s_i^K will occur and $(1 - \rho^a(s_j^K, s_j^K) - x)$ % chance one of the other states will occur. $\rho^a(s_i^K, s_j^K)$ is the value of x that makes the decision maker indifferent between spinning the wheel and performing the intervention.

Step 4 (repeat Step 3 for the other i s): The remaining $\rho^a(s_i^K, s_j^K)$ are elicited similarly in a sequential fashion, where for each subsequent elicitation, the number of outcomes of the influence wheel increases to include the states for which $\rho^a(s_i^K, s_j^K)$ are already elicited and their associated chance of occurring is held constant at $\rho^a(s_i^K, s_j^K)$. Repeat till $n - 1$ elicitations are obtained and then set the probability of the n^{th} state to $1 - \sum_{i=1}^{n-1} \rho^a(s_i^K, s_j^K)$.

Appendix 3-2: Proof for Theorem 1 (representation theorem for generalized generic controller)

\Leftarrow Suppose $P_a(S^K) = f_a(P(S^K))$ is a GGC. Let $P(S^K) = \alpha P'(S^K) + (1 - \alpha)P''(S^K)$ for some $P', P'', \alpha \in [0,1]$. $\forall s_i^K \in S^K$,

$$\begin{aligned}
 P_a(s_i^K) &= \sum_j P(s_j^K) \rho^a(s_i^K, s_j^K) && \text{definition of GGC} \\
 &= \sum_j (\alpha P'(s_j^K) + (1 - \alpha)P''(s_j^K)) \rho^a(s_i^K, s_j^K) && \text{assumption about } P, P', P'' \\
 &= \sum_j \alpha P'(s_j^K) \rho^a(s_i^K, s_j^K) + (1 - \alpha) \sum_j P''(s_j^K) \rho^a(s_i^K, s_j^K) && \text{by algebra} \\
 &= \alpha \sum_j P'(s_j^K) \rho^a(s_i^K, s_j^K) + (1 - \alpha) \sum_j P''(s_j^K) \rho^a(s_i^K, s_j^K) && \text{by algebra} \\
 &= \alpha P'_a(s_i^K) + (1 - \alpha)P''_a(s_i^K) && \text{GGC applied to } P', P''
 \end{aligned}$$

Therefore f_a is linear.

\Rightarrow Suppose $f_a: \Delta^{|S^K|} \rightarrow \Delta^{|S^K|}$ is a linear probability revision function.

To show that $P_a(S^K) = f_a(P(S^K))$ is a generalized generic controller it is sufficient to show that for some $0 \leq \rho^a(s_i^K, s_j^K) \leq 1$, $\sum_i \rho^a(s_i^K, s_j^K) = 1$;

$$P_a(s_i^K) = \sum_j P(s_j^K) \rho^a(s_i^K, s_j^K), \forall s_i^K \in S^K.$$

By iterated application of Definition 3-2 have, if $P(s_i^K) = \sum_j \alpha_j P^j(s_i^K)$, $\alpha_j \in [0,1]$ and

$$\sum_j \alpha_j = 1, \text{ then, } P_a(s_i^K) = \sum_j \alpha_j P_a^j(s_i^K).$$

Let $\mathbf{e}_i^K = \langle e(s_i^K, s_1^K), \dots, e(s_i^K, s_n^K) \rangle$ be the probability vector where $e(s_i^K, s_j^K) = 1$ for $i = j$, 0 otherwise and $f_a(\mathbf{e}_i^K) = \langle e_a(s_i^K, s_1^K), \dots, e_a(s_i^K, s_n^K) \rangle$. Since $f_a(\mathbf{e}_i^K)$ is a point in $\Delta^{|S^K|}$, we have $e_a(s_i^K, s_j^K) \in [0,1]$, $\sum_i e_a(s_i^K, s_j^K) = 1$.

$\forall s_i^K \in S^K$ we have,

$$P(s_i^K) = \sum_j P(s_j^K) e(s_i^K, s_j^K) \quad \text{definition of } e(s_i^K, s_j^K)$$

$$P_a(s_i^K) = \sum_j P(s_j^K) e_a(s_i^K, s_j^K) \quad \text{linearity assumption}$$

$$= \sum_j P(s_j^K) \rho^a(s_i^K, s_j^K) \quad \text{set } \rho^a(s_i^K, s_j^K) = e_a(s_i^K, s_j^K)$$

Therefore f_a is a generalized generic controller.

Appendix 3-3: The Dutch book for probability revision functions which are not Bayes conditionalization preserving

Here we present a general strategy for constructing a Dutch book for probability revision functions that do not satisfy property (ii) (Bayes conditionalization preserving), stated in section 3.4.1.

Statement of the Dutch book Theorem:

Suppose S^K is a state node, the probabilities of S^K depend on another state node S^R , $\mathbf{P}(S^K, S^R)$ is the joint probability distribution for $\{S^K, S^R\} = \{(s_i^K, s_j^R)\}$. If a probability revision function for intervention $f_a: \Delta^{|S^K||S^R|} \rightarrow \Delta^{|S^K||S^R|}$ is not Bayes conditionalization preserving (i.e for some $s_i^K, s_l^R, P_a(s_i^K | s_l^R) P(s_l^R) \neq P_a(s_i^K, s_l^R)$ where, $P_a(S^K, s_l^R) = f_a^l(\mathbf{P}(S^K, s_l^R))$, $P_a(S^K | s_l^R) = f_a^l(\mathbf{P}(S^K | s_l^R))$), then once the decision maker commits to performing the intervention, a bookie who knows no more or less than the decision maker can induce the decision maker to willingly buy and sell bets on the occurrence of S^R and S^K which will guarantee a net loss to the decision whatever happens.

In order to construct the proof we will assume that the subjective probabilities, $P(s_i^K)$, of a risk neutral decision maker represent the maximum (minimum) price he/she is willing to pay (accept) for a bet on S^K which pays 1 if s_i^K occurs and 0 otherwise. This bet is represented as follows:

$$\begin{aligned} s_i^K: & 1 \\ \bar{s}_i^K: & 0 \end{aligned}$$

Proof:

Suppose for some s_i^K , $P_a(s_i^K, s_i^R) \neq P_a(s_i^K | s_i^R) P(s_i^R)$ and $P_a(s_i^K, s_i^R) > P_a(s_i^K | s_i^R) P(s_i^R)$. Then, $P_a(s_i^K, s_i^R) = P_a(s_i^K | s_i^R) P(s_i^R) + \delta$, for some $\delta > 0$.

The bookie designs the following two bets on the occurrence of S^R and S^K :

Bet A	Bet B
$\overline{s_i^R}$: $P_a(s_i^K s_i^R)$	s_i^R, s_i^K : 1
s_i^R : 0	$\overline{s_i^R}, s_i^K$: 0
	$\overline{s_i^R}$: 0

Suppose a bookie sells Bet A and Bet B to the decision maker at the maximum price that the decision maker is willing to pay for the two bets, which is:

$$\begin{aligned} P_a(s_i^K | s_i^R) * P(\overline{s_i^R}) + P_a(s_i^K, s_i^R) * 1 &= P_a(s_i^K | s_i^R) * P(\overline{s_i^R}) + (P_a(s_i^K | s_i^R) P(s_i^R) + \delta) * 1 \\ &= P_a(s_i^K | s_i^R) + \delta \end{aligned}$$

If $\overline{s_i^R}$ occurs, then the payoff to the decision maker is $P_a(s_i^K | s_i^R) - (P_a(s_i^K | s_i^R) + \delta) = -\delta$.

If s_i^R occurs, then before the outcome of S^K is observed, the bookie buys back Bet B at the minimum price that the decision maker is willing to accept, which is $P_a(s_i^K | s_i^R)$. Then the payoff to the decision maker is $P_a(s_i^K | s_i^R) - (P_a(s_i^K | s_i^R) + \delta) = -\delta$. Thus, the decision maker incurs a net loss of $-\delta$ whatever happens.

If $P_a(s_i^K, s_i^R) < P_a(s_i^K | s_i^R) P(s_i^R)$, then the argument above can be reconstructed by having the bookie buy Bet A and Bet B from the decision maker and sell Bet B back to the decision maker depending on the outcome of state S^R .

Appendix 3-4: A necessary and sufficient condition for the basic properties of probability revision functions for interventions

Here we will prove that any probability revision function for interventions which commutes with Bayes rule satisfies *both* the basic properties stated in section 3.4.1 and vice-versa. First we will prove a necessary and sufficient condition for the first property (fixed-point at zero). We will then use this result to prove that ‘commutes with Bayes rule’ is a necessary and sufficient condition for the basic properties stated in section 3.4.1.

Theorem 3-3: Necessary and Sufficient condition for fixed-point at zero

Suppose S^K is a controllable state node with a direct predecessor $S^*(S^K) = S^R$, $\mathbf{P}(S^K, S^R)$ is any joint probability distribution for $\{S^K, S^R\} = \{(s_i^K, s_j^R)\}$, $f_a: \Delta^{|S^K||S^R|} \rightarrow \Delta^{|S^K||S^R|}$ is a probability revision function for an intervention on S^K and $g_{s_i^R}: \Delta^{|S^K||S^R|} \rightarrow \Delta^{|S^K||S^R|}$ represents Bayes conditioning with respect to

$s_i^R \in S^R$. Let $\mathbf{P}_a(s_i^K, s_j^R)$ represent the revised probabilities for $f_a^j \cdot \forall l$,
 $\sum_i P_a(s_i^K, s_l^R) = \sum_i P(s_i^K, s_l^R) \Leftrightarrow \forall l, f_a^l(\mathbf{0}) = \mathbf{0}$.

Proof:

\Rightarrow Suppose $\forall l, \sum_i P_a(s_i^K, s_l^R) = \sum_i P(s_i^K, s_l^R)$ but for some $m, f_a^m(\mathbf{0}) \neq \mathbf{0}$. Conditioning the distribution $P(S^K, S^R)$ on some $s_z^R \in S^R, z \neq m$ and using $P_{s_z^R}(s_i^K, s_j^R)$ to denote the revised probabilities, have $\mathbf{P}_{s_z^R}(S^K, s_m^R) = \mathbf{0}$. Now applying f_a to $\mathbf{P}_{s_z^R}(S^K, S^R)$ have
 $\mathbf{P}_{a, s_z^R}(S^K, s_m^R) = f_a^m(\mathbf{P}_{s_z^R}(S^K, s_m^R)) \neq \mathbf{0}$. $\sum_i P_{s_z^R}(s_i^K, s_m^R) = \mathbf{0} \neq \sum_i P_{a, s_z^R}(s_i^K, s_m^R)$,
 which is a contradiction. Therefore $\forall l, f_a^l(\mathbf{0}) = \mathbf{0}$.

\Leftarrow Suppose $\forall l, f_a^l(\mathbf{0}) = \mathbf{0}$ but for some $m, \sum_i P_a(s_i^K, s_m^R) \neq \sum_i P(s_i^K, s_m^R)$. Conditioning the distribution $P(S^K, S^R)$ on $s_m^R \in S^R$ and using $P_{s_m^R}(s_i^K, s_j^R)$ to denote the revised probabilities, have $\mathbf{P}_{s_m^R}(S^K, s_j^R) = \mathbf{0}$ for $j \neq m$. Since $P_{s_m^R}(s_i^K, s_j^R)$ are points in the probability simplex $\Delta^{|S^K||S^R|}$, $\sum_i P_{s_m^R}(s_i^K, s_m^R) = 1$. Applying f_a to $\mathbf{P}_{s_m^R}(S^K, S^R)$, have
 $\sum_i P_{a, s_m^R}(s_i^K, s_m^R) \neq \sum_i P_{s_m^R}(s_i^K, s_m^R) = 1$. This means for some $j \neq m, P_{a, s_m^R}(s_i^K, s_j^R) > 0$. But this is a contradiction since for $j \neq m, \mathbf{P}_{a, s_m^R}(S^K, s_j^R) = f_a^m(\mathbf{P}_{s_m^R}(S^K, s_j^R)) = f_a^m(\mathbf{0}) = \mathbf{0}$.

Theorem 3-3 states that when revising the probabilities of a controllable state node S^K with joint probabilities $P(s_i^K, s_j^R)$, the cumulative probability of each section of the joint distribution of $S_i^*(S^K), \forall i$ is conserved if and only if $\forall j, f_a^j(\mathbf{0}) = \mathbf{0}$.

Theorem 3-4: Equivalence between ‘commutes with Bayes rule’ and the basic properties of probability revision functions

Suppose S^K is a controllable state node, $S^*(S^K) = S^R$, $\mathbf{P}(S^K, S^R)$ is any joint probability distribution for $\{S^K, S^R\} = \{(s_i^K, s_j^R)\}$, $f_a: \Delta^{|S^K||S^R|} \rightarrow \Delta^{|S^K||S^R|}$ is a probability revision function for an intervention on S^K and $g_{s_i^R}: \Delta^{|S^K||S^R|} \rightarrow \Delta^{|S^K||S^R|}$ represents Bayes conditioning with respect to some $s_i^R \in S^R$.

$$\left(g_{s_j^R} \circ f_a\right)\left(\mathbf{P}(S^K, S^R)\right) = \left(f_a \circ g_{s_j^R}\right)\left(\mathbf{P}(S^K, S^R)\right), \forall s_j^R \in S^R \Leftrightarrow f_a^j(\mathbf{0}) = \mathbf{0},$$

$$\forall j \text{ and } f_a^j\left(\mathbf{P}(S^K, s_j^R)\right) = f_a^j\left(\mathbf{P}(S^K | s_j^R)\right) * P(s_j^R), \forall j.$$

Proof:

Let $P_a(s_i^K, s_j^R)$ represent the revised probabilities for f_a . We have $\mathbf{P}(S^K | s_i^R) = g_{s_i^R}(\mathbf{P}(S^K, S^R))$. Consider $(f_a \circ g_{s_i^R})$ and $(g_{s_i^R} \circ f_a)$ for any l .

$$(f_a \circ g_{s_l^R})(P(S^K, S^R)) = \begin{cases} f_a^j(\mathbf{P}(S^K | s_l^R)) & j = l \\ f_a^j(\mathbf{0}) & j \neq l \end{cases} \text{ and,}$$

$$(g_{s_l^R} \circ f_a)(P(S^K, S^R)) = \begin{cases} \frac{f_a^j(\mathbf{P}(S^K, s_l^R))}{\sum_i P_a(s_i^K, s_l^R)} & j = l \\ \mathbf{0} & j \neq l \end{cases}$$

For any l , $(g_{s_l^R} \circ f_a) = (f_a \circ g_{s_l^R}) \Leftrightarrow f_a^j(\mathbf{0}) = \mathbf{0}, \forall j \neq l$ and $f_a^j(\mathbf{P}(S^K | s_l^R)) = \frac{f_a^j(\mathbf{P}(S^K, s_l^R))}{\sum_i P_a(s_i^K, s_l^R)}$ for $j = l$.

Comparing the expression for $(g_{s_j^R} \circ f_a)$ and $(f_a \circ g_{s_j^R})$ for various s_j^R , have:

$$\Rightarrow (g_{s_j^R} \circ f_a) = (f_a \circ g_{s_j^R}), \forall j \Rightarrow f_a^j(\mathbf{0}) = \mathbf{0}, \forall j \quad \text{and} \quad f_a^j(\mathbf{P}(S^K | s_j^R)) = \frac{f_a^j(\mathbf{P}(S^K, s_j^R))}{\sum_j P_a(s_i^K, s_j^R)}, \forall j.$$

By Theorem 3-3, $\sum_i P_a(s_i^K, s_j^R) = \sum_i P(s_i^K, s_j^R) = P(s_j^R), \forall j$, therefore

$$f_a^j(\mathbf{P}(S^K | s_j^R)) P(s_j^R) = f_a^j(\mathbf{P}(S^K, s_j^R)), \forall j.$$

$$\Leftrightarrow f_a^j(\mathbf{0}) = \mathbf{0}, \forall j \quad \text{and} \quad f_a^j(\mathbf{P}(S^K | s_j^R)) P(s_j^R) = f_a^j(\mathbf{P}(S^K, s_j^R)), \forall j \Rightarrow f_a^j(\mathbf{0}) = \mathbf{0}, \forall j \quad \text{and}$$

$$f_a^j(\mathbf{P}(S^K | s_j^R)) = \frac{f_a^j(\mathbf{P}(S^K, s_j^R))}{\sum_j P_a(s_i^K, s_j^R)}, \forall j \Rightarrow (g_{s_j^R} \circ f_a) = (f_a \circ g_{s_j^R}), \forall j.$$

3.7 References for Chapter 3

1. Agarwal, S., 2014. The Moderating Effect of ‘Probability Control’ on Risk Taking.
2. Agarwal, S., Montibeller, G., 2014. A Quantitative Method for Measuring the Value of Operational Risk Mitigations.
3. Armendt, B., 1986. A Foundation for Causal Decision Theory. *Topoi* 5, 3–19.
4. Bolker, E., 1967. A simultaneous axiomatization of utility and subjective probability. *Philos. Sci.* 333–340.
5. Chelst, K., Bodily, S., 2000. Structured Risk Management: Filling a Gap in Decision Analysis Education. *J. Oper. Res. Soc.* 51, 1420–1432.
6. Courtney, H., 2001. *20/20 Foresight: Crafting Strategy in an Uncertain World*. Harvard Business School Press, Boston, MA.
7. Courtney, H., Kirkland, J., Viguerie, P., 1997. Strategy Under Uncertainty. *Harv. Bus. Rev.* 75, 67–79.
8. Gardenfors, P., 1982. Imaging and Conditionalization. *J. Philos.* 79, 747–760.
9. Gardenfors, P., 1988. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. MIT Press, Cambridge Massachusetts.
10. Gibbard, A., Harper, W., 1978. Counterfactuals and two kinds of expected utility, in: *Ifs*, (Eds.). Reidel, Dordrecht.

11. Heckerman, D., Shachter, R., 1994. A decision-based view of causality, in: Proceedings of the Tenth International Conference on Uncertainty in Artificial Intelligence. Morgan Kaufmann Publishers Inc., Seattle, WA, pp. 302–310.
12. Howard, R., 1990. From influence to relevance to knowledge., in: R.M.Oliver & J.Q. Smith, Influence Diagrams, Belief Nets and Decision Analysis. John Wiley & Sons Ltd.
13. Howard, R., Matheson, J., 1984. Influence Diagrams, in: R.A. Howard and J.E. Matheson (Eds.), The Principles and Applications of Decision Analysis. Strategic Decisions Group, Menlo Park, CA.
14. Howard, R., Matheson, J., 2005. Influence diagrams. *Decis. Anal.* 2, 127–143.
15. Jeffrey, R., 1965. The Logic of Decision, 2nd ed. University of Chicago Press, Chicago.
16. Joyce, J., 1999. The Foundations of Causal Decision Theory. Cambridge University Press, Cambridge Books Online.
17. Lewis, D., 1976. Probabilities of Conditionals and Conditional Probabilities. *Philos. Rev.* 85, 297–315.
18. Lewis, D., 1981. Causal decision theory. *Australas. J. Philos.* 59, 5–30.
19. Li, K., 2011. Preference towards control in risk taking: Control, no control, or randomize? *J. Risk Uncertain.* 43, 39–63.
20. Luce, D., Krantz, D., 1971. Conditional Expected Utility. *Econometrica* 39, 253–271.
21. March, J., Shapira, Z., 1987. Managerial perspectives on risk and risk taking. *Manag. Sci.* 33, 1404–1418.
22. Matheson, D., Matheson, J., 2005. Describing and Valuing Interventions That Observe or Control Decision Situations. *Decis. Anal.* 2, 165–81.
23. Matheson, J., 1990. Using influence diagrams to value information and control, in: R.M.Oliver & J.Q. Smith, Influence Diagrams, Belief Nets and Decision Analysis. John Wiley & Sons Ltd, pp. 25–48.
24. Pearl, J., 1994. A Probabilistic Calculus of Actions, in: Proceedings of the Tenth Annual Conference on Uncertainty in Artificial Intelligence (eds. R. Lopez de Mantaras and D. Poole). pp. 454–462.
25. Rosenzweig, P., 2013. What Makes Strategic Decisions Different. *Harv. Bus. Rev.* 89–93.
26. Rosenzweig, P., 2014. Left Brain, Right Stuff: How Leaders Make Winning Decisions. Public Affairs, USA.
27. Savage, L., 1954. The Foundations of Statistics. John Wiley & Sons, New York.
28. Shachter, R., 1986. Evaluating Influence Diagrams. *Oper. Res.* 34, 871–882.
29. Shachter, R., 2012. Effective Decision Analysis: A Cause to Believe In. *Decis. Anal. Today* 31, 29–33.

30. Shachter, R., Heckerman, D., 2010. Pearl Causality and Value of Control, in: R. Dechter, H. Geffner and J. Y. Halpern (Eds.), *Heuristics, Probability, and Causality: A Tribute to Judea Pearl*. College Publications, pp. 431–447.
31. Shephard, G., Kirkwood, C., 1994. Managing the judgmental probability elicitation process: a case study of analyst/manager interaction. *Eng. Manag. IEEE Trans. On* 41, 414–425.
32. Skyrms, B., 1982. Causal Decision Theory. *J. Philos.* 79, 695–711.
33. Stalnaker, R., 1968. A theory of conditionals, in: *Studies in Logical Theory*, American Philosophical Quarterly Monograph Series.
34. Sudarsanam, S., 2003. *Creating Value from Mergers and Acquisitions*, 2nd ed. Prentice-Hall.
35. Teller, P., 1973. Conditionalization and observation. *Synthese* 26, 218–258.
36. Von Neumann, J., Morgenstern, O., 1944. *Theory of Games and Economic Behavior*. Princeton University Press.
37. Von Winterfeldt, D., Kavet, R., Peck, S., Mohan, M., Hazen, G., 2012. The Value of Environmental Information without Control of Subsequent Decisions. *Risk Anal.* 32, 2113–2132.
38. Young, D., Goodie, A., Hall, D., 2011. Modeling the Impact of Control

3.8 Supplementary material for Chapter 3


S 3-1: A discussion of the conceptual issues with non-canonical influence diagrams

A normative objection of non-canonical influence diagrams is that the expected utility calculation of equation (3-1) to evaluate interventions is incompatible with the normative principles of classical expected utility paradigms (Savage, 1954; von Neumann and Morgenstern, 1944), which require the actions and states to be probabilistically independent (Matheson 1990). In particular, not all axioms of these classical expected utility paradigms, that are necessary to show that expected utility calculations of decisions represent qualitative preferences over decisions, are satisfied. From a Decision Analysis perspective, in influence diagrams which contain interventions, it is also not possible to represent value of information analysis on the state node that can be influenced as it creates a forbidden loop (Matheson 1990).

Some approaches have been proposed to resolve this conflict between normative foundations of expected utility calculations and influence diagrams that are not in the canonical form. One approach in Decision Analysis is to convert the influence diagram in the canonical form by making the ‘controllable’ state node deterministic and adding new chance

nodes (Figure 3-10), where each new node represents the probabilities conditional on each decision (Matheson 1990).

Figure 3-10: Revised influence diagram for the market share example (in the canonical form)



Value of information analysis on the controllable state node can then be represented by introducing arcs from the new nodes to the interventions. Conceptually, this means that the decision maker first commits to a decision and then evaluates the value of information given the selected decision.

Another method to make the state and decision node probabilistically independent of one another (i.e. remove the link between the decision node) is to repartition the state node (or revise how it is described) (Jeffrey 1965, p.22; Gibbard & Harper 1978). In the example in Figure 3-1, suppose there are three atomic states for the state node ($S^M = \{s_i^M\}$): gain no marketshare (s_o^M), gain partial marketshare (s_p^M), or gain full market share (s_f^M). Then the atomic states for the reformulated problem would be $S^M = \{s_{if\ a}^M\} \cup \{s_{if\ \neg a}^M\} \cup \{s_{if\ av\ \neg a}^M\}$, where $s_{if\ a}^M$ is the state that s_i^M will occur if the firm advertises and one of other states will occur if it does not advertise; and $s_{if\ av\ \neg a}^M$ is the state that s_i^M will occur regardless of what the firm does. From a practical perspective, this method is tedious — for example, compared to the approach of directly eliciting probabilities conditional on decisions, this new formulation requires a higher number of elicitations (nine instead of six). In particular, for n atomic states of the controllable node, k options and $m \geq 1$ atomic states of the parents of the controllable state node direct elicitation of post-intervention conditional probabilities $P(s_i|A)$ requires $m * (n - 1) * k$ elicitations, which is always fewer than repartitioning the state node which requires $m * (n^k - 1)$ elicitations. A numerical comparison is offered in Table 3-4.

Table 3-4: Number of elicitation required for each approach.

Atomic States n	Options k	Parent Atomic States m	Reformulated problem $m * (n^k - 1)$	Direct elicitation $(n - 1) * k * m$
3	2	1	8	4
3	2	2	16	8
3	2	3	24	12
3	2	4	32	16
3	3	1	26	6
3	3	2	52	12
3	3	5	130	30
3	3	10	260	60

S 3-2: The Dutch book argument for Bayes rule**Statement of the Dutch book Theorem:**

Suppose S^K is a state node, the probabilities of S^K depend on another state node S^R , $P(S^K, S^R)$ is the joint probability distribution for $\{S^K, S^R\} = \{(s_i^K, s_j^R)\}$. Upon receiving information about the state node S^R , if a decision maker's new beliefs about S^K , $P'(S^K)$, do not satisfy Bayes rule (i.e. for any s_i^R that is known to occur, for some s_i^K , $P'(s_i^K) \neq P(s_i^K | s_i^R) = P(s_i^K, s_i^R) / P(s_i^R)$), a bookie who knows no more or less than the decision maker can induce the decision maker to willingly buy and sell bets on the occurrence of S^R and S^K which will guarantee a net loss to the decision maker whatever happens.

In order to construct the proof we assume that the subjective probabilities, $P(s_i^K)$, of a risk neutral decision maker represent the maximum (minimum) price he/she is willing to pay (accept) for a bet on S^K which pays 1 if s_i^K occurs and 0 otherwise. This bet is represented as follows:

$$\begin{aligned} s_i^K: & 1 \\ \bar{s}_i^K: & 0 \end{aligned}$$

Proof:

Suppose for some s_i^K , $P'(s_i^K) \neq P(s_i^K | s_i^R)$ and $P'(s_i^K) < P(s_i^K | s_i^R)$. Then $P'(s_i^K) = P(s_i^K | s_i^R) - \delta$, for some $\delta > 0$.

Consider the following two bets on the occurrence of S^R and S^K :

Bet A	Bet B
$\bar{s}_i^R: P'(s_i^K)$	$s_i^R, s_i^K: 1$
$s_i^R: 0$	$\bar{s}_i^R, s_i^K: 0$
	$\bar{s}_i^K: 0$

Suppose a bookie sells Bet A and Bet B to the decision maker at the maximum price that the decision maker is willing to pay for the two bets, which is:

$$\begin{aligned} P'(s_i^K) * P(\overline{s_l^R}) + 1 * P(s_i^K, s_l^R) &= (P(s_i^K | s_l^R) - \delta) * P(\overline{s_l^R}) + P(s_i^K | s_l^R) * P(s_l^R) * 1 \\ &= (P(s_i^K | s_l^R) - \delta) * (1 - P(s_l^R)) + P(s_i^K | s_l^R) * P(s_l^R) * \\ &= -\delta (1 - P(s_l^R)) + P(s_i^K | s_l^R) = -\delta P(\overline{s_l^R}) + P(s_i^K | s_l^R) \end{aligned}$$

If $\overline{s_l^R}$ occurs, then the payoff to the decision maker is $P'(s_i^K) - (-\delta P(\overline{s_l^R}) + P(s_i^K | s_l^R)) = P(s_i^K | s_l^R) - \delta + \delta P(\overline{s_l^R}) - P(s_i^K | s_l^R) = -\delta P(s_l^R)$.

If s_l^R occurs, then before the outcome of S^K is observed, the bookie buys back Bet B at the minimum price that the decision maker is willing to accept, which is $P'(s_i^K)$. Then the payoff to the decision maker is $P'(s_i^K) - (-\delta P(\overline{s_l^R}) + P(s_i^K | s_l^R)) = -\delta P(s_l^R)$. Thus, the decision maker incurs a net loss of $\delta P(s_l^R)$ no matter what happens.

If $P'(s_i^K) > P'(s_i^K | s_l^R)$, then the argument above can be reconstructed by having the bookie buy Bet A and Bet B from the decision maker and sell Bet B back to the decision maker depending on the outcome of state S^R .

S 3-3: Expected utility calculations in influence diagrams

First we will discuss the expected utility calculation for an influence diagram where all decisions are made at the same time ($\mathcal{A}(A^K) = \emptyset, \forall A^K$) and decisions are not conditional ($\mathcal{S}^*(A^K) = \emptyset, \forall A^K$). An oriented influence diagram G , where the payoffs depend on both chance and decisions, i.e. $\mathcal{S}(V) \neq \emptyset, (\mathcal{A}(V) \cup \mathcal{A}(S^j)) \neq \emptyset$, for some S^j (and for which $(\mathcal{A}(A^K) \cup \mathcal{S}^*(A^K)) = \emptyset, \forall A^K$), can be evaluated directly using an expected utility calculation. Given the real value mappings for V and \mathcal{S} , the solution of the influence diagram is:

$$\max_{\mathcal{A}_i} EU(\mathcal{A}_i) = \sum_k P_{\mathcal{A}_i}(\mathcal{S}_k) * U(\mathcal{A}_i, \mathcal{S}_k) \quad (3-11)$$

where \mathcal{A}_i are the policies and \mathcal{S}_k are the events. The joint probabilities, $P_{\mathcal{A}_i}(\mathcal{S}_k)$, of the event \mathcal{S}_k can be determined from the conditional probabilities using the chain rule:

$$P_{\mathcal{A}_i}(\mathcal{S}_k) = \prod_j P_{\mathcal{A}_i^*(S^j)}(s_{k_j}^j | \mathcal{S}_k^*(S^j)) \quad (3-12)$$

For an influence diagram in the canonical form ($\mathcal{A}^*(S^j) = \emptyset, \forall j$), $\forall i, j, P_{\mathcal{A}_i^*(S^j)}(s_{k_j}^j | \mathcal{S}_k^*(S^j)) = P(s_{k_j}^j | \mathcal{S}_k^*(S^j))$. In Figure 3-2, consider a policy $\mathcal{A}_1 = (a_1^D, a_1^T)$ and the events $\mathcal{S} = \{(s_i^C, s_j^E, s_k^M) | i = 1, 2, \dots, |S^C|, j = 1, 2, \dots, |S^E|, k = 1, 2, \dots, |S^M|\}$. The

expected utility of this policy is given by

$$EU(\mathcal{A}_1) = EU(a_1^D, a_1^T) = \sum_{i,j,k} P_{\mathcal{A}_1}(s_i^C, s_j^E, s_k^M) * U(\mathcal{A}_1, (s_i^C, s_j^E, s_k^M))$$
 and

$$P_{\mathcal{A}_1}(s_i^C, s_j^E, s_k^M) = P_{a_1^D}(s_k^M | s_i^C) P_{a_1^T}(s_i^C | s_j^E) P(s_j^E).$$

Equation (3-12) can be generalized to calculate the optimal policies for influence diagrams which contain sequential decisions ($\mathcal{A}(A^K) \neq \emptyset$ for some A^K) and/or decisions that are conditional ($\mathcal{S}^*(A^K) \neq \emptyset$ for some A^K), i.e. the information state is different for various decisions. The procedure involves iterated application of equation (3-12) to restricted realizations of the influence diagram, corresponding to the known state of information at the time a particular decision is made. The optimal policy for such influence diagrams can be evaluated using the following procedure:

- i. First order the decisions in the set \mathcal{A} by temporal precedence $\mathcal{A}^1, \mathcal{A}^2, \dots, \mathcal{A}^n$ where \mathcal{A}^i is the set of decisions that is performed before the set of decisions \mathcal{A}^{i+1} .
- ii. Evaluate the *optimal conditional policy* for the *latest* set of decisions \mathcal{A}^x as follows: Let $\mathcal{A}' = \cup_{A^j \in \mathcal{A}^x} \mathcal{A}(A^j)$ represent the decisions and $\mathcal{S}' = \cup_{A^j \in \mathcal{A}^x} \mathcal{S}(A^j)$ represent the states that will be known at the time decisions \mathcal{A}^x need to be made. Let $\Omega = \mathcal{A}' \cup \mathcal{S}'$ represent the *information state* and $\Omega_i = \mathcal{A}'_i \cup \mathcal{S}'_i$ an event and policy in Ω . Let $\mathcal{S}^x = \mathcal{S} - \mathcal{S}'$ be the states that will be unknown at the time decisions \mathcal{A}^x need to be made. Define the optimal conditional policy for \mathcal{A}^x (conditional on Ω_j) as:

$$\max_{\mathcal{A}^x_i} EU(\mathcal{A}^x_i | \Omega_j) = \sum_k P_{\mathcal{A}^x_i, \mathcal{A}'_j}(\mathcal{S}^x_k | \mathcal{S}'_j) * U(\mathcal{A}^x_i \cup \mathcal{A}'_j, \mathcal{S}^x_k \cup \mathcal{S}'_j) \quad (3-13)$$


For each information state, Ω_j , let $EU(\mathcal{A}^x_{\Omega_j} | \Omega_j)$ represent the maximum expected utilities that corresponds to the optimal conditional policies $\mathcal{A}^x_{\Omega_j}$.

- iii. Once all the optimal conditional policies for \mathcal{A}^x (corresponding to each Ω_j) have been evaluated, revise the real value mapping for the value node to $U: \prod_{\mathcal{A}'} \mathcal{A}^j \times \prod_{\mathcal{S}'} \mathcal{S}^j \rightarrow \mathbb{R}$ where $U(\mathcal{A}'_i, \mathcal{S}'_j) = EU(\mathcal{A}^x_{\Omega_j} | \Omega_j)$.
- iv. Revise the influence diagram by removing the nodes \mathcal{S}^x and \mathcal{A}^x from the influence diagram and introduce arcs from their direct predecessors, to the value node, V , i.e. introduce $(N^j, V) \forall N^j \in (\mathcal{A}^*(N^k) \cup \mathcal{S}^*(N^k)), \forall N^k \in (\mathcal{A}^x \cup \mathcal{S}^x)$.
- v. Let Ω_j^x denote the information states for the optimal conditional policies $\mathcal{A}^x_{\Omega_j}$. The set of all optimal conditional policies are $\mathcal{A}^*_{\Omega_j} = \cup \{ \mathcal{A}^x_{\Omega_j^{(x+m)*}}; 0 \leq m \leq n-x \mid \mathcal{A}^x_{\Omega_j} \in \Omega_j^{x+m} \text{ for some } j \}$.

Repeat ii—v for remaining \mathcal{A}^i to determine the optimal policies corresponding to the other decisions.

Figure 3-11 illustrates this procedure for a hypothetical influence diagram.

Figure 3-11: Evaluation of a hypothetical influence diagram (extended version of Figure 3-2 with sequential decisions). The shaded nodes correspond to the information state Ω when the decision A^j needs to be made. The unshaded nodes are the variables for which the expected utility is calculated (for fixed Ω_j) in each iteration.




9(a): The original influence diagram.

$$\mathcal{A}^1 = A^E, \mathcal{A}^2 = A^T, \mathcal{A}^3 = A^D.$$

Iteration 1/3:


Since A^D is the latest decision, the optimal conditional policy will be evaluated for A^D first.



9(b): Modified influence diagram after conditional policies for A^D have been determined and the real value mapping for the value node is redefined.

Iteration 2/3:

Since A^T is the latest decision, the optimal conditional policy will be evaluated for A^T .



9(c): Influence diagram after conditional policies for A^D and A^T have been determined and the real value mapping for the value node is redefined.

Iteration 3/3:

Since A^E is the only decision, the optimal conditional policy will be evaluated for A^E .

S 3-4: Order invariance for Bayes rule

Suppose $\mathcal{S}(S^K) = \{S^R, A\}$, $S^R = \{s_i^R\}$, $A = \{a, \neg a\}$.

Let $f_a = \langle f_a^1, f_a^2, \dots, f_a^{|S^R|}, f_a^{|S^R|+1}, f_a^{|S^R|+2}, \dots, f_a^{2|S^R|} \rangle: \Delta^{2|S^K||S^R|} \rightarrow \Delta^{2|S^K||S^R|}$ represent Bayes conditioning with respect to the variable $a \in A$ and let

$g_{s_l^R} = \left\langle g_{s_l^R}^1, g_{s_l^R}^2, \dots, g_{s_l^R}^{|S^R|}, g_{s_l^R}^{|S^R|+1}, g_{s_l^R}^{|S^R|+2}, \dots, g_{s_l^R}^{2|S^R|} \right\rangle: \Delta^{2|S^K||S^R|} \rightarrow \Delta^{2|S^K||S^R|}$ represent

Bayes conditioning with respect to the variable s_l^R .

Suppose the functions f_a^i and $g_{s_l^R}^i$ are indexed so that $a \in \mathcal{S}_i(S^K)$ for $i \leq |S^R|$ and $\neg a \in \mathcal{S}_i(S^K)$ for $i > |S^R|$ and $s_l^R \in \mathcal{S}_i(S^K)$ and $\mathcal{S}_{|S^R|+i}(S^K)$.

$\forall (S^K, s_j^R, a_k), 1 \leq j \leq |S^R|,$

$$f_a(\mathbf{P}(S^K, s_j^R, a_k)) = \begin{cases} f_a^j(\mathbf{P}(S^K, s_j^R, a)) = \frac{\mathbf{P}(S^K, s_j^R, a)}{\sum_j \sum_i \mathbf{P}(s_i^K, s_j^R, a)} = \frac{\mathbf{P}(S^K, s_j^R, a)}{P(a)}; & a_k = a \\ f_a^{|S^R|+j}(\mathbf{P}(S^K, s_j^R, \neg a)) = 0; & a_k = \neg a \end{cases}$$

$$\begin{aligned} (g_{s_l^R} \circ f_a)(\mathbf{P}(S^K, s_j^R, a_k)) &= \begin{cases} g_{s_l^R} \left(\frac{\mathbf{P}(S^K, s_j^R, a)}{P(a)} \right); & a_k = a \\ 0; & a_k = \neg a \end{cases} \\ &= \begin{cases} \frac{\mathbf{P}(S^K, s_j^R, a)/P(a)}{\sum_i \mathbf{P}(s_i^K, s_j^R, a)/P(a)} = \frac{\mathbf{P}(S^K, s_j^R, a)}{P(s_j^R, a)}; & a_k = a, j = l \\ \mathbf{0}; & a_k = a, j \neq l \\ \mathbf{0}; & a_k = \neg a \end{cases} \end{aligned}$$

$$\underline{\text{For } j = l, a_k = a:} \frac{\mathbf{P}(S^K, s_j^R, a)}{P(s_j^R, a)} = \frac{\mathbf{P}(S^K, s_j^R, a)/P(s_j^R)}{\sum_i (\mathbf{P}(s_i^K, s_j^R, a)/P(s_j^R))} = f_a \left(\frac{\mathbf{P}(S^K, s_j^R, a)}{P(s_j^R)} \right) = (f_a \circ g_{s_l^R})(\mathbf{P}(S^K, s_j^R, a)).$$

$$\underline{\text{For } j \neq l, a_k = a:} g_{s_l^R}(\mathbf{P}(S^K, s_j^R, a)) = \mathbf{0}; \text{ therefore, } (f_a \circ g_{s_l^R})(\mathbf{P}(S^K, s_j^R, a)) = \mathbf{0}.$$

$$\underline{\text{For } a_k = \neg a:} (f_a \circ g_{s_l^R})(\mathbf{P}(S^K, s_j^R, \neg a)) = \mathbf{0}.$$

$$\text{Therefore, } (g_{s_l^R} \circ f_a)(\mathbf{P}(S^K, s_j^R, a)) = (f_a \circ g_{s_l^R})(\mathbf{P}(S^K, s_j^R, a)).$$

S 3-5: Illustrating the elicitation of the GGC through a conversation

The following hypothetical conversation between an analyst and a manager demonstrates the elicitation of the generalized generic controller for the market share example.

Part 1/3

Analyst: Suppose I told you I have information that your company will have no market share for this product if you don't advertise.

Manager: But if I launch a good advertising campaign then surely there is a chance that I will have some market share?

Analyst: True. In fact there is another way you can possibly avoid having no market share: it is by spinning the 'influence wheel'. You can see the influence wheel has two colours — red and blue. [Analyst sets the influence wheel so that exactly half the circle is coloured blue and the other half coloured red]. Upon spinning the wheel, if the pointer points in the blue area

then you will avoid no market share else you will have no market share. I can adjust the portion of the wheel that is red and blue, but let us begin with this setting, which is a fifty-fifty chance the pointer will point to a red or blue area.

Manager: So, the option is between spinning this wheel you are showing me or advertising?

Analyst: Correct. You can either spin the wheel at this current setting or advertise. What would you prefer to do?

Manager: I think the chances are better if I advertise, so I prefer to advertise.

[Analyst changes the setting of the influence wheel so that exactly one quadrant of the circle is coloured red].

Analyst: What about spinning the wheel with this new setting vs. advertising? Would you still prefer to advertise?

Manager: Probably not. I think I would try my luck with spinning the wheel at this new setting.

[Analyst changes the setting of the influence wheel so that exactly 65% is coloured blue and the rest is coloured red].

Analyst: What about the choice between spinning this wheel vs. advertising?

Manager: Difficult to say — guess I would advertise.

[Analyst stops. Coefficient $\rho^a(s_o^M, s_o^M)$ is between 0.35 and 0.25. Analyst picks mid-point: 0.3].

Analyst: Looks like you feel there is around seventy percent chance that advertising will lead your company to avoid no market share?

Manager: Yes, that sounds about right.

Analyst: Do you feel there is any chance that if you avoid no market share, you will be able to obtain full market share? Remember if you don't advertise you will have no market share.

Manager: Probably. But the chances are small if without advertising the outcome is going to be no market share.

Analyst: I have another influence wheel with three colours red, orange and green. If upon spinning the wheel, the pointer points to red your company gets no market share, if it points to orange your company gets partial market share and if it lands on green your company gets full market share. [Analyst sets the wheel so that 30% is red, 35 % is orange and 35% is green]. What do you think — spin the wheel or advertise?

Manager: I would definitely spin the wheel. If I advertise I think the chance of getting full market share is about a quarter that of getting partial market share.

[Analyst sets the wheel so that 30% is red, 55 % is orange and 15% is green].

Analyst: Now?

Manager: Not sure — I'd be fine with either but I guess it would be better to advertise.

[Analyst stops and records $\rho^a(s_o^M, s_o^M) = 0.30$, $\rho^a(s_o^M, s_p^M) = 0.55$ and $\rho^a(s_o^M, s_f^M) = 0.15$ as the coefficients corresponding to no market share for the GGC.]

Part 2/3

Analyst: Now suppose I tell you that there is new information—your company will definitely get partial market share if you don't advertise. How does that sound?

Manager: That's good news but we could probably get full market share if our advertising campaign was good.

Analyst: Sure. But do you think your ad-campaign could have a damaging effect, leading you to lose market share?

Manager: No, that will definitely not happen. The ad-campaign can only improve the chance of better market share.

Analyst: I see. So what are the chances the ad-campaign will lead to full market share?

Manager: I'd say the chances are fifty-fifty.

[Analyst sets the two-coloured wheel so that 50% is red and 50% is blue].

Analyst: So you are indifferent between spinning this wheel and advertising, if a red outcome on the wheel meant you would stay with partial market share?

Manager: Not sure. I think I'll spin the wheel.

[Analyst sets the two-coloured wheel so that 60% is red and 40% is blue].

Analyst: Now?

Manager: I guess I would be okay with either spinning the wheel or advertising.

Analyst: So you feel that the chances you will obtain full market share, if you advertise are slightly lower than fifty percent, around forty percent?

Manager: Yes.

[Analyst records $\rho^a(s_p^M, s_o^M) = 0$, $\rho^a(s_p^M, s_p^M) = 0.60$ and $\rho^a(s_p^M, s_f^M) = 0.40$ as the coefficients corresponding to partial market share for the GGC.]

Part 3/3

Analyst: Now suppose I tell you I have information that you will get full market share if you don't advertise, would you advertise at all?

Manager: Well, like I mentioned earlier it won't hurt to advertise. Depending on the budget, we may still launch a small ad-campaign to keep up the brand image and boost demand.

[Analyst records $\rho^a(s_o^M, s_f^M) = 0$, $\rho^a(s_p^M, s_f^M) = 0$ and $\rho^a(s_f^M, s_f^M) = 1$ as the coefficients corresponding to partial market share for the GGC.]

Analyst: Thank you. We now have all the numbers to calculate how beneficial it will be for your company to advertise to improve the chances of higher market share.

S 3-6: Bayes rule and the linearity condition

Bayes rule and GGC are distinct classes of probability revision functions since, in general, Bayes rule does not satisfy the linearity condition. However, under some special circumstances, Bayes rule satisfies the linearity property. This is discussed next.

Consider a joint probability distribution $P(s_i^K, s_j^R)$ for two variables S^K and S^R . Suppose $P(s_i^K, s_j^R) = \alpha P'(s_i^K, s_j^R) + (1 - \alpha)P''(s_i^K, s_j^R)$. The probability revision obtained by Bayes conditioning on s_l^R is:

lhs:

$$\begin{aligned} g_{s_l^R}(P(s_i^K, s_j^R)) &= P(s_i^K | s_l^R) \\ &= \frac{P(s_i^K, s_l^R)}{\sum_i P(s_i^K, s_l^R)} \\ &= \frac{\alpha P'(s_i^K, s_l^R) + (1 - \alpha)P''(s_i^K, s_l^R)}{\alpha \sum_i P'(s_i^K, s_l^R) + (1 - \alpha) \sum_i P''(s_i^K, s_l^R)} \\ &= \frac{\alpha P'(s_i^K, s_l^R) + (1 - \alpha)P''(s_i^K, s_l^R)}{\alpha P'(s_l^R) + (1 - \alpha)P''(s_l^R)} \end{aligned}$$

rhs:

$$\begin{aligned} \alpha g_{s_l^R}(P'(s_i^K, s_j^R)) + (1 - \alpha)g_{s_l^R}(P''(s_i^K, s_j^R)) \\ &= \alpha P'(s_i^K | s_l^R) + (1 - \alpha)P''(s_i^K | s_l^R) \\ &= \alpha \frac{P'(s_i^K, s_l^R)}{\sum_i P'(s_i^K, s_l^R)} + (1 - \alpha) \frac{P''(s_i^K, s_l^R)}{\sum_i P''(s_i^K, s_l^R)} \\ &= \alpha \frac{P'(s_i^K, s_l^R)}{P'(s_l^R)} + (1 - \alpha) \frac{P''(s_i^K, s_l^R)}{P''(s_l^R)} \end{aligned}$$

In general, the lhs (quotient of mixtures) will not be equal to the rhs (mixture of quotients). However when $P'(s_l^R) = P''(s_l^R)$, then the lhs will be equal to the rhs.

One natural situation where $P'(s_l^R) = P''(s_l^R)$ is as follows: Suppose in Figure 3-9 we replace the decision node A^D with a state node A^D (where $A^D = \{a^D, \neg a^D\}$) and describe the probability revision of S^M using Bayes rule. Setting $P'(s_i^M, s_j^C) = P(s_i^M, s_j^C | a^D)$, $P''(s_i^M, s_j^C) = P(s_i^M, s_j^C | \neg a^D)$ and $\alpha = P(a^D)$, we have Bayes rule satisfies the linearity condition when $\forall j, P'(s_j^C) = P(s_j^C | a^D) = P(s_j^C | \neg a^D) = P''(s_j^C)$ i.e. when A^D and S^C are independent. In an influence diagram this holds if S^C and A^D don't have any common parents i.e. $\mathcal{S}(S^C) \cap \mathcal{S}(A^D) = \emptyset$.

S 3-7: Post-intervention inference in a non-canonical influence diagram

In an influence diagram which is not in the canonical form, the reversibility of the arcs holds only between state nodes which are not controllable. If $\mathcal{A}^*(S^K) \neq \emptyset$, $S^R \in \mathcal{S}^*(S^K)$ and for simplicity suppose $\mathcal{A}^*(S^R) = \emptyset$, then the arc between S^K and S^R cannot be reversed because $P_{\mathcal{A}_i^*(S^K)}(s_i^K | s_j^R) * P(s_j^R)$ need not be equal to $P(s_j^R | s_i^K) * P_{\mathcal{A}_i^*(S^K)}(s_i^K)$ and the expected utility calculation will be different for the two influence diagrams. To see this, consider the joint distribution in Box 3 for the market share example and a probability revision function described by the generic controller in Box 1 ($q = 0.5$). The conditional probability distribution for the market share node, depending on whether or not the competitor advertises is given below in Table 3-5.

Table 3-5: Probabilities of market share conditional on the competitor’s decision to advertise or not.

<i>Conditional probabilities for market share</i>	Competitor does not advertise (s_{-d}^C)		Competitor advertises (s_d^C)		Marginal probability	
	$P_{-a^D}(\cdot s_{-d}^C)$	$P_{a^D}(\cdot s_{-d}^C)$	$P_{-a^D}(\cdot s_d^C)$	$P_{a^D}(\cdot s_d^C)$	$P_{-a^D}(\cdot)$	$P_{a^D}(\cdot)$
No market share (s_o^M)	0.10	0.05	0.30	0.15	0.2	0.1
Partial market share (s_p^M)	0.60	0.30	0.60	0.30	0.6	0.3
Full market share (s_f^M)	0.30	0.65	0.10	0.55	0.2	0.6

From Box 3, the probability that the competitor advertises is 0.5. The probability of the competitor node conditional on the market share node can be derived from the pre-intervention distribution $P(s_j^C | s_i^M) = P_{-a^D}(s_i^M | s_j^C) * P(s_j^C) / P_{-a^D}(s_i^M)$ and is shown in Table 3-6.

Table 3-6: Probabilities of competitor advertising, conditional on market share.

<i>Conditional probabilities for the competitor (pre-intervention)</i>	Competitor does not advertise ($\neg s_d^C$)	Competitor advertises (s_d^C)
No market share (s_o^M) $P(\cdot s_o^M)$	0.25	0.75
Partial market share (s_p^M) $P(\cdot s_p^M)$	0.50	0.50
Full market share (s_f^M) $P(\cdot s_f^M)$	0.75	0.25

In the original influence diagram, where the arc is from S^C to S^M , the post intervention joint distribution is given by: $P_{a^D}(S^M, S^C) = [P_{a^D}(S^M | s_d^C) * P(s_d^C), P_{a^D}(S^M | s_{-d}^C) * P(s_{-d}^C)]$. However, if the arc between S^M and S^C is reversed then, $P_{a^D}(S^M, S^C) = [P(S^C | s_o^M) *$

$P_{aD}(s_o^M), \mathbf{P}(S^C | s_p^M) * P_{aD}(s_p^M), \mathbf{P}(S^C | s_f^M) * P_{aD}(s_f^M)$]. The two approaches for calculating the joint distribution, $\mathbf{P}_{aD}(S^M, S^C)$, shown in Table 3-7, are not equal.

Table 3-7: Post-intervention joint distribution for $\{S^M, S^C\}$ depending on the direction of the arc between S^M and S^C .

$\mathbf{P}_{aD}(S^M, S^C)$	$\mathbf{P}_{aD}(S^M s_i^C) * \mathbf{P}(s_i^C)$		$\mathbf{P}(S^C s_i^M) * \mathbf{P}_{aD}(s_i^M)$	
	Competitor does not advertise ($\neg s_d^C$)	Competitor advertises (s_d^C)	Competitor does not advertise ($\neg s_d^C$)	Competitor advertises (s_d^C)
No market share (s_o^M)	0.025	0.075	0.05	0.15
Partial market share (s_p^M)	0.15	0.15	0.30	0.30
Full market share (s_f^M)	0.325	0.275	0.15	0.05

Thus, reversing the arc between two state nodes, where one node is controllable, may not preserve the utility calculation of the influence diagram.

PRELUDE TO CHAPTER 4

In the previous chapter, a general probability revision for modelling interventions was proposed. Its advantage in terms of alleviating the judgemental burden when modelling the potential effects of actions on uncertainty was exposed. The feasibility of using the proposed procedure in practice needs to be explored.

This chapter endeavours to leverage the proposed practical benefits of the proposed probability revision rule in the context of Operational Risk assessments. Operational Risk assessments are estimations of potentially extreme losses that can arise from the uncertainty about various events such as earthquakes, breakdown of IT systems, staff strikes etc. The uncertainty for some of these events (e.g. staff strikes, IT system breakdown) is, to some extent, controllable and managers endeavour to mitigate the likelihood or severity of these events to curtail the resulting losses.

In this chapter, it will be discussed how influence diagrams can be used to model and represent the Operational Risk for a given loss event. An operationalization of the probability revision rule discussed in chapter 3 will be presented and applied to model risk mitigations. Using a real world case study that was conducted for a leading insurance company, it is shown how this operationalization can be integrated with existing Operational Risk assessment methods to explicitly model the impact of mitigations on Operational Risk costs. The proposed integration can enable the impact of different mitigations to be compared tractably and also simplify the process of updating Operational Risk assessments to reflect changes in managerial policies.

Chapter 4

A Quantitative Method for Measuring the Value of Operational Risk Mitigations

Shweta Agarwal*, Gilberto Montibeller*

ABSTRACT

The estimation of costs associated with Operational Risks is a regulatory initiative to protect the solvency of financial organizations and to motivate better management processes for mitigating risks associated with ‘loss’ events. A sophisticated way to measure Operational Risk quantitatively is based on calculating the Value-at-Risk of the loss distribution using the Loss Distribution Approach (LDA). However, in this approach, it is difficult to capture the effect of all types of managerial mitigations on loss distributions effectively. While the LDA can deal with mitigations that reduce the impact of an Operational Risk event (e.g. purchasing of insurance), analysing the effect of mitigations that alter uncertainty or probabilities of the loss-generating event (e.g. managerial policies to prevent staff strike) can be tedious within the existing statistical framework. In this paper we thus propose a new approach for modelling and measuring the value of Operational Risk mitigations. We show how the concept of ‘controllers’ in Decision Analysis can be integrated with the LDA to model the impact of mitigations on uncertain events explicitly. We then operationalize this concept as a probability revision matrix, which encodes the effect of mitigations on probability distributions quantitatively. The theoretical properties that characterize this approach leads to a less burdensome way to calculate the effect of mitigation on Operational Risk costs, thereby providing a tractable way to compare the value of various mitigations. We also present a real world application of this approach for a specific Operational Risk scenario, staff strike, conducted for a leading insurance company.

Key words: Operational Risk, risk mitigation, probability revision, value of control

*Management Science Group, Department of Management, London School of Economics and Political Science. London, UK. E-mail: s.agarwal@lse.ac.uk; g.montibeller@lse.ac.uk.

4.1 Introduction

In recent times, Operational Risk management and measurement has become an important regulatory requirement that needs to be performed by all financial organizations. Operational Risk (OpRisk) is the ‘residual’ business risk after systemic risks such as market, credit, & interest risks have been considered (Basel Committee, 2003; Cummins and Embrechts, 2006) and is defined¹ as: “the risk of direct or indirect loss resulting from inadequate or failed internal processes, people and systems or from external events” (Chaudhury, 2010; Frachot et al., 2001). Examples of OpRisks are business interruptions and losses due to natural disasters (e.g. earthquakes, hurricanes, etc.) as well mishaps caused by human factors (e.g. manual processing errors, fraud), failed ‘processes’ (e.g. employment practices and workplace safety, client practices) or ‘systems’ (e.g. hardware failures, break down of technical systems). As a part of the Basel 2 mandate (Basel Committee, 2005), businesses are required to quantify their OpRisk and set aside monetary reserves (also known as capital charge or economic capital) to be able to absorb the costs of loss events that may result from operational failures (Cummins and Embrechts, 2006; Neil et al., 2005). The mandate describes various methods (qualitative and quantitative) to calculate capital charge.² An important body of literature develops quantitative models for Operational Risk assessments (Chaudhury, 2010; Chavez-Demoulin et al., 2006; Chavez-Demoulin and Embrechts, 2004; Embrechts et al., 2004, 2003; Frachot et al., 2001) which can be used when organizations have the capabilities to quantify OpRisk using their own internal data and build OpRisk models specific to the nature of their business. One of the most sophisticated (Frachot et al., 2001) statistical/actuarial method for modelling losses due to Operational Risks described in the Basel mandate, which is becoming increasingly popular and therefore also the one employed in this paper, is the Loss Distribution Approach (LDA).

The main purpose of assessing OpRisk costs and setting aside capital reserves is to incentivize banks to improve management processes (Basel Committee, 2011) and therefore mitigate the chance or impact of a loss event. As part of business continuity planning, organizations are encouraged to have specific policies and processes to control or mitigate risks (e.g. early warning indicators, timely performance reviews) as well as devise strategies or contingency plans of actions they will undertake to mitigate the severity of a loss event if it were to occur (Basel Committee, 2003). OpRisk management is important not only to

¹ Basel Committee on Banking Supervision 2006, paragraph 644, p. 144.

² This is ‘Pillar 1’ of the 3 pillar approach which is taken over by Solvency Project Insurance (Solvency 2) set up in 2001.

prevent an organizational failure, but also because announcements of operational losses have a negative impact on the firm's market value, can strain stakeholder relationships and can affect financial ratings of the firm (Cummins and Embrechts, 2006). Furthermore, unlike other types of risk, higher OpRisk does not mean higher returns for an organization, and therefore there is a strong motivation to minimize OpRisk and thus enhance the value of a firm.

One of the key challenges that OpRisk measurement poses is how to quantitatively model various risk mitigation policies and explicitly calculate the effect of these managerial measures on capital charge assessments. Given the rarity of OpRisk events and changing management policies, these internal mitigations cannot directly be captured statistically. Ideally, a framework is required to link managerial mitigations on operational processes to the statistical methods that are used to calculate Operational Risks. Although the need for such a framework is conceptualized in Basel (as risk control self-assessment) (Lubbe and Snyman, 2010) and mentioned in the literature (Frachot et al., 2001) there has been little academic exploration on how it can be formalized quantitatively. The goal of this paper is to address this challenge of modelling risk mitigations quantitatively and better integrating them with advanced OpRisk assessment methods (such as the LDA) to explicitly calculate the effect of mitigations on capital charge assessments.

A distinction between uncertain events in terms of probability 'controllability' is instrumental for modelling mitigations and understanding *how* they affect loss distributions. The probability of some OpRisk events, such as those associated with natural disasters (e.g. hurricanes, floods) is 'uncontrollable' in the sense that the probability of the loss event cannot be altered by managerial actions, and thus these risks are usually mitigated by managing the impact of these events on business losses (e.g. by purchasing insurance). In contrast, the probability of other OpRisk events, such as those caused by people, processes or systems (e.g. staff strike), can often be controlled by managerial actions and, thus, these risks can be mitigated by affecting the probability of the loss event. While the effect of managing the consequences (e.g. purchasing insurance) is amenable to quantitative analysis and can be incorporated within the LDA (Franzetti, 2011, p. 292), explicit calculations of the effect of managerial measures to modify the probability of the loss event are challenging, as these measures are usually taken to be implicit in the model assumptions.

In this paper, we will provide a method for modelling mitigations that manage uncertainty or influence the probability of an OpRisk event *explicitly*. We will discuss how concepts from Decision Analysis, such as the Expected Value of Perfect Control (Matheson, 1990) and the generic controller of Matheson and Matheson (Matheson and Matheson,

2005), can be used to formally model the effect of such mitigations quantitatively as percentage reduction in capital charge. In particular, we propose the *probability revision matrix* (PRM) method as a way to operationalize generalizations of these Decision Analysis concepts (Agarwal et al., 2014) and show how it can offer a way to quantitatively model beliefs about mitigations against OpRisk events.

This paper also discusses an application of the PRM approach in a real world case study that was conducted for one of several Operational Risk scenarios of a leading insurance company. For this scenario, we illustrate how the PRM can serve as a powerful tool to model internal mitigations within the LDA. The problem considered relates to employment practices and workplace safety and models Operational Risks that are attributed to staff strikes. The mitigation modelled is measures to reduce the length of a strike, if it occurs. The detailed study of this mitigation exposes the practical benefits of the PRM method in terms of requiring fewer inputs, reducing the computational burden and allowing statistical risk measurements to be combined with subjective views about the mitigation.

The paper is organized as follows: In section 2, we will review how OpRisk is measured using the LDA and its limitations. In section 3, we will show how existing Decision Analysis methods can be extended to model OpRisk mitigations and propose a new approach for modelling such mitigating actions. In section 4, we use the staff strike case study to illustrate how this proposed approach can be used in the LDA to calculate capital charges with and without mitigation, as well as perform several sensitivity tests. The final section contains some reflections, concluding remarks and directions for further research.

4.2 Measuring and quantifying Operational Risk

According to Basel 2, risk reporting should provide a clear understanding of the key Operational Risks, the related drivers and the effectiveness of internal controls (Basel Committee, 2011). The mandate defines three approaches for organizations to calculate their Operational Risks – 1) Basic indicator Approach (BIA), 2) Standardized Approach (SA) 3) Advanced Measurement Approach (AMA) — which vary in the extent to which an organization can customize an OpRisk model to better reflect their business practices (Frachot et al., 2001). An advanced statistical/actuarial method of calculation under the AMA is the Loss Distribution Approach (LDA) (Chaudhury, 2010; Frachot et al., 2001). As we will describe our model for mitigation within this framework, in this section we will briefly review this approach.

4.2.1 Loss Distribution Approach

In the LDA, the loss distribution and capital charge is estimated for each event type in a given business line using internal data, over a one year horizon, and then aggregated to calculate the capital charge for the organization as a whole (Chaudhury, 2010; Frachot et al., 2001). Typically, the potential losses associated with an event type are characterized by two factors — frequency (per year) of the loss event type and impact or severity of the loss event if it occurs, usually measured directly in terms of monetary losses (Frachot et al., 2003). Both, frequency and severity are taken to be stochastic variables and typically described by a parametric probability distribution. Often the Poisson distribution is used to model frequency and one of the extreme value distributions (e.g. Lognormal, Weibull, Gumbel) is used to model severity (Moosa, 2007). The compound loss distribution for a one year time horizon is a convolution of the frequency distribution and the severity distribution, which is estimated numerically using Monte Carlo methods since it does not have an analytical expression (Frachot et al., 2001). Each simulated loss value is determined by sampling one value from the frequency distribution and k values from the severity distribution (where k is the value simulated from the frequency distribution).

Once the loss distribution has been estimated for an event type in a business line, the capital charge is typically calculated from the α -quantile of the distribution for the event type, where $\alpha > 99\%$ depends on the rating of the company (Frachot et al., 2001). The expected loss (mean of the loss distribution) and unexpected loss (difference between α -quantile and mean), of the event type can also be estimated. According to the method described in the Basel Handbook, the capital charge for the bank as a whole is computed by simply adding the capital charges across all event types and business lines. Computationally this approach is less burdensome than calculating the aggregate loss distribution for the whole bank (as a convolution of the individual loss distributions) and then estimating the charges from the aggregate distribution. Alternative methods for estimating the aggregate capital charge entails summing the squares of individual capital charges (square root rule). The theoretical assumptions that underpin these methods are discussed in Frachot et al. (2001).

4.2.2 Limitations of the Loss Distribution Approach

Using the LDA in practice poses some challenges. Firstly, often there is inadequate historical data that is required by statistical methods to assess OpRisks (Chaudhury, 2010; Chavez-Demoulin et al., 2006; Frachot et al., 2003) and therefore it is difficult to get reliable estimations of the frequency and severity distributions. While external data can be used to

supplement internal data, this may not always be suitable since operational accidents depend not only on historical observations but organizational features such as checks on fallibility, management practice, organizations' exposure to a loss event (Neil et al., 2005), which change over time. Alternatively, distributions can be estimated from subjective assessments of some statistics of the distribution (such as percentiles or moments). These can be obtained using qualitative methods, for example from scores for certain drivers of the loss event, and then can be used in combination with known statistics (Neil et al., 2005). Sensitivity analysis of capital charge to the input parameters of the distribution can be insightful and help determine the robustness of capital charge estimates.

A second limitation of the LDA is that it is based on numerical methods and the accuracy of capital charge assessment might be compromised, since the results are sensitive to the assumptions inherent to sampling methods (e.g. number of simulations). Frachot et al. (2001) discuss methods to control errors due to numerical estimation, such as comparing the theoretical moments of the distribution with the estimated ones or checking for convergence of statistics.

In this paper, we address a third limitation of the LDA: incorporating the effect of mitigation measures on the loss distribution and subsequently calculating the reduction in capital charge. Incorporating such effects may not be straightforward if the severity and frequency distributions are estimated from historical observations. In addition, if the effect of mitigation is implicit in the parameter estimates of the distributions, reassessments can be time consuming and tedious every time a mitigation is reviewed. In the next section, we will describe a procedure that can be integrated with the LDA and demonstrate how it offers a promising way to address the challenge of explicitly calculating the impact of a mitigation on the loss distribution.

4.3 Modelling mitigations when dealing with Operational Risks


In this section, we will discuss how methods from Decision Analysis to model 'control' as influences on probability distributions, can be applied to quantitatively model managerial mitigations that reduce the probability of an OpRisk event. We then suggest a new approach for modelling mitigating actions in this context.

The benefits realized from integrating Decision Analysis methods with the LDA has been recognized before. Neil et al. (2005) discuss the advantages of using graphical tools, such as Bayes nets, to formulate and analyse OpRisk problems as they offer visual convenience, can deal with the challenges of combining qualitative assessments about the effectiveness of business processes with statistical estimates, as well as facilitate 'what if'

analysis to test sensitivity of end results to model inputs. They emphasize that graphical methods help decompose the loss distribution into various causal components, which can make it cognitively easier for experts to supply probability judgements. They also show how dependence between frequency and severity can be handled by adding a common cause variable to the Bayes net (e.g. process effectiveness) and eliciting conditional probability distributions for frequency and severity. King (2001) describes how the machinery of causal models to infer causal relations from data can be used to analyse operational errors, which can help predict the effects of intervention that alter values of particular variables. These features of graphical tools are particularly useful for a risk analyst, as it allows him/her to exploit the expertise in an organization by obtaining inputs to the model from different sources, separately, and then combine them using the mathematics that underpins the graphs to get a more realistic estimate of the loss distribution.


In this paper, our choice of modelling tool for analysing OpRisk mitigations is the influence diagram (Howard, 1990; Howard and Matheson, 2005), which is a directed acyclic graph that differentiates between decision, chance and value variables and, therefore, is useful for problem structuring and clearly defining the problem variables. The components of an OpRisk event can be described in terms of decision analytic primitives — actions (mitigations), uncertain states of the world (loss events), and consequences (impact in terms of monetary losses). In an influence diagram (Figure 4-1), mitigating actions can be represented explicitly as a decision node with an arc into the target chance variable (e.g. severity of loss event). The chance variable then has different distributions associated with different decisions (or mitigations) and can also be conditional on the other chance variables (drivers of the loss event).

Figure 4-1: Generic influence diagram for OpRisk of a loss event.



In Figure 4-1, OpRisk is captured by the loss distribution (probability distribution of ‘expected losses’ node), which is determined by the probability distribution of the loss *events* and the possible monetary *consequences* associated with each outcome of the events. Notice that the loss distribution can be altered by either managing uncertainty (mitigation to control probabilities of events), or affecting the consequences (mitigation to control impact), but in *both* cases the OpRisk is altered. The statistical difference between the two types of mitigation is shown in Figure 4-2.

Figure 4-2: Statistical difference between managing uncertainty (distribution of an OpRisk event) and managing the impact of an OpRisk event (consequences).



Probabilities are assigned (col 1) to the outcome of an event (e.g. number of days the event continues for) and the loss distribution is determined by assigning a monetary loss to each outcome of the event (e.g. €1 million per day). When uncertainty is managed (col 2), the probability distribution for the length of the event is altered (grey line), which in turn affects the loss distribution (grey line). On the other hand, when the impact is managed (col 3), for example by buying an insurance which limits the loss to €2 million, the probability distribution remains the same and only the payoffs associated with the event outcomes are affected, and, hence only the loss distribution is altered (grey line).

Agarwal et al. (2014) describe how in the special case when the probability distribution of the uncertain node depends on whether or not an action is performed (i.e. ‘no action’ is a possible outcome of the decision node), actions can be regarded as interventions that alter the probability distribution of the state node. Within this perspective, OpRisk mitigations can be regarded as interventions that alter the probability distribution of a loss event which subsequently alters the capital charge assessments.

Agarwal et al. (2014) extend existing methods to model interventions in Decision Analysis and propose a general way to analyse the effect of interventions on probability distribution in a manner that is similar to analysing the effect of information on probability distribution using Bayes Rule: as a revision of a prior probability distribution. This probability revision method, when applied to OpRisk mitigations, offers a less burdensome way to compare capital charge with and without mitigation, while holding other model

inputs fixed. This provides an effective way to assess the effect of a mitigation on OpRisk costs. In the rest of the section, we will briefly review this procedure and then propose a way to operationalize it, using OpRisk mitigations as an example.

4.3.1 Modelling interventions in Decision Analysis: The generalized generic controller

A common way to model interventions in Decision Analysis is as a perfect (Matheson, 1990) or deterministic intervention (Heckerman and Shachter, 1994; Pearl, 1994) on a target uncertain state, which sets the probability of the desired event to one. The ‘value’ of exerting control is then assessed by comparing the optimal payoff before and after an intervention is performed. Matheson and Matheson (2005) extend the concept of perfect control to imperfect control. They propose the *generic controller* to model interventions that cause the most desirable state to be achieved with probability q (with $0 \leq q \leq 1$), where q reflects the *quality* of the intervention. For a target uncertainty S with n outcomes s_1, s_2, \dots, s_n and respective probabilities p_1, p_2, \dots, p_n , if s_1 is the most desirable state outcome, the revised probabilities p_i' , based on the generic controller, are a linear mixture of p_i and deterministic distribution (that guarantees s_1), given by:

$$p_i' = \begin{cases} (1 - q)p_i, & i > 1 \\ q + (1 - q)p_i, & i = 1 \end{cases} \quad (4-1)$$

Agarwal et al. (2014) suggest an alternative interpretation for the probability revision described by the generic controller in equation (4-1): as probability mass transfer from less desirable states to more desirable states. A probability revision rule based on this interpretation — *imaging* — was first proposed in the 1970s (Lewis, 1976) as an alternative to Bayes rule for evaluating the probability of conditional statements. Subject areas of causality and causal decision theory have advocated imaging as a way to calculate probabilities conditional on actions (Joyce, 1999; Pearl, 1994) since, unlike Bayes conditionalization, imaging has the advantage that the probability revision is described even when the prior probability of the conditioning event is 0 or 1.

Agarwal et al. (2014) propose the *generalized generic controller* (GGC) to describe interventions that allow arbitrary movements of probability mass between states. The corresponding probability revision rule is given by:

$$p_i' = \sum_j \rho_{ji} p_j; \quad \sum_i \rho_{ji} = 1; \quad \rho_{ji} \in [0,1] \quad (4-2)$$

where the coefficients ρ_{ij} represent the probability mass transfer from state s_j to s_i .³ They show that the GGC coincides with imaging, as both of these methods are characterized by the same linearity property (Gärdenfors, 1982) defined as follows:

Definition 4-1: Linear Probability Revision Rule: For any probability function P , let $f(P(s_i))$ be the post-intervention probabilities for a target state node $S = \{s_i\}$. A probability revision function $f: \Delta^n \rightarrow \Delta^n$ is linear if $\forall i, P(s_i) = \alpha P'(s_i) + (1 - \alpha)P''(s_i)$, for some probability functions P', P'' , $\alpha \in [0,1]$, then $f(P(S)) = \alpha f(P'(s_i)) + (1 - \alpha)f(P''(s_i))$.

Agarwal et al. (2014) expose the properties of the GGC in terms of how it interacts with Bayes rule and prove that when the target node S has a parent node R then the post-intervention probability distribution of S is the same regardless of the order in which the effect of information received about R (Bayes rule) and effect of intervention performed (GGC) is incorporated. They also discuss how this order invariance condition can be interpreted as a coherence condition: if the decision maker's pre-intervention beliefs are coherent (immune to a Dutch book) then the GGC guarantees that the post-intervention beliefs are also coherent.

There are some relevant benefits of representing OpRisk mitigations as a GGC. Firstly, the probability revision approach can be less burdensome than treating the effect of mitigations as being implicit in the parameter estimates of the frequency/severity distribution, since it avoids a complete reassessment of the model inputs when a change in mitigation policy is announced. Secondly, the linearity property of GGC imposes the constraint that the inputs of the GGC (ρ_{ji}) ought not to depend on the probability distribution of the loss event. This is crucially advantageous for our purpose as it allows an analyst to obtain inputs from different business divisions in large organizations, where frequently the personnel responsible for risk assessments are different from those managing the risks. In the next subsection, we will propose a more compact formulation of the GGC, which guarantees that its characterizing linearity property is always preserved and apply it to model OpRisk mitigations.

³ For the generic controller, if i^* is the target state, $\forall i \neq i^*$, $\rho_{ii} = 1 - q$, $\rho_{ii^*} = q$, $\rho_{i^*i^*} = 1$ and 0 otherwise.

4.3.2 The probability revision matrix

Throughout this section, we will use bold characters to represent vectors. We observe that the GGC in equation (4-2) is in fact a system of linear equations in p_1, p_2, \dots, p_n and can alternatively be formulated as a product of matrices (Meyer, 2001), given by:

$$\begin{bmatrix} \rho_{11} & \rho_{12} & \dots & \rho_{1n} \\ \vdots & \vdots & & \vdots \\ \rho_{n1} & \rho_{n2} & \dots & \rho_{nn} \end{bmatrix} \begin{bmatrix} p_1 \\ \vdots \\ p_n \end{bmatrix} = \begin{bmatrix} p'_1 \\ \vdots \\ p'_n \end{bmatrix} \quad (4-3)$$

where $\sum_i \rho_{ij} = 1$. The matrix $\boldsymbol{\rho} = [\rho_{ij}]$ represents a linear map (Lang, 1966, p. 80) of the standard Euclidean simplex $\Delta_e^{n-1} \subset \mathcal{R}^n$ into $\Delta_e^{n-1} \subset \mathcal{R}^n$ with respect to the standard basis vectors $\mathbf{e}^j = \{(e_{1j}, e_{2j}, \dots, e_{nj}) \mid e_{ij} = 1 \text{ if } j = i; e_{ij} = 0 \text{ if } j \neq i\}$ defined by $f(\mathbf{e}^j) = \boldsymbol{\rho}^j$ where $\boldsymbol{\rho}^j$ is the column vector $(\rho_{1j}, \rho_{2j}, \dots, \rho_{nj})$. We will refer to the matrix $\boldsymbol{\rho} = [\rho_{ij}]$ as the *probability revision matrix* (PRM), which is defined as:

Definition 4-2: Probability revision matrix for GGC:

A probability revision matrix for the GGC, $\boldsymbol{\rho}$, is an $n \times n$ matrix of real numbers such that: (i) $\rho_{ij} \geq 0$ (ii) $\sum_i \rho_{ij} = 1, \forall j$ (iii) for any two distributions $\mathbf{q} = (q_1, q_2, \dots, q_n)$ and $\mathbf{p} = (p_1, p_2, \dots, p_n)$, if $\boldsymbol{\rho}_p$ and $\boldsymbol{\rho}_q$ describe the probability revision matrix for \mathbf{p} and \mathbf{q} respectively and for the same GGC, then $\boldsymbol{\rho}_p = \boldsymbol{\rho}_q$.

The first two conditions guarantee that $\mathbf{p}' = (p'_1, p'_2, \dots, p'_n)$ satisfies the conditions of a probability distribution. The third condition is needed to relate the definition of linear maps to the linearity condition in definition 1, which is achieved if $\boldsymbol{\rho}_{ij}$ is not determined by the probabilities \mathbf{p} . The proof that PRM satisfies the linearity property in Definition 4-1 can be found in Appendix 1.

We will now elaborate on how the PRM operationalizes the GGC. Note that the pre-intervention probabilities can be described as:

$$\begin{bmatrix} 1 & & \\ & \ddots & \\ & & 1 \end{bmatrix} \begin{bmatrix} p_1 \\ \vdots \\ p_n \end{bmatrix} = \begin{bmatrix} p_1 \\ \vdots \\ p_n \end{bmatrix} \quad (4-4)$$

where the columns of the matrix represent the possible outcomes of the target state node. Column j of the identity matrix represents the possibility that s_j will occur; row i in column j represents the proportion of times s_i will occur instead of s_j . By definition, the matrix that represents no mitigation (identity matrix) is one where for $i \neq j$, s_i will not occur instead of s_j . An intervention can be regarded as causing a desired state s_{i^*} in a world where $s_j, j \neq i$, would have occurred (without the mitigation). ρ_{i^*j} in equation (4-3) describes the proportion of times s_{i^*} can be made to occur instead of s_j or equivalently the probability mass that gets transferred from s_j to s_{i^*} .

4.3.3 The probability revision matrix for OpRisk mitigations

OpRisk mitigations are often directed towards preventing an undesirable state from occurring and therefore correspond to a leftward shift of probability mass. If for a given target chance variable $S = \{s_i\}$, the outcome s_i is preferred to s_{i+1} for every i , the corresponding PRM is an upper triangular matrix. For $S = \{s_1, s_2, s_3\}$, where s_1 is the most desirable state and s_3 the least desirable state, the general PRM for an OpRisk mitigation is given by:

$$\begin{array}{c} s_1 \quad s_2 \quad s_3 \\ \begin{array}{l} s_1 \\ s_2 \\ s_3 \end{array} \begin{bmatrix} 1 & q_{12} & q_{13} \\ 0 & 1 - q_{12} & q_{23} \\ 0 & 0 & 1 - q_{23} - q_{13} \end{bmatrix} \end{array} \quad \text{where } \sum_{i=1}^{j-1} q_{ij} = 1 \quad (4-5)$$

where q_{ij} represents the proportion of probability mass that is shifted from the state s_j to s_i . Eliciting the coefficients for the PRM in the general form can be tedious. It may therefore be desirable to describe the mitigation by a single parameter, as in the generic controller. For a mitigation described by the generic controller, the single parameter q can be interpreted as the probability of the mitigation preventing all undesirable states from occurring (equation (4-6)). Alternatively, a mitigation that reduces severity can be described by transferring a fixed proportion (single parameter q) of probability mass from a state outcome to its 'nearest' less severe outcome (equation (4-6)).

$$\begin{array}{c} s_1 \quad s_2 \quad s_3 \\ \begin{array}{l} s_1 \\ s_2 \\ s_3 \end{array} \begin{bmatrix} 1 & q & q \\ 0 & 1 - q & 0 \\ 0 & 0 & 1 - q \end{bmatrix} \end{array} \quad \begin{array}{c} s_1 \quad s_2 \quad s_3 \\ \begin{array}{l} s_1 \\ s_2 \\ s_3 \end{array} \begin{bmatrix} 1 & q & 0 \\ 0 & 1 - q & q \\ 0 & 0 & 1 - q \end{bmatrix} \end{array} \quad (4-6)$$

(i) The generic controller


(ii) Probability mass transferred to 'nearest' state

Depending on the nature of the state variable that is mitigated, other single parameter PRMs can be formulated.

Agarwal et al. (2014) offer a geometric interpretation of the GGC as transformation of the probability simplex. The algebraic correspondence of the PRM to simplex transformations is discussed in Appendix 2. Here we will graphically illustrate the simplex transformations for the PRMs in equation (4-6). Equation (4-6) is a special simplex contained in Δ_{n-1}^e , where for $q > 1$ the number of vertices are the same as the original simplex, but the length of edges is modified to $1 - q < 1$. On the other hand, equation (4-6) encodes a more general transformation of the original simplex, where the number of vertices in the new simplex can be fewer than those in the original simplex and the length of edges can vary. The corresponding graphical representation for different values of q is shown in Figure 4-3.

Figure 4-3: Each of the 4 panels shows the original simplex (solid line), the simplex corresponding to the generic controller (4-6i) (dotted dash line) and the simplex corresponding to equation (4-6ii) (dashed line) for various values of q . The vertices represent the atomic state probabilities $P(s_i) = 1$.

$q = 1$ (The generic controller is the point s_1 and equation (4-6ii) is the edge between s_1 and s_2)



In the next section, we will describe how a single parameter PRM was formulated for modelling a mitigation on a loss event (length of a staff strike) and show how the PRM approach can be integrated with the Loss Distribution Approach to calculate the resulting impact on capital charge calculations.

4.4 Case Study: Modelling mitigations for the ‘Staff Strike’ scenario

An important criterion for any approach that is proposed for OpRisk assessments is the convenience with which it can be implemented. Despite the emerging literature on sophisticated methods to calculate OpRisk, institutions are often faced with practical challenges when attempting to apply them in practice (Chaudhury, 2010). We tested the feasibility of the PRM method to estimate the effect of mitigation on capital charge within the LDA in a case study for one loss event, *staff strike* (classified under the ‘employment relations, diversity and discrimination risk category and belonging to the Human Resources business unit), that was conducted for a leading insurance company. The aim of the study was to test the intuitive appeal of the PRM as a model for mitigations and to engage in


analyses that can offer potentially useful insights on how mitigations encoded as PRM affect capital charge assessments.

We will first give an overview of the strike problem and discuss how we adopted a Decision Analysis approach to identify the various drivers of the loss distribution for this event. Then we will describe how the loss distribution was estimated, and subsequently illustrate how we explored the effect of risk mitigation on the length of strike using the PRM method. (Although not computationally most efficient, it was convenient to build the loss distribution model in Excel so that it could be shared easily over email with the company for whom the case study was considered).

4.4.1 Problem description

A high level description of the strike event based on interviews with the Risk Assessment (RA) team and Human Resources (HR) Team is as follows: typically employee strikes may be triggered when employment official bodies disagree with management decisions (e.g. annual appraisal of salary changes). The frequency (per year) with which a strike occurs depends on the probability of the trigger event and the relationship or ‘negotiability’ of the management with the unions and workers. The severity of this event is the monetary loss incurred due to the business disruption and is a function of the number of days a strike lasts. We used an influence diagram (Figure 4-4) to structure the problem and obtained inputs for the model from both the RA and HR teams.

Figure 4-4: Influence diagram for the OpRisk of staff strike (with mitigation on length of strike).



Notice from the influence diagram that there is a difference between our implementation of the LDA to the approach described in section 4.2: rather than modelling the severity of a strike directly in terms of monetary losses we model it as one of the key drivers of losses — length of strike, measured in days — given a strike occurs. The losses

are calculated separately for each possible length of strike based on various contributing factors, such as cost due to delays in processing, additional workforce/overtime rates, loss of new business, etc. (For some of these factors the cost was a linear function of length of strike whereas for others it was a convex function).

There are a few advantages of not measuring severity directly in terms of the monetary losses, but instead a more fundamental variable on which the costs depend. Firstly, it adds transparency to the model and the assumptions that go into estimating the OpRisk costs and loss distribution. Secondly, the cost structure in a firm is often constantly reviewed and subject to frequent changes. When the parameters of the severity distribution do not depend directly on the monetary losses, the extent of effort required to update OpRisk models to reflect any changes in the cost scheme is considerably lower. Thirdly, it is possible to model the effect of specific managerial mitigations targeted towards one of the drivers of the loss distribution more reliably, since often those planning the specific mitigation may not be properly acquainted with cost implications. In a questionnaire filled out by the HR managers, it was found that the firm has contingency and mitigation plans in place to reduce the length of strike and they believe that if these mitigations are executed they alter the probabilities associated with the length of the strike.


4.4.2 Estimating the loss distribution

We assumed that both frequency x_f (number of strikes per year) and severity x_l (length of strike, days) follow a parametric distribution and the loss distribution was sampled from the assumed distributions using Monte-Carlo simulation. We assumed no dependence between the parameters of the frequency and severity distributions. A possible dependence (modelled as an arc from the 'negotiability' node to the 'length of strike' node (Neil et al., 2005)) is explored later on, in section 4.4.4.4.

The frequency distribution was assumed to be a Poisson distribution, whereas for severity two different choices of distribution — Lognormal and Weibull — were considered. Due to lack of historical data, the parameters of the distribution were obtained from the company experts' judgments. For strike frequency, negotiability (obtained using a scorecard) and probability of trigger event were considered to be the key drivers and the mean for the Poisson distribution was estimated from these scores. The parameters of the severity distribution were estimated from the company experts' opinions of what constitutes a 'typical', 'serious' and 'extreme' case of the length of strike, which correspond to the mode m_l , α_1 percentile and α_2 percentile of the distribution for severity, respectively. The distribution was then determined using the weighted least-squared approach discussed in

Appendix 3. The non-parametric loss distribution was estimated by sampling 10,000 observations from the frequency distribution and severity distribution. The capital charge was taken to be the 99.5 percentile of the distribution (total computation time was approximately 85 seconds in Excel). Figure 4-5 shows the distributions for frequency, severity and the loss distribution. The size of the simulation was constrained by the capabilities of building the model in Excel. We performed convergence tests to measure the stability of the capital charge estimates (as will be discussed in section 4.4.4.2).

Figure 4-5: Distributions for frequency ($x_f \sim Pois(\lambda = 0.02)$) and length of strike ($x_l \sim LogN(\mu = 1.06, \sigma = 0.74)$; $x_l \sim W(\alpha = 0.94, \beta = 1.49)$).



4.4.3 Modelling risk mitigation

Having estimated the loss distribution, we explored the effect of a risk mitigation on the loss distribution. We considered a mitigation on the length of strike (e.g. by giving in to the workers' demands, taking prompt actions to initiate a dialogue, etc.) and modelled it as a single parameter PRM, discussed in section 4.3.3, where the parameter q is described as follows: once a strike occurs, the firm has contingency plans in place which has $q\%$ efficiency or that there is $q\%$ less chance that a strike will continue to the next day. While hypothetical values for q are assumed in this paper, following are some techniques that can be used for obtaining the parameter q :


- i) from an efficiency score card that may already be in place to assess the efficiency of mitigations;
- ii) elicit by asking the following question: "Suppose a strike has lasted n days. What is the chance that the mitigation will prevent the strike from continuing to the next day?"

The variable q can be interpreted as the proportion of probability mass transferred from days that the strike will continue beyond day n to day n . Let $n = 0$ indicate that a strike will not occur and let us also assume that if a strike occurs it will last at least $n = 1$ day. Let p_1, p_2, \dots, p_n denote the prior probabilities that the strike will last 1 day, 2 days, etc. The corresponding PRM (given a strike occurs) and post mitigation distribution is given by:

$$\begin{matrix} \text{Strike days} \\ \mathbf{1} \\ \mathbf{2} \\ \mathbf{3} \\ \mathbf{4} \\ \mathbf{5} \\ \vdots \end{matrix} \begin{matrix} \mathbf{1} & \mathbf{2} & \mathbf{3} & \mathbf{4} & \mathbf{5} & \dots \\ \mathbf{1} & \mathbf{q} & \mathbf{q} & \mathbf{q} & \mathbf{q} & \mathbf{q} \\ \mathbf{0} & \mathbf{1-q} & \mathbf{q(1-q)} & \mathbf{q(1-q)} & \mathbf{q(1-q)} & \mathbf{q(1-q)} \\ \mathbf{0} & \mathbf{0} & \mathbf{(1-q)^2} & \mathbf{q(1-q)^2} & \mathbf{q(1-q)^2} & \mathbf{q(1-q)^2} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{(1-q)^3} & \mathbf{q(1-q)^3} & \mathbf{q(1-q)^3} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{(1-q)^4} & \mathbf{q(1-q)^4} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{(1-q)^5} \end{matrix} \begin{matrix} \left[\begin{matrix} \mathbf{p_1} \\ \mathbf{p_2} \\ \mathbf{p_3} \\ \mathbf{p_4} \\ \mathbf{p_5} \\ \vdots \end{matrix} \right] \end{matrix} = \begin{matrix} \left[\begin{matrix} \mathbf{p_1 + qG(1)} \\ \mathbf{(1-q)(p_2 + qG(2))} \\ \mathbf{(1-q)^2(p_3 + qG(3))} \\ \mathbf{(1-q)^3(p_4 + qG(4))} \\ \mathbf{(1-q)^4(p_5 + qG(5))} \\ \vdots \end{matrix} \right] \end{matrix} \quad (4-7)$$


where $G(n)$ is the decumulative distribution function of x_l , i.e. $P(x_l) > n$. The transformation of the probability simplex corresponding to this mitigation is shown in Figure 4-6.

Figure 4-6: The transformation of the simplex corresponding to various values of q in equation (4-7) (modified for at most three strike days).



Once the post-mitigation distribution for the length of strike was calculated, the loss distributed was re-estimated (sample: 10,000 observations) and the effect of the mitigation on the capital charge was assessed (Figure 4-7).

Figure 4-7: Severity (Lognormal) and Loss distribution without mitigation and with mitigation ($q = 0.5$ in equation (4-7)).



In general, the value of mitigation can be assessed by comparing various statistics of the pre-mitigation and the post-mitigation distribution. For example, the difference in means of the distributions before and after ‘control’ corresponds to the notion of Expected Value of

Control. An industry measure for valuing mitigations, also the one presented here, is the percentage reduction in capital charge due to mitigation.

4.4.4 Sensitivity analysis

In this section, we will discuss a few exploratory tests that were performed to test the sensitivity of the capital charge estimates to various model assumptions and input parameters.

4.4.4.1 Effect of varying q

We explored the effect of a more general mitigation on the capital charge estimate where we allowed the quality of mitigation to increase or decrease over time. This model requires *two* parameter inputs: the initial quality, q_1 , of terminating a strike on the first day and the change in quality, δ . We can determine the quality ($0 \leq q_n \leq 1$) for the remaining days using the recursion:

$$q_n = \min(1, q_{n-1} + \delta) + \max(0, q_{n-1} + \delta) - (q_{n-1} + \delta); \quad -q_1 \leq \delta \leq 1 - q_1 \quad (4-8)$$


The corresponding PRM is given by:

Strike days	1	2	3	4	...	prior	post-mitigation
1	q_1	q_1	q_1	q_1	q_1	$\begin{bmatrix} p_1 \\ p_2 \\ p_3 \\ p_4 \\ \vdots \end{bmatrix}$	$= \begin{bmatrix} p_1 + q_1 G(1) \\ (1 - q_1)(p_2 + q_2 G(2)) \\ \prod_1^2 (1 - q_i) (p_3 + q_3 G(3)) \\ \prod_1^3 (1 - q_i) (p_4 + q_4 G(3)) \\ \vdots \end{bmatrix}$
2	$1 - q_1$	$q_2(1 - q_1)$	$q_2(1 - q_1)$	$q_2(1 - q_1)$	$q_2(1 - q_1)$		
3	0	$\prod_1^2 (1 - q_i)$	$q_3 \prod_1^2 (1 - q_i)$	$q_3 \prod_1^2 (1 - q_i)$	$q_3 \prod_1^2 (1 - q_i)$		
4	0	0	$\prod_1^3 (1 - q_i)$	$q_4 \prod_1^3 (1 - q_i)$	$q_4 \prod_1^3 (1 - q_i)$		
\vdots	0	0	$\prod_1^3 (1 - q_i)$	$q_4 \prod_1^3 (1 - q_i)$	$q_4 \prod_1^3 (1 - q_i)$		
0	0	0	0	0	\ddots		

(4-9)

Figure 4-8 show the percentage reduction in capital charge for various values of q_1 and δ in equation (4-8).

Figure 4-8: Percentage reductions in capital charge (99.5 percentile) for various qualities of intervention in equation (4-9). The solid line shows the percentage reduction if the quality is fixed over time (equation (4-8)). The grey bars indicate the extent to which percentage reduction can change depending on how quality changes over time.




The sensitivity of percentage reduction in capital charge to the quality of intervention (Figure 4-8) shows that the marginal value of mitigation decreases as quality increases. Furthermore, as the quality of intervention increases, changes in quality for each additional day of strike has a lower effect on the resulting capital charge estimates. When reliable estimates for quality of intervention cannot be obtained, sensitivity analysis provides useful insights on the desired levels of quality of intervention and given a quality, whether it is worth taking the effort to improve the quality of intervention as the event continues.

4.4.4.2 Effect of chosen distribution for severity and convergence tests

As mentioned in section 4.2.2, when using numerical methods to estimate the loss distribution, it is important to test for the stability of the sample distribution. Figure 4-9 shows the cumulative density function for the 99.5 percentile of the sample distribution in Figure 4-7 (based on 10,000 observations) for 5,000 replications.

Figure 4-9: Convergence tests for capital charge estimates and percentage reduction: The black line corresponds to a Lognormal distribution and grey line corresponds to a Weibull distribution for severity ($p < 0.01$ for all figures, based on two sample Kolmogorov-Smirnov test). For Lognormal distribution the mean capital charge was found to be €3.13 million (\pm €3750) without mitigation, and 1.13 million (\pm €1150) with mitigation. The mean percentage reduction was found to be 64%. (\pm .03%).




The capital charge estimates were found to be more robust when a Lognormal distribution is assumed for severity (significantly lower variance than Weibull) whereas the measured percentage reduction in capital charge is higher when a Weibull distribution is assumed. One reason for these differences is that the Weibull distribution has heavier tails.

4.4.4.3 Effect of model inputs

We also performed extensive tests to check the sensitivity of the capital charge calculations to various model inputs and assumptions. The left panel in Figure 4-10 shows the sensitivity of the capital charge to the frequency (mean of Poisson distribution) with a fixed severity distribution. The right panel in Figure 4-10 shows the sensitivity of the capital charge

estimates to the scores from the score cards for negotiability — social dialogue adequacy and exposure to conflict.

Figure 4-10: Sensitivity of capital charge estimates (without mitigation) to various inputs of the score.




A regression was conducted to estimate the slope of the plane and it was found that the sensitivity of capital charge estimates to ‘social dialogue adequacy’ score is around 1.5 (2) times the sensitivity to ‘exposure to conflict’ score when a Lognormal (Weibull) distribution is assumed for the length of strike.

4.4.4.4 Effect of assuming correlation between the frequency and severity distribution

We also looked at the effect of assuming dependence of the distribution for the length of strike on negotiability. This dependence introduces a correlation between the frequency and severity distribution (Neil et al., 2005). We found that at a fixed level of probability of trigger event, introducing this correlation increases the spread of capital charges (standard deviation of capital charge for various for scores of the score cards) by 32% (9%) for Lognormal (Weibull) distribution. From Figure 4-11 we can see that for lower values of negotiability score, the capital charge is higher when the parameters of the length of strike are conditional on the ‘negotiability’ score when using a Lognormal distribution for the length of the strike.

Figure 4-11: Capital charge estimates using parameters conditional on ‘negotiability’ score vs. unconditional parameters. (Based on rank sum test, $p \leq 0.05$ for Lognormal distribution, when negotiability < 0.25).



Although not presented here, the sensitivity tests to model inputs (in sections 4.4.4.3 and 4.4.4.4) can be easily extended to include the mitigations within our proposed framework.

4.5 Summary and conclusion

Managers frequently contemplate risk management as being able to affect the odds of chance variables and, thus, the notion of ‘control’ can be potentially useful for analysing such decisions where uncertainty is regarded as being ‘controllable’. OpRisk management is an area where the probabilities of many OpRisk events are intrinsically linked to managerial practices in an organization. In fact, one of the aims of formalizing the OpRisk assessments is that, over time, it will converge towards effective risk management practices (Basel Committee, 2011), thus leading to fewer business disruptions. Explicitly calculating the reduction in capital charge due to mitigations can provide natural incentives for enterprises to take active measures to mitigate risks.

In this article we proposed a quantitative method to assess the value of Operational Risk (OpRisk) mitigations that alter the probability distribution of the loss event as percentage reduction of capital charge within a Loss Distribution Approach (LDA) framework. Since the LDA can be expressed in terms of Decision Analysis primitives, Decision Analysis methods can be extended to support the input estimations needed in the LDA. The concept of a controller and Expected Value of Control has been a part of the Decision Analysis repertoire but, unlike Expected Value of Information, its application has not been discussed extensively. One reason for this is that, in Decision Analysis, often the probabilities of uncertain variables are formulated as being outside the decision maker’s influence. Thus, there has been little research on developing a proper framework to encode views about interventions that are not deterministic, and a systematic method to define post intervention distributions as a probability revision of the prior. In this article, we proposed the probability revision matrix (PRM) method as one way to quantitatively model non-deterministic interventions and thus broaden the applicability of controllers. This approach operationalizes the concept of controllers in Decision Analysis and, as we showed, can be integrated with the Loss Distribution Approach as a separate input. OpRisk mitigations modelled as a PRM provide a more consistent and tractable comparison of various mitigation policies on a given probability distribution (modelled using LDA) and their effects on capital charge.

In the case study that was conducted and presented here, we were able to test the advantages of integrating Decision Analysis methods and the PRM with traditional OpRisk assessment methods, such as the LDA. We found graphical tools useful for engaging experts

in a discussion on their beliefs about various assessments and dependencies between variables. When there was lack of consensus whether a dependency between variables should be modelled or discretionary choice of distribution for modelling length of strike, it was possible to evaluate the extent to which the assumptions affected the end estimates by altering the assumptions about one variable and repeating the calculations, while keeping all other assumptions intact. Our sensitivity tests revealed that some of these decisions can lead to significantly different estimates for the capital charge. We were also able to demonstrate the relationship between quantitative assessment of mitigations and capital charge reduction, which can be insightful when designing mitigation procedures. Furthermore, the tractability of the proposed method enables comparisons of different mitigations, permits calculations to measure the combined effect of multiple mitigations, and allows the analyst to investigate mitigations that have the required effect on the probabilities of a loss event.

Experts from the OpRisk team endorsed the usefulness of sensitivity analysis and robustness tests in the Advanced Measurement Approach, as they demonstrate rigour and reliability of the estimates, which makes it easier to get approval from the regulators. When mitigations are a part of the OpRisk model, the PRM method can make these calculations less tedious. The ability to analytically integrate the expertise between the Operational Risk team and Human Resources team was also a novel contribution to the OpRisk calculation procedures that were already in place. This type of analysis can be extended to other OpRisk scenarios, whenever uncertainty can be ‘controlled’ and the scenarios are modelled statistically using the LDA.

Despite the computational advantages, eliciting the inputs for the PRM so that it is a reasonable quantitative description of the mitigation may not always be straightforward. In general a clear rationale and unbiased process for obtaining the inputs is crucial if an institution wants to use the PRM method to negotiate a lower capital charge on the basis of certain mitigation plans it has in place. (A thorough procedure for eliciting the parameter ρ_{ji} for a general PRM, based on an approach that is similar to the probability wheel (Shephard and Kirkwood, 1994), is explained in Agarwal et al. (2014)). In this paper, we formulated PRMs with at most two parameters. One area for future research is to develop simpler and more intuitive ways to measure the numeric inputs of the PRM. The feasibility and usefulness of this approach can also merit from further exploration.

Acknowledgements

This research is sponsored by the AXA Research Fund. We would like to thank Christian Vie from the AXA Human Resources Team and Neil Ramsay, Dora Elamri and Fabien

Chabanon from the AXA Operational Risk team for insightful discussions on Operational Risk assessments in practice and for sharing internal data for empirical testing of the model proposed in this paper. Unless stated explicitly, the opinions and content of this paper are of its authors only and are not endorsed or represent those of AXA.

4.6 Appendices for Chapter 4

Appendix 4-1: Proof that the probability revision matrix satisfies the linearity property

Let $f(P(s_i))$ be the post-intervention probabilities for a target state node $S = \{s_i\}$ for any probability function P . Suppose $\forall i \ P(s_i) = \alpha P'(s_i) + (1 - \alpha)P''(s_i)$ for some $P'(s_i), P''(s_i)$ and $0 \leq \alpha \leq 1$. Let $[\rho_{ij}]$ represent the probability revision matrix corresponding to the probability revision function f . Then $\forall i, f(P(s_i)) = \sum_j \rho_{ij} P(s_j)$. $\forall i,$

$$\begin{aligned} f(\alpha P'(s_i) + (1 - \alpha)P''(s_i)) &= \sum_j \rho_{ij} (\alpha P'(s_j) + (1 - \alpha)P''(s_j)) \\ &= \sum_j \rho_{ij} \alpha P'(s_j) + \sum_j \rho_{ij} (1 - \alpha)P''(s_j) \\ &= \alpha \sum_j \rho_{ij} P'(s_j) + (1 - \alpha) \sum_j \rho_{ij} P''(s_j) \\ &= \alpha f(P'(s_i)) + (1 - \alpha)f(P''(s_i)) \end{aligned}$$

Appendix 4-2: Geometric interpretation of probability revision matrix

Suppose S has n outcomes s_1, s_2, \dots, s_n and the probability distribution for S is given by $\mathbf{p} = (p_1, p_2, \dots, p_n)$. Let $\mathbf{p}' = (p'_1, p'_2, \dots, p'_n)$ be the post-mitigation probabilities as described by equation (4-2). When no mitigation is performed, as in equation (4-4), the columns of the PRM are the standard basis vectors \mathbf{e}^j and geometrically correspond to the vertices of the Euclidean $n - 1$ simplex $\Delta_e^{n-1} \subset \mathcal{R}^n$, defined as (Wallace, 1957, p. 96):

Definition 4-3: Euclidean $n - 1$ simplex

Let $\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^n$ be points of \mathcal{R}^n where $\mathbf{x}^j = (x_{1j}, x_{2j}, \dots, x_{nj})$. Then the Euclidean $n - 1$ simplex Δ_x^{n-1} is the set of all points $\mathbf{z} = (z_1, z_2, \dots, z_n)$, given by:

$$\Delta_x^{n-1} = \{(z_1, z_2, \dots, z_n) \in \mathcal{R}^n \mid z_i = \sum_{j=1}^n \lambda_j x_{ij}, \lambda_j \geq 0, \sum_{j=1}^n \lambda_j = 1, \forall i = 1, 2, \dots, n\}.$$

λ_j are uniquely defined for each point \mathbf{z} (Wallace, 1957, p. 97). Setting $\mathbf{x}^j = \mathbf{e}^j$ and $\lambda_j = p_j$ we have the pre-mitigation probability distribution $\mathbf{p} = \sum_{j=1}^n p_j \mathbf{e}^j$ and is a point on Δ_e^{n-1} . In general, Δ_e^{n-1} can be considered to be the set of all possible pre-mitigation probabilities. Consider the post-mitigation probabilities $\mathbf{p}' = (p'_1, p'_2, \dots, p'_n)$ where $\mathbf{p}' = \sum_{j=1}^n p_j \boldsymbol{\rho}^j$. By

Definition 4-3, \mathbf{p}' is a point on the simplex Δ_{ρ}^{n-1} described by the vertices ρ^j and $\lambda_j = p_j$. The probabilities \mathbf{p} and \mathbf{p}' are related by the linear map $f: \Delta_e^{n-1} \rightarrow \Delta_{\rho}^{n-1}$ such that $f(\mathbf{e}^j) = \rho^j$ and $\forall \mathbf{y} = (y_1, y_2, \dots, y_n) \in \Delta_e^{n-1}, \mathbf{y} = \sum_{j=1}^n \mathbf{e}^j y_j, f(\mathbf{y}) = \sum_{j=1}^n f(\mathbf{e}^j) y_j = \sum_{j=1}^n \rho^j y_j \in \Delta_{\rho}^{n-1[30]}$. This linear map f can be described by the PRM $\rho = [\rho_{ij}]$. Thus, the PRM describes a transformation of Δ_{n-1}^e where $\Delta_{\rho}^{n-1} \subset \Delta_e^{n-1}$ is the transformed simplex and \mathbf{p}' is the projection of \mathbf{p} on a new simplex with vertices ρ^j . To ensure that \mathbf{p}' satisfies the linearity property that characterizes the GGC we need that *all* $\mathbf{y} \in \Delta_e^{n-1}$ get projected to the *same* simplex and therefore the specification of ρ_{ij} should not depend on \mathbf{y} .

Appendix 4-3: Distribution fitting functions (weighted least square method) for Lognormal and Weibull distribution

Given estimates of mode (m), events corresponding to the ‘serious’ percentile (y_1) and ‘extreme’ percentile (y_2), and precision w_i (also obtained from experts), the function to be minimized is:

$$w_1(m - m_l)^2 + w_2(y_1 - F^{-1}(\alpha_1))^2 + w_3(y_2 - F^{-1}(\alpha_2))^2$$

where m_l is the mode of the distribution, α_1, α_2 the percentiles that reflect ‘serious’ and ‘extreme’ cases, respectively, and F is the cumulative distribution of the severity function. The corresponding objective function for a Lognormal distribution and Weibull distribution are:

Lognormal (with parameters μ, σ)

$$w_1(m - e^{\mu - \sigma^2})^2 + w_2 \left(y_1 - \frac{1}{2} \left(1 + \operatorname{erf} \left(\frac{\ln(\alpha_1) - \mu}{\sigma\sqrt{2}} \right) \right) \right)^2 + w_3 \left(y_2 - \frac{1}{2} \left(1 + \operatorname{erf} \left(\frac{\ln(\alpha_2) - \mu}{\sigma\sqrt{2}} \right) \right) \right)^2$$

Weibull (with parameters α, β)

$$w_1 \left(m - \beta * \left(1 - \frac{1}{\alpha} \right)^{1/\alpha} \right)^2 + w_2 (y_1 - \beta * (-\ln(1 - \alpha_1))^{1/\alpha})^2 + w_3 (y_2 - \beta * (-\ln(1 - \alpha_2))^{1/\alpha})^2$$

4.7 References for Chapter 4

1. Agarwal, S., Montibeller, G., Morton, A., 2014. Stochastic Interventions: A New Probability Revision Method to Model the Effects of Actions on Uncertainties.
2. Basel Committee, 2003. Sound Practices for the Management and Supervision of Operational Risk.
3. Basel Committee, 2005. International Convergence of Capital Measurement and Capital Standards: A Revised Framework. Bank for International Settlements, Basel, Switzerland.

4. Basel Committee, 2011. Operational Risk – Supervisory Guidelines for the Advanced Measurement Approaches.
5. Chaudhury, M., 2010. A review of the key issues in operational risk capital modeling. *J. Oper. Risk* 5, 37–66.
6. Chavez-Demoulin, V., Embrechts, P., 2004. Advanced extremal models for operational risk. Prepr. Dep. Math. ETH-Zent.
7. Chavez-Demoulin, V., Embrechts, P., Nevšlehová, J., 2006. Quantitative models for operational risk: extremes, dependence and aggregation. *J. Bank. Finance* 30, 2635–2658.
8. Cummins, J., Embrechts, P., 2006. Introduction: Special section on operational risk. *J. Bank. Finance* 30, 2599–2604.
9. Embrechts, P., Furrer, H., Kaufmann, R., 2003. Quantifying regulatory capital for operational risk. *Deriv. Use Trading Regul.* 9, 217–233.
10. Embrechts, P., Kaufmann, R., Samorodnitsky, G., 2004. Ruin theory revisited: stochastic models for operational risk. *Risk Manag. Cent. Bank Foreign Reserv.* 243–261.
11. Frachot, A., Georges, P., Roncalli, T., 2001. Loss distribution approach for operational risk. Available SSRN 1032523.
12. Frachot, A., Georges, P., Roncalli, T., 2001. Loss distribution approach for operational risk. Available SSRN 1032523.
13. Frachot, A., Moudoulaud, O., Roncalli, T., 2003. Loss distribution approach in practice. *Basel Handb. Guide Financ. Pract.*
14. Franzetti, C., 2011. Operational risk modelling and management. Chapman and Hall/CRC.
15. Gardenfors, P., 1982. Imaging and Conditionalization. *J. Philos.* 79, 747–760.
16. Heckerman, D., Shachter, R., 1994. A decision-based view of causality, in: *Proceedings of the Tenth International Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers Inc., Seattle, WA, pp. 302–310.
17. Howard, R., 1990. From influence to relevance to knowledge., in: R.M.Oliver & J.Q. Smith, *Influence Diagrams, Belief Nets and Decision Analysis*. John Wiley & Sons Ltd.
18. Howard, R., Matheson, J., 2005. Influence diagrams. *Decis. Anal.* 2, 127–143.
19. Joyce, J., 1999. *The Foundations of Causal Decision Theory*. Cambridge University Press, Cambridge Books Online.
20. King, J., 2001. *Operational Risk: Measurement and Modelling*. John Wiley & Sons Ltd, Chichester.
21. Lang, S., 1966. The matrix associated with a linear map, in: *Linear Algebra*. Addison-Wesley Publishing Co., USA.


-
22. Lewis, D., 1976. Probabilities of Conditionals and Conditional Probabilities. *Philos. Rev.* 85, 297–315.
 23. Lubbe, J., Snyman, F., 2010. The advanced measurement approach for banks. *IPM17 Implement.* 1993 *Syst. Natl. Acc.* 141.
 24. Matheson, D., Matheson, J., 2005. Describing and Valuing Interventions That Observe or Control Decision Situations. *Decis. Anal.* 2, 165–81.
 25. Matheson, J., 1990. Using influence diagrams to value information and control, in: R.M.Oliver & J.Q. Smith, *Influence Diagrams, Belief Nets and Decision Analysis.* John Wiley & Sons Ltd, pp. 25–48.
 26. Meyer, C., 2001. *Matrix Analysis and Applied Linear Algebra.* Society for Industrial and Applied Mathematics.
 27. Moosa, I., 2007. *Operational risk management.* Palgrave Macmillan.
 28. Neil, M., Fenton, N., Tailor, M., 2005. Using Bayesian Networks to Model Expected and Unexpected Operational Losses. *Risk Anal.* 25, 963–972.
 29. Pearl, J., 1994. A Probabilistic Calculus of Actions, in: *Proceedings of the Tenth Annual Conference on Uncertainty in Artificial Intelligence* (eds. R. Lopez de Mantaras and D. Poole). pp. 454–462.
 30. Shephard, G., Kirkwood, C., 1994. Managing the judgmental probability elicitation process: a case study of analyst/manager interaction. *Eng. Manag. IEEE Trans.* On 41, 414–425.
 31. Wallace, A., 1957. The Homology groups, in: *An Introduction to Algebraic Topology,* International Series of Monographs in Pure and Applied Mathematics. Pergamon Press.

4.8 Supplementary Material for Chapter 4

S 4-1: Screenshots from the Excel tool showing the model inputs and outputs for data

	A	B	C	D	E	G	H	I	J	K	
1	Model INPUTS					Nature of Variable	Model OUTPUTS				
2	1) Likelihood of restructuring	10%			entered by user (given by AXA): <i>Enter value</i>						
3										Recalculate	
4										<i>Calculation Completed.</i>	
5	SCORES					Value at Risk	99.50%	i			
6	2.1) Negotiability 1	0.4			entered by user (given by AXA): <i>Enter values between 0.2 and 0.9 with 1 decimal place</i>	9.1) Capital charge without mitigation	€	3,296,940	€	3,296,940	
7	2.2) Negotiability 2	0.5				9.2) Capital charge with mitigation	€	1,098,191	€	1,000,000	
8	3) Cost function	Non-Linear, Cost Table in GS European s		2012	entered by user (given by AXA):	Savings	€	2,198,749	€	2,296,940	
9											
10	4) Frequency distribution	2.00%	Mean of Poisson distribution		Auto calculated from (1) and (2) based on AXA's methodology	Simulation indicators from 10000 scenarios					
11						Max frequency		2 per year			
12	5.1) Severity (days of strike) distribution known values	Mode (Typical) 2	80%tile (Serious) 5	95%tile (Extreme) 10	Re-fit distribution entered by user (2012 estimates given by AXA): <i>User can change</i>	Max length of strike		18 days			
13	5.2) Select Distribution	Weibull			selected by user: <i>Select from list box</i>	Max Cost		11,903,616 €			
14	5.3) Dependence on Negotiability 2?	Yes			selected by user: <i>Select from list box</i>						
15											
16	5.4) Parameters of selected distribution	mean 1.06	std. dev. 0.74		Calculated in the model based on	Implied by the fitted distributions	Mode Typical 2	80% Serious 5	95% Extreme 10		
17											
18	6) Mitigation	Initial probability of Mitigation 0.5	Step 0	Maximum Mitigation Probability 1	Probability revision matrix						
19											

S 4-2: Algorithm for simulating the loss distribution



S 4-3: Sample Questionnaire used to understand AXAs model and risk mitigation policies (not used in the paper)

Important Note:

The question corresponding to the impact of mitigation is worded differently in this questionnaire (Part 2 Q5 and Q6). We phrased a less confusing wording for numerical elicitations of probability revision matrix when writing the paper (described in this paper in section 4.4.3).

Responses to the questionnaire are not included for confidentiality reasons.

Short Questionnaire for AXA on Operational Risks due to Strikes

Shweta Agarwal (s.agarwal@lse.ac.uk)

Supervisor: Gilberto Montibeller (g.montibeller@lse.ac.uk)

Department of Management, London School of Economics, London, WC2A 2AE

Any information provided by AXA in this questionnaire will be kept strictly confidential. This information will be used for research only and will not be reproduced without explicit consent from AXA.

5) In your experience at AXA, what is the longest number of days that a strike has lasted?

0 days	1 day	2 days	3 days	4 days	5 days	6 days	7 days	8 days	9 days or more

6) We welcome any comments, thoughts or related information you would like to share with us on estimating the duration of strike. Please feel free to add them below.

PART 2

In this part of the questionnaire, we would like to ask AXA a few questions on their approach to mitigate the risks due to strikes

1) Please indicate whether you ‘agree’ or ‘disagree’ with the following statements.

	Agree	Disagree
i. AXA has a plan of mitigation options before a strike occurs.		
ii. AXA takes initiatives to mitigate how long a strike lasts, once a strike occurs.		
iii. If a strike has occurred in the year, the estimated chance of a strike occurring (again) will not change.		
iv. AXA takes initiatives to reduce the chance of a strike occurring even when there is no threat of a strike.		
v. If a strike has occurred in the year, it is more likely that a strike will occur again.		
vi. AXA does not take initiatives to mitigate how long a strike lasts, once a strike occurs.		
vii. AXA takes initiatives to reduce the chance of a strike occurring only when there is threat of a strike.		
viii. If a strike has occurred in the year, it is less likely that a strike will occur again.		

7) It would be helpful for us to know the type of mitigations that are used by managers at AXA. Please add any information you can share on initiatives taken by AXA

i. To mitigate the risk of a strike occurring

ii. To reduce the number of days a strike lasts for, if it occurs

iii. Any other information

Chapter 5

Concluding Remarks

The motivating idea of this research was that, when contemplating risky decisions, individuals not only seek to optimize expected payoffs but also endeavour to actively mitigate the risks while retaining the potential rewards (e.g. Huber, 2002, 2007). Risks can be mitigated by influencing either the uncertainties (e.g. advertising to improve the chance of a higher demand) and/or the impact (e.g. purchasing insurance). While the notion of managing impact (dependence between actions and consequences) is embedded in most formalizations of decisions (e.g. Savage, 1954), the notion of managing uncertainty (dependence between actions and probabilities), conceptualized in the literature as ‘control’ (Goodie, 2003; e.g. Matheson, 1990) or as ‘interventions’ (e.g. Matheson and Matheson, 2005; Pearl, 2000) has been comparatively less explored. The empirical research on decision making under risk has also predominantly been in the context where actions cannot directly influence probabilities attached to uncertainties. In contrast, corporate perspectives on risk management and risk taking are closely linked to the notion of managing the uncertainty inherent to the description of risks, tantamount to modifying the probability distribution associated with uncertain events, when possible.

Thus, as argued in this thesis, in order to make the formal study of decision making under uncertainty better resonate with managerial thinking, the evaluation of risky decisions should focus not only on the effect of actions on consequences but, when relevant, also build upon the contingent influence of actions on uncertainty. In relation, this research focused on deepening the understanding of the behavioural effect of uncertainty control on risk taking and also attempted to expand the capability of Decision Analysis to quantitatively model the effect of interventions on probability distributions.

Following previous suggestions (see Brandstätter and Schwarzenberger, 2001; Goodie, 2003), first a distinction of uncertainty based on its ‘controllability’, i.e. whether or not a decision maker can modify the probabilities attached to uncertain events, was proposed. This dual construct of uncertainty mirrors the distinctions between uncertainties based on the nature of probabilities (e.g. Morgan and Henrion, 1990; Von Mises, 1949, pp. 107–115) or information (e.g. Fox and Ulkumen, 2011; Chatterton, 2001; Hoffman and Hammonds, 1994; Chernoff and Moses, 1959) (discussed briefly in section 1.1.1). It was noted that while there is ample evidence that subjective characteristics of uncertainty, such as ambiguity, can play an essential role in understanding why some options are preferred

over others and the willingness to accept risk (Ellsberg, 1961; Ulkumen et al., 2014), the prevailing understanding of how ‘controllability’ of uncertainty affects risk taking is somewhat limited (see Huber, 2007). This thesis revisited some of the prevailing conceptualizations of control (Brandstätter and Schwarzenberger, 2001; Goodie, 2003; Langer, 1975; Li, 2011; Matheson and Matheson, 2005; e.g. Matheson, 1990; Young et al., 2011) and identified the unaddressed behavioural, procedural and conceptual issues (section 1.2).

From a behavioural perspective, it was argued that the commonly observed finding that control, illusory or perceived, increases willingness to accept risk was potentially incomplete, owing to the design of the studies that have been conducted. Specifically, the gambles employed in the previous studies have not explored the full range of probabilities and payoffs, and were skewed towards cases where typically risk aversion is observed in the absence of control. Thus, the effect of control in cases when otherwise risk taking tendencies prevail (e.g. gambles with predominantly negative valued payoffs) has not been captured prior to this study. In this thesis, it was suggested that for a more comprehensive understanding of the effect of control on risk taking, its effect on the fourfold pattern of risk attitudes, within the context Cumulative Prospect Theory (Tversky and Kahneman, 1992), needs to be investigated. It was also proposed that, from a methodological perspective, a study which endows subjects with actual control, rather than leaving it to their perception, would constitute a novel test of control on risk taking. Both these research gaps were addressed by the study presented in this thesis (Chapter 2), with the intent of enhancing the behavioural insights on how control affects decision making and better tying laboratory findings to managerial risk taking tendencies.

From an analytic perspective, too, it emerged that work needs to be done on developing the tools of Decision Analysis for modelling and analysing control. The thesis discussed that while the notion of control is easily represented in influence diagrams as interventions (using an arc from a decision node to state node) (Howard and Matheson, 2005, 1984), the analysis of interventions has mostly been from the perspective that they bring about a desired state with certainty (Matheson, 1990; Pearl, 2000). Although Matheson and Matheson (2005) have proposed a procedure for modelling imperfect control, as discussed in section 3.3.2, the scope of this procedure is limited by the type of interventions it can model. Furthermore, unlike Bayes rule — the probability revision rule used to model the effect of information — which is supported by theoretical foundations such as the coherence criterion, the existing Decision Analysis procedures for modelling the impact of interventions on probabilities are procedural and lack a similar theoretical basis. Thus, the

thesis argued that existing methods to model interventions as a probability revision process in Decision Analysis are underdeveloped. Additionally, it was suggested that there is also scope to explore the relationship between the concept of control in Decision Analysis and a belief revision rule, called imaging, in Philosophy. In this thesis, these research gaps were addressed and developments were proposed to enhance the capability of Decision Analysis to formally deal with the concept of control (Chapter 3). The feasibility of applying the proposed procedure in practice was also explored (Chapter 4). The intent of pursuing these directions was that it can lead to the development of Decision Analytic methods that better resonate with managerial rationale for making strategic risky choices and thus expand the applicability of Decision Analysis to corporate risk management. The next subsection summarizes the research aims of this thesis and key results.

5.1. A summary of the areas of research covered in this thesis

The study of control from behavioural, theoretical and prescriptive perspectives of decision making was distributed correspondingly over three papers, which constitute this thesis. In the context of behavioural decision making, the primary aim of the related paper (Chapter 2) was to conduct a more comprehensive and novel laboratory study to test the effect of probability control on risk taking. A secondary aim was to explore any systematic relation between exerting control and prior probabilities of ‘success’. In the context of theoretical decision making, the primary goal of the related paper (Chapter 3) was twofold — 1) to extend the existing methods of modelling interventions in Decision Analysis so that more general interventions can be modelled; and 2) to ground the existing and proposed procedures for modelling the effect of interventions on probabilities in theoretical foundations, similar to the coherence criterion that supports Bayes rule as a procedure for modelling the effect of information on probabilities. A secondary aim was to explore potential relationships between Decision Analysis methods to model control and imaging. Finally, in the context of prescriptive decision making, the goal of the related paper (Chapter 4) was to test the feasibility of applying the developed procedure in practice, especially in areas of risk management. Each of the aims in the respective areas was addressed as follows:

5.1.1. Understanding the behavioral effects of control

A novel within subject randomized controlled experiment was designed to test the behavioural effect of control on risk taking. The study differed from previous experimental work in three crucial ways, to enable a test of the effect of control on the fourfold pattern of risk attitudes, leading to a broader understanding of the effect of control on risk taking.

Firstly, controlling probabilities was manipulated as a fixed change in probability mass from the worse outcome to the better outcome and therefore was an objective contingency of the task and not subject to judgement. Secondly, the tasks were designed so that they cover both negative and positive valued payoffs. The majority of previous research has only looked at positive valued payoffs. Finally, in this research the entire spectrum of the probability scale was explored, and this was possible because controllable probabilities were not associated with confidence in general knowledge answers.

Findings from the experiments suggest control has a *moderating effect* on risk attitudes. Although this finding does not conflict with extant insights (obtained mostly in the context of gains-only or mixed gambles for moderate probabilities of gains), they differ from their conclusions that control *increases* risk taking (Brandstätter and Schwarzenberger, 2001; Goodie and Young, 2007; Langer, 1975). Nonetheless, it was discussed in section 2.3.3 that the findings of the present study can be explained using arguments similar to the ones used in previous studies, which are based on optimism and positive affect that control induces (see Young et al., 2011) and the known effects of positive affect on risk taking (Isen and Patrick, 1983; Seo et al., 2010). Alternative cognitive explanations based on risk bearing, which help tie the observed effects of control to managerial risk taking, were also offered (Wiseman and Gomez-Mejia, 1998).

The data analysis on understanding how exerting control is valued revealed that ‘certainty’ plays a crucial role in ascertaining the perceived value of control. Specifically, eliminating the certainty of a loss or guaranteeing some gain is valued more than simply modifying probabilities. When exerting control does not introduce or eliminate uncertainty faced, perceived value of control is found to not vary with prior probability of success except in situations where a loss is likely (greater than 0.5). Where a loss is likely, control is valued differentially (directed towards mitigating the chance of the worst outcome or improving the chance of the best outcome) depending on the adversity of the situation. While most of these findings are in line with the precepts of Cumulative Prospect Theory, they highlight that boundary effects may be more salient drivers of probability distortion in the presence of control, compared to when control is absent (e.g. Wu and Gonzalez, 1996).

5.1.2. Modelling control in Decision Analysis

To address the theoretical aims, existing discussions on modelling ‘imperfect’ control (Matheson and Matheson, 2005) were developed further, and a new general probability revision procedure for interventions, the *generalized generic controller*, was developed in section 3.3.2. Its close correspondence with belief revision rules in Philosophy that are

considered suitable to evaluate counterfactual statements ('imaging' (Gärdenfors, 1982; Lewis, 1976)) was also exposed in section 3.3.3. Additionally, two key properties — 1) fixed-point at zero; and 2) Bayes conditionalization preserving — that probability revision rules for interventions should satisfy were stated, and the anomalies resulting from the violation of these properties were demonstrated in section 3.4.1. The first property ensures that a probability revision rule does not assign positive probabilities to states that are known to not occur. The second property, which requires that post intervention probabilities satisfy Bayes rule, can be regarded as a coherence criterion for probability revision rules for interventions and mirror the coherence criterion that supports Bayes rule, the probability revision rule for information.

5.1.3. Application of Decision Analysis to areas of risk assessment and risk management

To test the feasibility of using the proposed probability revision procedure, the generalized generic controller was first operationalized as a *probability revision matrix*. Using this operationalization, its algebraic properties as a transformation of a probability simplex were exposed. The convenience with which the probability revision matrix can be integrated with risk assessment techniques in Operational Risks, such as the Loss Distribution Approach (Frachot et al., 2001) was demonstrated using a real world case study. This case study was conducted in collaboration with executives at a leading insurance company and was an endeavour to apply Decision Science methods to calculate the impact of risk mitigations on Operational Risk costs.

5.2. Main contributions of this research revisited

This research potentially makes three main contributions to the core literature on theoretical and behavioural research on decision making and offers two peripheral insights. Additionally, by demonstrating how the proposed developments of Decision Analysis can be used in practice to analyse managerial risk mitigations in the real world, the research possibly contributes to an area of risk management — Operational Risks — which was explored in this thesis. The main contributions of this research are:

- 1. A new understanding of the behavioral effects of control on risk taking:** By studying the effect of control on a wider range of probabilities and payoffs in this research, a main contribution of the experiment conducted in this research is that it offers a more complete understanding of how control affects risk taking within the context of fourfold pattern of risk attitudes predicted by Cumulative Prospect

Theory. The findings from the present study, that control has a moderating effect on risk taking, were not entirely expected and therefore cast a new perspective on the effect of control on risk taking. Using affective and cognitive explanations, it was discussed how these findings mirror discussions in the agency literature about how managers respond to risk. These findings can have implications for influencing or mitigating risk taking by introducing selective control in managerial tasks. From a methodological perspective, this study highlights the crucial mediating effect of decision frame and probability magnitudes on any general conclusion that can be made on the effect of one or more contextual factors on risk taking. The research also offers some peripheral insights on the close correspondence between the perceived value of exerting control and certainty effect. The enhanced perceived value of exerting control when it eliminates certain losses or guarantees some gains demonstrates that certainty in and of itself has value. The notion of valuing certainty is absent in most formal treatments and analysis of decisions and control (e.g. Expected Value of Control concept). This opens a direction for future study related to refining the analytic techniques used to compute the value of control to explicitly incorporate the impact of any certainty achieved.

2. **A tractable probability revision rule for modelling the effects of interventions on probabilities:** The proposed probability revision procedure for modelling the effect of actions on uncertainty — the generalized generic controller (GGC) — is a new addition to the Decision Analysis toolkit which can hopefully enhance the applicability of Decision Analysis in practice, especially in areas of risk management. As illustrated in section 3.3.3, a probability revision rule can significantly reduce the number of inputs required when analyzing interventions on uncertainties which depend on other uncertain variables. Furthermore, it enables an analyst to obtain judgments about probabilities and mitigations separately from different sources and use the proposed calculus to combine these judgments and thus compute the post-intervention probabilities. Consequently, if pre-intervention probability judgments about uncertain events are revised, then the post-intervention probabilities can be updated without requiring new inputs. These are all the desirable benefits of Bayes rule that were briefly summarized in section 1.2.2. These benefits, however, were only partially available when analyzing interventions using existing approaches. By developing the proposed probability revision procedure, this research has helped enhance the procedural convenience with which interventions are analyzed. A peripheral contribution relates to the interesting link that was

established between Decision Analysis and more theoretic areas of decision making such as Philosophy, which paves the way to explore potentially useful synergies between the two disciplines.

- 3. Theoretical foundations for probability revision rules for interventions:** One of the most compelling reasons for using Bayes rule for modelling the effect of actions on uncertainties is that it is grounded in consistency principles which are exposed using Dutch book arguments. By establishing two theoretical properties that probability revision rules for interventions should satisfy to avoid generating inconsistent beliefs, this research is a step towards giving the analysis of interventions the same normative anchorage as Bayes rule. Beyond being a subject of theoretical interest, these properties help constrain the possible probability revision rules for interventions, thus eliminating some arbitrariness from the proposed and other potential recommendations for probability revision rules for interventions.
- 4. Demonstrating an innovative and potentially useful integration of Decision Analysis methods with risk management:** The third paper, which was an application of Decision Analysis to areas of Operational Risk management, is a useful contribution to the methods for modelling Operational Risks. Existing methods, such as the Loss Distribution Approach, are relatively inflexible to the introduction of or changes in risk mitigation policies, since: 1) they rely on historical data; and 2) typically the effect of mitigations is reflected in the input parameters and a reassessment of the probability distributions can be tedious. The benefits and flexibility of using the proposed probability revision rule for modelling interventions, operationalized as a matrix, was unveiled using a real world case study, and it was shown that the procedure can be easily integrated with existing Operational Risk assessment techniques. This enables multiple mitigations to be analyzed quantitatively without much effort. Furthermore, the power of influence diagrams as a suitable front end tool to model the various relationships between the uncertain states and process them probabilistically was demonstrated, thus emphasizing the usefulness of Decision Analysis methods in practice.

In summary, the developments of this research offer novel insights on how control affects risk taking and contributes to the Decision Analysis toolkit to enable the role of control in decision making to be incorporated explicitly. It was exposed how these developments can help bridge the gap between managerial motivations for risk taking and behaviour observed

in laboratory settings as well as be useful for modelling real world mitigations, more conveniently and realistically. It is the hope that the proposed techniques offer a way to structure and analyse decisions so that they correspond more closely to managerial thinking.

5.3. Research limitations

The present research attempts to make some new contributions to the study of decision making under uncertainty from a number of perspectives, with the intent of better aligning its methods with how managers think about decisions and contemplate risk taking. However, given the breadth of topics covered across a range of disciplines in decision making, the depth of discussion for each topic was limited. Some of the main limitations of this research are:

- 1. Limitations of the empirical insights on how control affects risk taking:** One of the major limitations of the study presented in this thesis relates to the restricted selection of payoff and probability magnitudes employed as well as how the gambles were presented in the two conditions. These issues, however, are inherent to most laboratory studies on attitudes to risk and the need to impose discretionary constraints when designing the stimuli is almost indispensable. Another limitation is that, as they stand, the explanations offered for the observed findings of this study also lack direct empirical support, even though they draw on previous empirical work. A crude attempt was made to extract the self-reported cognitive explanations of subjects regarding the basis for their decisions in this study, but it was found to be only remotely informative.
- 2. Limitations of how control was manipulated in the empirical study:** The manipulation of control in this study — as a fixed probability change — is, arguably, somewhat artificial and divorced from real world conceptions of control. In the real world, decision makers experience varied levels of control depending on their circumstances and perception; and the relationship between risk taking and control is probably also linked to the level of control individuals feel they have. Any systematic relationship between levels of control experienced and the moderating effect of control discovered in this study ought to be examined for obtaining a more complete understanding of the effect of control on risk taking.
- 3. Input elicitation for a probability revision rule for interventions:** A potential criticism and practical limitation of the probability revision procedure presented in section 3.3.2 is that eliciting the inputs for the generalized generic controller (GGC) is not very different from eliciting probabilities and therefore this procedure inherits

all the cognitive biases that probability judgments suffer from (e.g. von Winterfeldt and Edwards, 1986a). Although a procedure analogous to eliciting probabilities was presented, this procedure is still tedious and simpler methods will need to be devised to make the use of the generalized generic controller more appealing. Another limitation of the research is that it only exposed the procedural similarities between the generalized generic controller and imaging, and not much was said about the cognitive or information processing advantages of this equivalence. Imaging has been proposed as a way to assign probabilities to counterfactual statements and there have been passing claims that this approach fitly applies to the analysis of the effect of actions on uncertainties (Joyce, 1999; Pearl, 2000). A deeper discussion of how this underpinning feature of imaging can be utilized to analyze interventions in the context of Decision Analysis is needed.

- 4. Strengthening the theoretical foundations for a probability revision rule for interventions:** While this research attempted to ground probability revision rules for interventions in consistency principles by stating a couple of basic properties they should satisfy, there is scope to develop these theoretical foundations further. Note that although it was shown that the GGC satisfies the stated basic properties of probability revision functions for interventions, the converse may not hold. This leaves open the possibility of other probability revision rules for modelling interventions that could also be suitable.
- 5. The probability revision matrix and its applicability:** In this research, the discussion of the GGC and its operationalization in section 4.3 focused only on discrete probability distributions. This limits the use of the GGC as it is not immediately clear how it can be applied to model influences on probability distributions described over continuous random variables. Developing an extension of the GGC which can be applied to continuous probability distributions is a natural direction for future research. The practical application of the probability revision matrix that was explored assumed hypothetical inputs. While an attempt was made to elicit the inputs using a questionnaire, the reliability of the procedure remains disputable. One of the challenges experienced was suitable framing of the elicitation question to properly fit with the definition of the inputs of the probability revision matrix and also ensure that it is not confusing. While performing sensitivity analysis (Figure 4-8), can address some of the problems related to the ambiguity regarding inputs, a rigorous procedure for de-biasing the inputs is desirable and ought to be developed.

5.4. Directions for further research

This research is a small step in the direction of developing the formal study of control in Decision Sciences. Promising avenues for further research emerge from questions that remain unaddressed and the discovery of new questions.

- 1. Further testing of the empirical observations and explanations of the effect of control on risk taking:** The present study is one of the first attempts at explaining the effect of control in the context where risk seeking behavior prevails when uncertainties cannot be controlled (e.g. losses only payoffs and probability levels). It paves the way for a number of interesting empirical studies. First, variants of this study need to be conducted to check for the robustness of the moderating effect of control on risk attitudes that was revealed in this study. One variation that merits research is preserving the structure of gambles (i.e. two gambles each with two outcomes) in the two conditions (with and without control), as the number of outcomes in a ‘gamble’ can affect risk taking (Brooks et al., 2013; Payne, 2005). Future work on variations of the design presented here should also endeavor to go beyond simple gambles and use narratives that describe corporate risks, when testing for the effect of control on decision making. Another important area for further research, which has emerged from this study, is a direct test of the explanations offered for the observed effects of control on risk taking. Future studies should seek to capture the affective moods and cognitive explanations in a more sophisticated manner to enable a clearer understanding of *why* control affects risk taking.
- 2. Conducting empirical research on perceived levels of exerting control:** Given that, in the real world, the extent to which uncertainty can be controlled is often a matter of perception, it would be interesting to explore how levels of control experienced affect risk taking and the perceived value of exerting control. There is also scope to take the study on perceived value of control further to investigate preferences for exerting control using alternative dependent variables, such as willingness to pay for exerting control. Finally it will be interesting to explore how perceived control affects not only risk taking but if it is a determining factor of choices people make or why some options are preferred over others. For example, are people willing to accept riskier options with some control over uncertainty instead of safe options with no control over uncertainty? Taking this line of inquiry further, it would be useful to develop a quantitative way to directly measure the risk offsetting effect of a given level of control (see related concept in Huber, 2007).

- 3. Tying the input elicitation for a probability revision rule for interventions to human reasoning:** One limitation of the present research which can impede its usability is that eliciting the inputs for the GGC is not straightforward. Therefore, more creative methods which align the elicitation with how individuals naturally think about interventions need to be devised. One direction in which this can be explored is by building on the relationship between the GGC and imaging exposed in this research, to find a way to relate the input elicitation to counterfactual reasoning. Establishing this relationship has prescriptive value as, in addition to facilitating the elicitation, it can provide analytic support to corporate decisions, such as those pertaining to identifying and developing new opportunities, where counterfactual reasoning plays a salient role (e.g. Gaglio, 2004).
- 4. Strengthening the theoretical foundations for a probability revision rule for interventions:** As mentioned, the theoretical foundations proposed in this research, are still short of providing a compelling case for using the GGC or, equivalently, imaging as the probability revision rule for interventions. One area for further research is to find more properties that probability revision functions for interventions should satisfy so that the GGC is both a necessary and sufficient condition for the established properties to be satisfied.
- 5. Leveraging the probability revision matrix:** It is worth observing that the operationalization of the GGC as a probability revision matrix has a close correspondence with the techniques of Markov Decision Processes (MDP) (Howard, 1960). Exploring the similarities or differences between the methods of MDP constitutes an interesting area for future research and can also inform the extension of the GGC to continuous probability distributions. Another interesting and potentially useful area of research builds upon the algebraic description of the GGC. It was mentioned earlier that the formal analysis of control such as the Expected Value of Control concept is typically based on the impact exerting control has on the ‘expected utility’ and is not sensitive to the value attributed to mitigating uncertainty itself. The simplex representation can be useful for developing metrics that assess the value of mitigating uncertainty (in the desired direction). Some premature suggestions on this front entail developing metrics that are based on the change in volume of the simplex or the distance of the centroid from the most favorable vertex. The development of metrics that capture pure uncertainty mitigation can be helpful when it is difficult to assign numeric scores to the consequences (e.g. impact of climate change).

6. Normative/Philosophical issues related to non-canonical influence diagrams:

Finally, it is worth exploring an issue which was mentioned at the outset of the research but was left unaddressed: the issue pertaining to the normative underpinnings of an influence diagram which is not in the canonical form. While normative theories of decision making which allow for probabilities of uncertain states to depend on an action have been proposed (Jeffrey, 1965; Armendt, 1986; Luce and Krantz, 1971; Joyce, 1999), the interpretation of the probabilities is subject to debate. For a given decision problem, different interpretations can result in different expected utility calculations (e.g. Armendt, 1986; Gibbard and Harper, 1978; Maher, 1990). Understanding which of these decision theories mesh with expected utility calculations that support influence diagrams constitutes an interesting direction for future research. Developments in this direction can help strengthen the normative status of influence diagrams, thus making its use in practice more compelling.

On a closing note, it is worth drawing attention to a subject area that the field of Decision Analysis has steered clear of, and was therefore also avoided in this thesis: the notion of *causality* (Shachter, 2012). Causality has been an active area of discussion among some (causal) decision theorists who advocate that a recommendation procedure for actions should be based on the good that an agent can *cause* and not on the evidential or ‘symptomatic’ properties (e.g. Joyce, 1999; Maher, 1990). In other words, the way an agent reasons about the potential effects of his or her actions on uncertainty should be different from how he or she contemplates the effect of information and this distinction is crucially important when the outcomes of an uncertain variable are correlated with the an action. The notion of causal influence is embedded in influence diagrams but has not been explored sufficiently (Shachter, 2012). Exploring the unaddressed subtleties between Decision Analysis methods and causal thinking (Shachter, 2012), and better formalizing this association to reflect how managers think about their decisions and capabilities as *causing* outcomes, merits future research. This can help extend the scope of Decision Analysis tools and make them pertinent in situations where corporate risk management goes beyond gearing choices towards reducing risks, to converting them into an opportunity or “making things happen” (March and Shapira, 1987).

5.5. References for Chapter 5

1. Armendt, B., 1986. A Foundation for Causal Decision Theory. *Topoi* 5, 3–19.
2. Brandstätter, E., Schwarzenberger, H., 2001. Beyond the gambling paradigm: Internal controllability in decision-making. *Psychol. Rep.* 89, 259–266.
3. Brooks, P., Peters, S., Zank, H., 2013. Risk behavior for gain, loss, and mixed prospects. *Theory Decis.* 1–30.
4. Chatterton, J., 2001. Uncertainty: Types and Sources (Weak Links project: Case Studies.), Weak Links project: Case Studies. Building Research Establishment and Association of British Insurers.
5. Chernoff, H., Moses, L., 1959. *Elementary Decision Theory*. Wiley, New York.
6. Ellsberg, D., 1961. Risk, ambiguity and the Savage axioms. *Q. J. Econ.* 75, 643–669.
7. Fox, C., Ulkumen, G., 2011. Distinguishing two dimensions of uncertainty.
8. Frachot, A., Georges, P., Roncalli, T., 2001. Loss distribution approach for operational risk. Available SSRN 1032523.
9. Gaglio, C., 2004. The Role of Mental Simulations and Counterfactual Thinking in the Opportunity Identification Process*. *Entrep. Theory Pract.* 28, 533–552.
10. Gardenfors, P., 1982. Imaging and Conditionalization. *J. Philos.* 79, 747–760.
11. Gibbard, A., Harper, W., 1978. Counterfactuals and two kinds of expected utility, in: *Ifs*, (Eds.). Reidel, Dordrecht.
12. Goodie, A., 2003. The effects of control on betting: Paradoxical betting on items of high confidence with low value. *J. Exp. Psychol. Learn. Mem. Cogn.* 29, 598–610.
13. Goodie, A., Young, D., 2007. The skill element in decision making under uncertainty: Control or competence? *Judgm. Decis. Mak.* 2, 189–203.
14. Hoffman, O., Hammonds, J., 1994. Propagation of uncertainty in risk assessments: The need to distinguish between uncertainty due to lack of knowledge and uncertainty due to variability. *Risk Anal.* 14.
15. Howard, R., 1960. *Dynamic Programming and Markov Processes*. MIT Press.
16. Howard, R., Matheson, J., 1984. Influence Diagrams, in: R.A. Howard and J.E. Matheson (Eds.), *The Principles and Applications of Decision Analysis*. Strategic Decisions Group, Menlo Park, CA.
17. Howard, R., Matheson, J., 2005. Influence diagrams. *Decis. Anal.* 2, 127–143.
18. Huber, O., 2002. Beyond gambles and lotteries. *Decis. Mak. Cogn. Models Explan.* 145.
19. Huber, O., 2007. Behavior in risky decisions: Focus on risk defusing, in: *Uncertainty and Risk*. Springer, pp. 291–306.
20. Isen, A., Patrick, R., 1983. The effect of positive feelings on risk taking: When the chips are down. *Organ. Behav. Hum. Perform.* 31, 194–202.

21. Jeffrey, R., 1965. *The Logic of Decision*, 2nd ed. University of Chicago Press, Chicago.
22. Joyce, J., 1999. *The Foundations of Causal Decision Theory*. Cambridge University Press, Cambridge Books Online.
23. Langer, E., 1975. The illusion of control. *J. Pers. Soc. Psychol.* 32, 311–328.
24. Lewis, D., 1976. Probabilities of Conditionals and Conditional Probabilities. *Philos. Rev.* 85, 297–315.
25. Li, K., 2011. Preference towards control in risk taking: Control, no control, or randomize? *J. Risk Uncertain.* 43, 39–63.
26. Luce, D., Krantz, D., 1971. Conditional Expected Utility. *Econometrica* 39, 253–271.
27. Maher, P., 1990. Symptomatic Acts and the Value of Evidence in Causal Decision Theory. *Philos. Sci.* 57, 479–498.
28. March, J., Shapira, Z., 1987. Managerial perspectives on risk and risk taking. *Manag. Sci.* 33, 1404–1418.
29. Matheson, D., Matheson, J., 2005. Describing and Valuing Interventions That Observe or Control Decision Situations. *Decis. Anal.* 2, 165–81.
30. Matheson, J., 1990. Using influence diagrams to value information and control, in: R.M. Oliver & J.Q. Smith, *Influence Diagrams, Belief Nets and Decision Analysis*. John Wiley & Sons Ltd, pp. 25–48.
31. Morgan, M., Henrion, M., 1990. The Nature and Sources of Uncertainty, in: *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge University Press, Cambridge, pp. 47–56.
32. Payne, J., 2005. It is Whether You Win or Lose: The Importance of the Overall Probabilities of Winning or Losing in Risky Choice. *J. Risk Uncertain.* 30, 5–19.
33. Pearl, J., 2000. *Causality: Models, Reasoning and Inference*. Cambridge University Press, Cambridge.
34. Savage, L., 1954. *The Foundations of Statistics*. John Wiley & Sons, New York.
35. Seo, M.-G., Goldfarb, B., Barrett, L., 2010. Affect and the framing effect within individuals over time: Risk taking in a dynamic investment simulation. *Acad. Manage. J.* 53, 411–431.
36. Shachter, R., 2012. Effective Decision Analysis: A Cause to Believe In. *Decis. Anal. Today* 31, 29–33.
37. Tversky, A., Kahneman, D., 1992. Advances in prospect theory: cumulative representation of uncertainty. *J. Risk Uncertain.* 5, 297.
38. Ulkumen, G., Erner, C., Fox, C., Tannenbaum, D., Walters, D., 2014. Perceived Nature of Market Uncertainty Predicts Investment Behaviors. Presented at the Behavioural Decision Research in Management, London.

-
39. Von Mises, L., 1949. *Human Action: A Treatise on Economics*, 4th ed. Fox & Wilkes, San Francisco.
 40. Von Winterfeldt, D., Edwards, W., 1986. Cognitive illusions, in: *Decision Analysis and Behavioral Research*. Cambridge University Press, Cambridge.
 41. Wiseman, R., Gomez-Mejia, L., 1998. A Behavioral Agency Model of Managerial Risk Taking. *Acad. Manage. Rev.* 23, 133–153.
 42. Wu, G., Gonzalez, R., 1996. Curvature of the Probability Weighting Function. *Manag. Sci.* 42, 1676–1690.
 43. Young, D., Goodie, A., Hall, D., 2011. Modeling the Impact of Control on the Attractiveness of Risk in a Prospect Theory Framework. *J. Behav. Decis. Mak.* 24, 47–70.

Chapter 6

Consolidated List of References

1. Abdellaoui, M., Baillon, A., Placido, L., Wakker, P., 2011. The Rich Domain of Uncertainty: Source Functions and Their Experimental Implementation. *Am. Econ. Rev.* 101, 695–723.
2. Agarwal, S., 2014. The Moderating Effect of ‘Probability Control’ on Risk Taking.
3. Agarwal, S., Montibeller, G., 2014. A Quantitative Method for Measuring the Value of Operational Risk Mitigations.
4. Agarwal, S., Montibeller, G., Morton, A., 2014. Stochastic Interventions: A New Probability Revision Method to Model the Effects of Actions on Uncertainties.
5. Anderson, R., Fraser, D., 2000. Corporate control, bank risk taking, and the health of the banking industry. *J. Bank. Finance* 24, 1383–1398.
6. Armendt, B., 1986. A Foundation for Causal Decision Theory. *Topoi* 5, 3–19.
7. Arrow, K., 1951. Alternative Approaches to the Theory of Choice in Risk-Taking Situations. *Econometrica* 19, 404–437.
8. Arrow, K., 1971. *Essays in the Theory of Risk Bearing*, Markham economics series. Markham Pub. Co., Chicago.
9. Basel Committee, 2003. *Sound Practices for the Management and Supervision of Operational Risk*.
10. Basel Committee, 2005. *International Convergence of Capital Measurement and Capital Standards: A Revised Framework*. Bank for International Settlements, Basel, Switzerland.
11. Basel Committee, 2011. *Operational Risk – Supervisory Guidelines for the Advanced Measurement Approaches*.
12. Bedford, T., Cooke, R., 2001. *Probabilistic Risk Analysis: Foundations and Methods*. Cambridge University Press, Cambridge.
13. Bell, D., 1982. Regret in decision making under uncertainty. *Oper. Res.* 30, 961–981.
14. Bell, D., Raiffa, H., Tversky, A., 1988. Descriptive, Normative, and Prescriptive Interactions in Decision Making, in: D. Bell, H. Raiffa, A. Tversky. *Decision Making: Descriptive, Normative and Prescriptive Interactions*. Cambridge University Press, Cambridge.
15. Bernoulli, D., 1738. Exposition of a New Theory on the Measurement of Risk. *Econometrica* 22, 23–36.
16. Bernstein, P., 1998. *Against the Gods: The Remarkable Story of Risk*. John Wiley & Sons, New York.

17. Bolker, E., 1967. A simultaneous axiomatization of utility and subjective probability. *Philos. Sci.* 333–340.
18. Brandstätter, E., Kühberger, A., Schneider, F., 2002. A cognitive-emotional account of the shape of the probability weighting function. *J. Behav. Decis. Mak.* 15, 79–100.
19. Brandstätter, E., Schwarzenberger, H., 2001. Beyond the gambling paradigm: Internal controllability in decision-making. *Psychol. Rep.* 89, 259–266.
20. Brooks, P., Peters, S., Zank, H., 2013. Risk behavior for gain, loss, and mixed prospects. *Theory Decis.* 1–30.
21. Brooks, P., Zank, H., 2005. Loss averse behavior. *J. Risk Uncertain.* 31, 301–325.
22. Bryant, P., Dunford, R., 2008. The Influence of Regulatory Focus on Risky Decision-Making. *Appl. Psychol.* 57, 335–359.
23. Buede, D., 2005. Influence diagrams: A practitioner’s perspective. *Decis. Anal.* 2, 235–237.
24. Byrnes, J., Miller, D., Schafer, W., 1999. Gender differences in risk taking: A meta-analysis. *Psychol. Bull.* 125, 367.
25. Caliendo, M., Fossen, F., Kritikos, A., 2009. Risk attitudes of nascent entrepreneurs—new evidence from an experimentally validated survey. *Small Bus. Econ.* 32, 153–167.
26. Camerer, C., 1989. An experimental test of several generalized utility theories. *J. Risk Uncertain.* 2, 61–104.
27. Camerer, C., 1992. Recent tests of generalizations of expected utility theory, in: *Utility Theories: Measurements and Applications*. Springer, pp. 207–251.
28. Camerer, C., 2004. Prospect theory in the wild: Evidence from the field. Colin F Camerer George Loewenstein Matthew Rabin Eds *Adv. Behav. Econ.* 148–161.
29. Camerer, C., Ho, T.-H., 1994. Violations of the betweenness axiom and nonlinearity in probability. *J. Risk Uncertain.* 8, 167–196.
30. Carbone, E., Hey, J., 1994. Estimation of expected utility and non-expected utility preference functionals using complete ranking data, in: *Models and Experiments in Risk and Rationality*. Springer, pp. 119–139.
31. Cavagnaro, D., Pitt, M., Gonzalez, R., Myung, J., 2013. Discriminating among probability weighting functions using adaptive design optimization. *J. Risk Uncertain.* 47, 255–289.
32. Charness, G., Gneezy, U., 2012. Strong evidence for gender differences in risk taking. *J. Econ. Behav. Organ.* 83, 50–58.
33. Chatterton, J., 2001. Uncertainty: Types and Sources (Weak Links project: Case Studies.), Weak Links project: Case Studies. Building Research Establishment and Association of British Insurers.

34. Chau, A., Phillips, J., 1995. Effects of perceived control upon wagering and attributions in computer blackjack. *J. Gen. Psychol.* 122, 253.
35. Chaudhury, M., 2010. A review of the key issues in operational risk capital modeling. *J. Oper. Risk* 5, 37–66.
36. Chavez-Demoulin, V., Embrechts, P., 2004. Advanced extremal models for operational risk. Prepr. Dep. Math. ETH-Zent.
37. Chavez-Demoulin, V., Embrechts, P., Nevlehová, J., 2006. Quantitative models for operational risk: extremes, dependence and aggregation. *J. Bank. Finance* 30, 2635–2658.
38. Chelst, K., Bodily, S., 2000. Structured Risk Management: Filling a Gap in Decision Analysis Education. *J. Oper. Res. Soc.* 51, 1420–1432.
39. Chernoff, H., Moses, L., 1959. *Elementary Decision Theory*. Wiley, New York.
40. Chew, S.H., Sagi, J., 2008. Small worlds: Modeling attitudes toward sources of uncertainty. *J. Econ. Theory* 139, 1–24.
41. Chow, C., Sarin, R., 2002. Known, Unknown, and Unknowable Uncertainties. *Theory Decis.* 52, 127–138.
42. Cohen, J., 1960. *Chance, Skill and Luck*. Penguin books, Middlesex.
43. Cohen, J., Hansel, C.E.M., 1959. Preferences for different combinations of chance and skill in gambling.
44. Courtney, H., 2001. *20/20 Foresight: Crafting Strategy in an Uncertain World*. Harvard Business School Press, Boston, MA.
45. Courtney, H., Kirkland, J., Viguerie, P., 1997. Strategy Under Uncertainty. *Harv. Bus. Rev.* 75, 67–79.
46. Crowe, E., Higgins, E.T., 1997. Regulatory Focus and Strategic Inclinations: Promotion and Prevention in Decision-Making. *Organ. Behav. Hum. Decis. Process.* 69, 117–132.
47. Cubitt, R., Sugden, R., 2001. On money pumps. *Games Econ. Behav.* 37, 121–160.
48. Cummins, J., Embrechts, P., 2006. Introduction: Special section on operational risk. *J. Bank. Finance* 30, 2599–2604.
49. Edwards, W., 1955. The prediction of decisions among bets. *J. Exp. Psychol.* 50, 201.
50. Ellsberg, D., 1961. Risk, ambiguity and the Savage axioms. *Q. J. Econ.* 75, 643–669.
51. Embrechts, P., Furrer, H., Kaufmann, R., 2003. Quantifying regulatory capital for operational risk. *Deriv. Use Trading Regul.* 9, 217–233.
52. Embrechts, P., Kaufmann, R., Samorodnitsky, G., 2004. Ruin theory revisited: stochastic models for operational risk. *Risk Manag. Cent. Bank Foreign Reserv.* 243–261.
53. Fennema, H., Wakker, P., 1997. Original and cumulative prospect theory: a discussion of empirical differences. *J. Behav. Decis. Mak.* 10, 53–64.

54. Fishburn, P., 1982. Nontransitive measurable utility. *J. Math. Psychol.* 26, 31–67.
55. Fox, C., Ulkumen, G., 2011. Distinguishing two dimensions of uncertainty.
56. Frachot, A., Georges, P., Roncalli, T., 2001. Loss distribution approach for operational risk. Available SSRN 1032523.
57. Frachot, A., Moudoulaud, O., Roncalli, T., 2003. Loss distribution approach in practice. *Basel Handb. Guide Financ. Pract.*
58. Franzetti, C., 2011. *Operational risk modelling and management*. Chapman and Hall/CRC.
59. Fung, M., 2014. Cumulative Prospect Theory and Managerial Incentives for Fraudulent Financial Reporting. *Contemp. Account. Res.*
60. Gaglio, C., 2004. The Role of Mental Simulations and Counterfactual Thinking in the Opportunity Identification Process*. *Entrep. Theory Pract.* 28, 533–552.
61. Gardenfors, P., 1982. Imaging and Conditionalization. *J. Philos.* 79, 747–760.
62. Gardenfors, P., 1988. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. MIT Press, Cambridge Massachusetts.
63. Gibbard, A., Harper, W., 1978. Counterfactuals and two kinds of expected utility, in: *Ifs*, (Eds.). Reidel, Dordrecht.
64. Gilboa, I., 1999. Can free choice be known?, in: *Bicchieri et Al., The Logic of Strategy*. Oxford University Press, Oxford, pp. 163–174.
65. Gilboa, I., Postlewaite, A., Schmeidler, D., 2012. Rationality of belief or: why Savage's axioms are neither necessary nor sufficient for rationality. *Synthese* 187, 11–31.
66. Goldstein, W., Einhorn, H., 1987. Expression theory and the preference reversal phenomena. *Psychol. Rev.* 94, 236.
67. Gonzalez, R., Wu, G., 1999. On the Shape of the Probability Weighting Function. *Cognit. Psychol.* 38, 129–166.
68. Goodie, A., 2003. The effects of control on betting: Paradoxical betting on items of high confidence with low value. *J. Exp. Psychol. Learn. Mem. Cogn.* 29, 598–610.
69. Goodie, A., Young, D., 2007. The skill element in decision making under uncertainty: Control or competence? *Judgm. Decis. Mak.* 2, 189–203.
70. Hacking, J., 1975. *The Emergence of Probability*. Cambridge University Press, Cambridge.
71. Hájek, A., 2008. Dutch book arguments. *Oxf. Handb. Ration. Soc. Choice* 173–195.
72. Harless, D., Camerer, C., 1994. The Predictive Utility of Generalized Expected Utility Theories. *Econometrica* 62, 1251–1289.
73. Heath, C., Tversky, A., 1991. Preference and belief: Ambiguity and competence in choice under uncertainty. *J. Risk Uncertain.* 4, 5–28.

74. Heckerman, D., Shachter, R., 1994. A decision-based view of causality, in: *Proceedings of the Tenth International Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers Inc., Seattle, WA, pp. 302–310.
75. Heckerman, D., Shachter, R., 1995. Decision Theoretic Foundations for Causal Reasoning. *J. Artif. Intell. Res.* 3, 405–430.
76. Higgins, E.T., 1996a. Emotional Experiences: The Pains and Pleasures of Distinct Regulatory Systems. *Emot. Interdiscip. Perspect.* 203.
77. Higgins, E.T., 1996b. Ideals, oughts, and regulatory focus. *Psychol. Action Link. Cogn. Motiv. Behav.* 91–114.
78. Higgins, E.T., 1997. Beyond pleasure and pain. *Am. Psychol.* 52, 1280.
79. Higgins, E.T., 1998. Promotion and prevention: Regulatory focus as a motivational principle. *Adv. Exp. Soc. Psychol.* 30, 1–46.
80. Hoffman, O., Hammonds, J., 1994. Propagation of uncertainty in risk assessments: The need to distinguish between uncertainty due to lack of knowledge and uncertainty due to variability. *Risk Anal.* 14.
81. Holton, G., 2002. History of Value-at-Risk. Working paper. Contingency Analysis, Boston.
82. Holton, G., 2004. Defining risk. *Financ. Anal. J.* 60.
83. Horswill, M., McKenna, F., 1999. The Effect of Perceived Control on Risk Taking. *J. Appl. Soc. Psychol.* 29, 377–391.
84. Howard, R., 1960. *Dynamic Programming and Markov Processes*. MIT Press.
85. Howard, R., 1988. Decision analysis: practice and promise. *Manag. Sci.* 34, 679–695.
86. Howard, R., 1990. From influence to relevance to knowledge., in: R.M.Oliver & J.Q. Smith, *Influence Diagrams, Belief Nets and Decision Analysis*. John Wiley & Sons Ltd.
87. Howard, R., Matheson, J., 1984. Influence Diagrams, in: R.A. Howard and J.E. Matheson (Eds.), *The Principles and Applications of Decision Analysis*. Strategic Decisions Group, Menlo Park, CA.
88. Howard, R., Matheson, J., 2005. Influence diagrams. *Decis. Anal.* 2, 127–143.
89. Howell, W., 1971. Uncertainty from internal and external sources: A clear case of overconfidence. *J. Exp. Psychol.* 89, 240.
90. Howell, W., Burnett, S., 1978. Uncertainty measurement: A cognitive taxonomy. *Organ. Behav. Hum. Perform.* 22, 45–68.
91. Huber, O., 1997. Beyond gambles and lotteries: Naturalistic risky decisions. *Decis. Mak. Cogn. Models Explan.* (Eds.) 145–162.
92. Huber, O., 2007. Behavior in risky decisions: Focus on risk defusing, in: *Uncertainty and Risk*. Springer, pp. 291–306.

93. Humphreys, P., Berkeley, D., 1985. Handling uncertainty: levels of analysis of decision problems, in: G. Wright, Behavioral Decision Making.
94. Idson, L.C., Liberman, N., Higgins, E.T., 2000. Distinguishing Gains from Nonlosses and Losses from Nongains: A Regulatory Focus Perspective on Hedonic Intensity. *J. Exp. Soc. Psychol.* 36, 252–274.
95. Isen, A., Patrick, R., 1983. The effect of positive feelings on risk taking: When the chips are down. *Organ. Behav. Hum. Perform.* 31, 194–202.
96. Jeffrey, R., 1965. *The Logic of Decision*, 2nd ed. University of Chicago Press, Chicago.
97. Joyce, J., 1999. *The Foundations of Causal Decision Theory*. Cambridge University Press, Cambridge Books Online.
98. Kahneman, D., Tversky, A., 1979. Prospect Theory: An analysis of decision under risk. *Econometrica* XLVII, 263–291.
99. Keynes, J., 1936. *The General Theory of Employment, Interest and Money*. Macmillan Cambridge University Press, Cambridge.
100. King, J., 2001. *Operational Risk: Measurement and Modelling*. John Wiley & Sons Ltd, Chichester.
101. Kluger, A., Stephan, E., Ganzach, Y., HersHKovitz, M., 2004. The effect of regulatory focus on the shape of probability-weighting function: Evidence from a cross-modality matching method. *Organ. Behav. Hum. Decis. Process.* 95, 20–39.
102. Knight, F., 1921. *Risk, Uncertainty and Profit*. Houghton Mifflin Co., Boston.
103. Kocher, M., Pahlke, J., Trautmann, S., 2013. Tempus Fugit: Time Pressure in Risky Decisions. *Manag. Sci.*
104. Lang, S., 1966. The matrix associated with a linear map, in: *Linear Algebra*. Addison-Wesley Publishing Co., USA.
105. Langer, E., 1975. The illusion of control. *J. Pers. Soc. Psychol.* 32, 311–328.
106. Leavens, D., 1945. Diversification of investments. *Trusts Estates* 80, 469–473.
107. Lefebvre, M., Vieider, F., 2014. Risk taking of executives under different incentive contracts: Experimental evidence. *J. Econ. Behav. Organ.* 97, 27–36.
108. Lerner, J., Keltner, D., 2001. Fear, anger, and risk. *J. Pers. Soc. Psychol.* 81, 146.
109. Lewis, D., 1976. Probabilities of Conditionals and Conditional Probabilities. *Philos. Rev.* 85, 297–315.
110. Lewis, D., 1981. Causal decision theory. *Australas. J. Philos.* 59, 5–30.
111. Li, K., 2011. Preference towards control in risk taking: Control, no control, or randomize? *J. Risk Uncertain.* 43, 39–63.
112. Loewenstein, G., Weber, E., Hsee, C., Welch, N., 2001. Risk as feelings. *Psychol. Bull.* 127, 267–286.

113. Loomes, G., Sugden, R., 1986. Disappointment and dynamic consistency in choice under uncertainty. *Rev. Econ. Stud.* 53, 271–282.
114. Lubbe, J., Snyman, F., 2010. The advanced measurement approach for banks. IPM17 Implement. 1993 Syst. Natl. Acc. 141.
115. Luce, D., Krantz, D., 1971. Conditional Expected Utility. *Econometrica* 39, 253–271.
116. MacCrimmon, K., Wehrung, D., 1986. *Taking Risks: The Management of Uncertainty*. Macmillan, London.
117. MacCrimmon, K., Wehrung, D., 1990. Characteristics of Risk Taking Executives. *Manag. Sci.* 36, 422–435.
118. Machina, M., 1987. Choice Under Uncertainty: Problems Solved and Unsolved. *Econ. Perspect.* 1, 121–154.
119. Maher, P., 1990. Symptomatic Acts and the Value of Evidence in Causal Decision Theory. *Philos. Sci.* 57, 479–498.
120. March, J., Shapira, Z., 1987. Managerial perspectives on risk and risk taking. *Manag. Sci.* 33, 1404–1418.
121. Markowitz, H., 1952. Portfolio Selection. *J. Finance* 7, 77–91.
122. Matheson, D., Matheson, J., 2005. Describing and Valuing Interventions That Observe or Control Decision Situations. *Decis. Anal.* 2, 165–81.
123. Matheson, J., 1990. Using influence diagrams to value information and control, in: R.M.Oliver & J.Q. Smith, *Influence Diagrams, Belief Nets and Decision Analysis*. John Wiley & Sons Ltd, pp. 25–48.
124. Meyer, C., 2001. *Matrix Analysis and Applied Linear Algebra*. Society for Industrial and Applied Mathematics.
125. Moosa, I., 2007. *Operational risk management*. Palgrave Macmillan.
126. Morgan, M., Henrion, M., 1990. The Nature and Sources of Uncertainty, in: *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge University Press, Cambridge, pp. 47–56.
127. Neil, M., Fenton, N., Tailor, M., 2005. Using Bayesian Networks to Model Expected and Unexpected Operational Losses. *Risk Anal.* 25, 963–972.
128. Nozick, R., 1969. Newcomb’s Problem and Two Principles of Choice, in: N. Rescher (Ed.), *Essays in Honor of Carl G. Hempel: A Tribute on His Sixty-Fifth Birthday*. D.Reidel, Dordrecht-Holland, pp. 114–146.
129. Olsen, R., 1997. Prospect theory as an explanation of risky choice by professional investors: Some evidence. *Rev. Financ. Econ.* 6, 225–232.
130. Pate-Cornell, M.E., 1996. Uncertainties in risk analysis: Six levels of treatment. *Reliab. Eng. Syst. Saf.* 54, 95–111.
131. Payne, J., 2005. It is Whether You Win or Lose: The Importance of the Overall Probabilities of Winning or Losing in Risky Choice. *J. Risk Uncertain.* 30, 5–19.

132. Payne, J., Bettman, J., Schkade, D., 1999. Measuring Constructed Preferences: Towards a Building Code. *J. Risk Uncertain.* 19, 243–270–270.
133. Pearl, J., 1994. A Probabilistic Calculus of Actions, in: *Proceedings of the Tenth Annual Conference on Uncertainty in Artificial Intelligence* (eds. R. Lopez de Mantaras and D. Poole). pp. 454–462.
134. Pearl, J., 2000. *Causality: Models, Reasoning and Inference*. Cambridge University Press, Cambridge.
135. Pearl, J., 2010. Physical and Metaphysical Counterfactuals, in: *Technical Report R-359*.
136. Pratt, J., 1964. Risk aversion in the small and in the large. *Econometrica* 32, 122–136.
137. Prelec, D., 1998. The Probability Weighting Function. *Econometrica* 66, 497–527.
138. Presson, P., Benassi, V., 1996. Illusion of control: A meta-analytic review. *J. Soc. Behav. Personal.* 11, 493–510.
139. Quiggin, J., 1982. A theory of anticipated utility. *J. Econ. Behav. Organ.* 3, 323–343.
140. Rosenzweig, P., 2013. What Makes Strategic Decisions Different. *Harv. Bus. Rev.* 89–93.
141. Rosenzweig, P., 2014. *Left Brain, Right Stuff: How Leaders Make Winning Decisions*. Public Affairs, USA.
142. Rottenstreich, Y., Hsee, C., 2001. Money, kisses, and electric shocks: On the affective psychology of risk. *Psychol. Sci.* 12, 185–190.
143. Roy, A., 1952. Safety first and the holding of assets. *Econom. J. Econom. Soc.* 431–449.
144. Savage, L., 1954. *The Foundations of Statistics*. John Wiley & Sons, New York.
145. Segal, U., 1987. Some remarks on Quiggin’s anticipated utility. *J. Econ. Behav. Organ.* 8, 145–154.
146. Seo, M.-G., Goldfarb, B., Barrett, L., 2010. Affect and the framing effect within individuals over time: Risk taking in a dynamic investment simulation. *Acad. Manage. J.* 53, 411–431.
147. Shachter, R., 1986. Evaluating Influence Diagrams. *Oper. Res.* 34, 871–882.
148. Shachter, R., 2012. Effective Decision Analysis: A Cause to Believe In. *Decis. Anal. Today* 31, 29–33.
149. Shachter, R., Heckerman, D., 2010. Pearl Causality and Value of Control, in: R. Dechter, H. Geffner and J. Y. Halpern (Eds.), *Heuristics, Probability, and Causality: A Tribute to Judea Pearl*. College Publications, pp. 431–447.
150. Shephard, G., Kirkwood, C., 1994. Managing the judgmental probability elicitation process: a case study of analyst/manager interaction. *Eng. Manag. IEEE Trans. On* 41, 414–425.

151. Sitkin, S., Pablo, A., 1992. Reconceptualizing the determinants of risk behavior. *Acad. Manage. Rev.* 17, 9–38.
152. Skyrms, B., 1982. Causal Decision Theory. *J. Philos.* 79, 695–711.
153. Sokol-Hessner, P., Camerer, C., Phelps, E., 2013. Emotion regulation reduces loss aversion and decreases amygdala responses to losses. *Soc. Cogn. Affect. Neurosci.* 8, 341–350.
154. Stalnaker, R., 1968. A theory of conditionals, in: *Studies in Logical Theory*, American Philosophical Quarterly Monograph Series.
155. Starmer, C., Sugden, R., 1989. Violations of the independence axiom in common ratio problems: An experimental test of some competing hypotheses. *Ann. Oper. Res.* 19, 79–102.
156. Stott, H., 2006. Cumulative prospect theory's functional menagerie. *J. Risk Uncertain.* 32, 101–130.
157. Strickland, L., Lewicki, R., Katz, A., 1966. Temporal orientation and perceived control as determinants of risk-taking. *J. Exp. Soc. Psychol.* 2, 143–51.
158. Sudarsanam, S., 2003. *Creating Value from Mergers and Acquisitions*, 2nd ed. Prentice-Hall.
159. Teller, P., 1973. Conditionalization and observation. *Synthese* 26, 218–258.
160. Thaler, R., Johnson, E., 1990. Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Manag. Sci.* 36, 643–660.
161. Thompson, S., 1981. Will it hurt less if I can control it? A complex answer to a simple question. *Psychol. Bull.* 90, 89–101.
162. Trepel, C., Fox, C., Poldrack, R., 2005. Prospect theory on the brain? Toward a cognitive neuroscience of decision under risk. *Cogn. Brain Res.* 23, 34–50.
163. Tsoukiàs, A., 2008. From decision theory to decision aiding methodology. *Eur. J. Oper. Res.* 187, 138–161.
164. Tversky, A., Fox, C., 1995. Weighing risk and uncertainty. *Psychol. Rev.* 102, 269–283.
165. Tversky, A., Kahneman, D., 1981. The framing of decisions and the psychology of choice. *Science* 211, 453.
166. Tversky, A., Kahneman, D., 1986. Rational Choice and the Framing of Decisions. *J. Bus., The Behavioral Foundations of Economic Theory* 59, S251–S278.
167. Tversky, A., Kahneman, D., 1992. Advances in prospect theory: cumulative representation of uncertainty. *J. Risk Uncertain.* 5, 297.
168. Tversky, A., Wakker, P., 1995. Risk Attitudes and Decision Weights. *Econometrica* 63, 1255–1280.

169. Ulkumen, G., Erner, C., Fox, C., Tannenbaum, D., Walters, D., 2014. Perceived Nature of Market Uncertainty Predicts Investment Behaviors. Presented at the Behavioural Decision Research in Management, London.
170. Von Mises, L., 1949. *Human Action: A Treatise on Economics*, 4th ed. Fox & Wilkes, San Fransisco.
171. Von Mises, R., 1957. *Probability, Statistics and Truth*, 2nd ed. Macmillan, New York.
172. Von Neumann, J., Morgenstern, O., 1944. *Theory of Games and Economic Behavior*. Princeton University Press.
173. Von Winterfeldt, D., Edwards, W., 1986a. Cognitive illusions, in: *Decision Analysis and Behavioral Research*. Cambridge University Press, Cambridge.
174. Von Winterfeldt, D., Edwards, W., 1986b. Theory of value and utility measurement, in: *Decision Analysis and Behavioral Research*. Cambridge University Press, Cambridge.
175. Von Winterfeldt, D., Kavet, R., Peck, S., Mohan, M., Hazen, G., 2012. The Value of Environmental Information without Control of Subsequent Decisions. *Risk Anal.* 32, 2113–2132.
176. Vroom, V., Pahl, B., 1971. Relationship between age and risk taking among managers. *J. Appl. Psychol.* 55, 399.
177. Wakker, P., 2010. *Prospect theory: For risk and ambiguity*. Cambridge University Press Cambridge.
178. Wakker, P., Erev, I., Weber, E., 1994. Comonotonic independence: The critical test between classical and rank-dependent utility theories. *J. Risk Uncertain.* 9, 195–230.
179. Wallace, A., 1957. The Homology groups, in: *An Introduction to Algebraic Topology*, International Series of Monographs in Pure and Applied Mathematics. Pergamon Press.
180. Weber, E., Kirsner, B., 1997. Reasons for rank-dependent utility evaluation. *J. Risk Uncertain.* 14, 41–61.
181. Weber, E., Milliman, R., 1997. Perceived Risk Attitudes: Relating Risk Perception to Risky Choice. *Manag. Sci.* 43, 123–144.
182. Weinstein, N., 1980. Unrealistic optimism about future life events. *J. Pers. Soc. Psychol.* 39, 806.
183. Winkler, R., 1996. Uncertainty in probabilistic risk assessment. *Reliab. Eng. Syst. Saf.* 54, 127–132.
184. Wiseman, R., Gomez-Mejia, L., 1998. A Behavioral Agency Model of Managerial Risk Taking. *Acad. Manage. Rev.* 23, 133–153.
185. Wu, G., Gonzalez, R., 1996. Curvature of the Probability Weighting Function. *Manag. Sci.* 42, 1676–1690.

186. Wu, G., Markle, A., 2008. An empirical test of gain-loss separability in prospect theory. *Manag. Sci.* 54, 1322–1335.
187. Young, D., Goodie, A., Hall, D., 2011. Modeling the Impact of Control on the Attractiveness of Risk in a Prospect Theory Framework. *J. Behav. Decis. Mak.* 24, 47–70.

ANNEXURE: Brief Background of some areas of Decision Sciences

In this section some preliminary concepts and methods of Decision Sciences that were assumed in the chapters (but not covered in the chapters) are briefly reviewed. This is mainly for the benefit of a reader who is not versed with the basic concepts of Expected Utility Theory, Cumulative Prospect Theory and principles of Decision Analysis.

7.1 The basic theoretical framework for evaluating decisions

The formal analysis of decisions under uncertainty is based on a three part decomposition of decision problems — actions, uncertain events and consequences. These components are related as follows: an agent chooses amongst uncertain prospects (*actions*) which lead to some *outcomes* (or consequences) that depend on the (uncertain) *states of the world*.

An agent's preference over the possible consequences and beliefs about the uncertainty that underlies these consequences is central to the analysis of decisions. A common feature of various theories that have been put forth to analyze decisions is suitable quantification of both the preferences and beliefs about uncertainty. Tracing its roots to the eighteenth century, *probability* remains a widely accepted metric for measuring uncertainty. With regards to the consequences, the moral worth of the consequences — 'utilities' — are scored between 0 and 1 (Bernoulli, 1738) and can be inferred from preferences among simple binary bets (introduced by von Neumann and Morgenstern, 1944). The recommendation of a suitable action is then evaluated based on an 'Expected Utility' (EU) calculation which combines utilities of the consequences (rewards or losses) with the corresponding probabilities of obtaining them. If p_1, p_2, \dots, p_n denote the probabilities attached to the consequences x_1, x_2, \dots, x_n , $x_i < x_{i+1}$ with utilities $U(x_i)$, the expected utility of a decision $X = \{x_1, p_1; x_2, p_2; \dots; x_n, p_n\}$ is given by:

$$EU(X) = \sum_{i=1}^n p_i U(x_i) \quad (\mathbf{A-1})$$

This expectation procedure is underpinned by certain assumptions (or axioms) about the nature of decision makers' preferences over consequences and beliefs about the uncertainty.¹ The justification for some of the axioms as reasonable assumptions about a 'rational' decision maker corresponds to a form of consistency and is demonstrated by constructing

¹ As well as some technical axioms (e.g. continuity, completeness) which are required to enable the required quantification.

money pump arguments and Dutch books: if a decision maker violates these underlying axioms then he/she can be forced to accept a series of fair bets which will guarantee a net loss whatever happens (e.g. Cubitt and Sugden, 2001; Gilboa et al., 2012; Hájek, 2008). Some of the commonly accepted principles of rationality are transitivity, the sure thing (or independence) principle and that subjective probabilities should obey basic properties of probability calculus (see von Winterfeldt and Edwards, 1986b, pp. 321–25 for a complete list). Any person who accepts these principles of rationality commits to evaluating decisions based on EU calculations (equation A-1), which reflect the (unknown) ordinal preference ranking of the agent over the decisions. A representation theorem is a formal claim of this equivalence between an agent's hidden preference for decisions and the numerical values obtained using equation A-1.


The extent to which quantitative assessments of uncertainty and consequences alone can serve the prediction of the wide range of real world decisions has remained a point of concern. There is ample evidence from the field of Behavioural Decision Making that decisions people make and risks they take depend also on qualitative contextual factors such as the frame of the decision problem (e.g. Kahneman and Tversky, 1979; Tversky and Kahneman, 1981), type of uncertainty (Ellsberg, 1961; Abdellaoui et al., 2011), emotions they experience (e.g. Isen and Patrick, 1983; Loewenstein et al., 2001; Rottenstreich and Hsee, 2001; Seo et al., 2010). In what follows next, the adequacy of EU calculations in terms of modelling and capturing the variety of real world decisions is discussed.

7.2 Modelling the psychology of decision making

Over the years, the Expected Utility (EU) framework has been adapted and applied in various ways. The shape of an agent's utility function reflects not only an agent's preferences for consequences but can also be used to capture his or her attitudes to risk (Arrow, 1971; Pratt, 1964; Arrow, 1951). Attitudes to risk can be inferred from choices people make between a bet containing uncertain payoffs and a sure payoff which is equal to the expected value of the bet (Wakker, 2010, p. 52). Individuals who opt for the sure payoff are regarded as risk averters and the bet rejection can be explained on the basis that the 'subjective' value of the uncertain bet, adjusted for uncertainty, is lower than its 'objective' expected value. Conversely, those who opt for the uncertain bet are regarded as risk seekers (subjective value is greater than objective value). Those who are indifferent conform to the normative assumption of risk neutrality. Attitude to risk is reflected in the curvature of the utility function and measured as the ratio of first and second derivative of the utility function (Pratt, 1964). A concave utility function captures the preference for the sure amount to risk

aversion and a convex utility function implies risk seeking behaviour. This is illustrated in Figure A-1 for a hypothetical choice between a bet which yields £60 with probability 0.5 and £0 otherwise (denoted by $[\text{£}60, 0.5; \text{£}0, 0.5]$) and a sure amount of £30. The subjective value of the bet is given by $U^{-1}(0.5 * U(60))$.

Figure A-1 Illustration of functions capturing risk aversion and risk seeking behaviour for a hypothetical choice between a bet $[\text{£}60, 0.5; \text{£}0, 0.5]$ and sure amount of £30.




The degree of risk aversion, modelled as a parameter of the utility function, can be estimated by eliciting the subjective value of the bet which is the sure payoff that makes the individual indifferent between the bet and sure payoff (e.g. Howard, 1988).² Given a bet $[\text{£}p, p; \text{£}0, 1 - p]$, and elicited sure payoff (CE) the parameter is obtained by solving the equation $CE = U^{-1}(p * U(x))$.

Behavioural scientists have investigated the extent to which individuals conform to the underpinning assumptions of EU theory and how well it predicts actual choices people make and their preference for taking or averting risk. Well known examples of violations of the axioms are Allais Paradox or common consequence effect, the common-ratio effect (see Machina, 1987) and Ellsberg Paradox. Rather than attribute violations of the aforementioned principles of rationality to irrational behaviour, the descriptive theorists endeavour to explain the inconsistencies and construct formal theories or models that rationalize and predict the observed decisions. These theories either relax some of the controversial assumptions of EU theory or introduce a formal correction in the recommendation procedure to reflect the

² Also known as an individual's 'certainty equivalence' for the bet.

cognitive processing of choices. In response to the empirical violations of the axioms, refinements of EU theory, such as non-expected utility (NEU) theories (e.g. Edwards, 1955; Kahneman and Tversky, 1979; Quiggin, 1982; Tversky and Kahneman, 1992), regret theory (Bell, 1982; Fishburn, 1982; Loomes and Sugden, 1986), have been proposed. One of the widely accepted theory of choice under risk and uncertainty which adequately captures the observed choices of individuals, and is used in this research, is Cumulative Prospect Theory (CPT) (Tversky and Kahneman, 1992). Specifically, it predicts a fourfold pattern of attitudes to risk — risk aversion for moderate-to-large probabilities of gains and small probabilities of losses and, risk seeking for moderate-to-large probabilities of gains and small probabilities of losses (Tversky and Kahneman, 1992; Tversky and Fox, 1995). From a procedural perspective, it entails a non-linear transformation of probabilities (akin to the transformations of consequences captured by the utility function). This transformation captures individuals tendencies to overweigh small probabilities of gains and under-weigh large probabilities of gain. This is illustrated in Figure A-2 using some hypothetical weighting functions ($w(p)$).

Figure A-2: Illustration of probability weighting ($w(p)$) typically exhibited by individuals.



It is worth noting that risk acceptance can also be captured by attitudes to uncertainty, i.e. optimism or pessimism (Gonzalez and Wu, 1999; Wakker, 2010, p. 147). To see this, consider the bet [£60, 0.5; £0, 0.5]. An optimistic individual who overweighs the probability of £60 ($w(0.5) > 0.5$) will evaluate the subjective value of the bet to be $w(0.5) * 60$ and hence find the bet more attractive than a sure amount of £30. Conversely, a pessimistic individual who underweighs the probability of £60 ($w(0.5) < 0.5$) will find the bet less attractive than a sure amount of £30. Other aspects of this theory and more details of the functional forms is offered in Appendix 2-1. A comprehensive discussion of Cumulative Prospect Theory can be found in (Wakker, 2010).

Researchers have attempted to obtain a better understanding of the underlying motivations of the psychophysics of the Prospect Theory functions. In particular, the robustness of the ‘risk biases’ predicted by Cumulative Prospect Theory (Kahneman and Tversky, 1979) to the effect of objective contingencies of a gamble (losses/gains frame and magnitude of gain/loss probability), hedonic, emotional, cognitive and other situational factors has been explored extensively. There is evidence that affective components underlie and mediate the Prospect Theory phenomenon (see Trepel et al., 2005 for a review): probability distortion at the boundary (impossibility to possibility, possibility to certainty) is heightened for affect laden outcomes (Camerer, 1992; Rottenstreich and Hsee, 2001), anger induces risk taking whereas fear reduces risk taking (Lerner and Keltner, 2001) and risk taking or aversive biases are reduced if the decision maker experiences pleasant feelings (Seo et al., 2010). Alternative explanations are based on cognitive factors (e.g. Sitkin and Pablo, 1992) and regulatory focus theory (e.g. Higgins, 1996b, 1997, 1998). These theories are briefly discussed in section 2.3.3.


7.3 Decision support tools for modelling decisions in practice

Broadly, the field of Decision Analysis concerns itself with applying the theoretical precepts of Decision Sciences to aid decision making in the real world (Howard, 1988). It provides a formal language for encoding various aspects of a vaguely defined decision problem which helps better understand, analyse and recommend a solution (e.g. Tsoukiàs, 2008). Many of its procedures involve developing models and techniques, such as graphical representation of decision problem, preference and probability elicitation techniques, robustness analysis which are sensitive to the decision makers’ natural thinking. These procedures not only serve as a conduit for translating the decision makers’ beliefs, preferences, reasoning and mental construction to quantitative inputs or relationships in a decision model, but are endowed with the mathematical structure required to perform the corresponding analysis. Some of the widely used tools to represent and model decision problems are decision trees, Bayes nets and influence diagrams. The three part decomposition of decision problems into actions states and outcomes constitute the building blocks for decision trees and influence diagrams and the EU approach for analysing decisions is embedded in the analysis that support these frameworks for representing decisions.

The modelling tool used and discussed in this research is the influence diagram (Howard and Matheson, 1984, 2005). An influence diagram is a graphical representation of a decision problem which describes the relationship between the decisions, uncertainties and consequences. These decision components are represented as nodes in a graph and the

relationship is described by directed arcs. Some illustrations of simple influence diagrams and the relationships they represent are shown in Figure A-3. Figure A-3a shows the prototypical influence diagram for decision problems where payoffs (V) depend on actions (A) and uncertain states (S). Figure A-3b describes decision problems where probabilities attached to uncertain states depend on actions and payoffs depend on the uncertain states; and Figure A-3c illustrates decision problems where both the payoffs and probabilities depend on actions and the payoffs also depend on the uncertain states.

Figure A-3 Influence diagrams for different decision problems.



A thorough technical definition of influence diagrams is presented in sections 3.2 and S 3-3. Influence diagrams offer a convenient and unambiguous way to describe the various relationships between the components of a decision problem and have been found to be widely useful in practice (Howard, 1988; Buede, 2005). In the influence diagram, the notion of control which is the central topic of this thesis is represented by the arc from a decision node to a state node Figure A-3a&b. This arc represents the probabilistic dependence of uncertainties on actions; and the corresponding decisions can be regarded as *interventions* on uncertainty when one of the decision outcomes is a do nothing option.

7.4 References for Annexure

1. Abdellaoui, M., Baillon, A., Placido, L., Wakker, P., 2011. The Rich Domain of Uncertainty: Source Functions and Their Experimental Implementation. *Am. Econ. Rev.* 101, 695–723.
2. Arrow, K., 1951. Alternative Approaches to the Theory of Choice in Risk-Taking Situations. *Econometrica* 19, 404–437.
3. Arrow, K., 1971. *Essays in the Theory of Risk Bearing*, Markham economics series. Markham Pub. Co., Chicago.
4. Bell, D., 1982. Regret in decision making under uncertainty. *Oper. Res.* 30, 961–981.
5. Bernoulli, D., 1738. Exposition of a New Theory on the Measurement of Risk. *Econometrica* 22, 23–36.

6. Bernstein, P., 1998. *Against the Gods: The Remarkable Story of Risk*. John Wiley & Sons, New York.
7. Buede, D., 2005. Influence diagrams: A practitioner's perspective. *Decis. Anal.* 2, 235–237.
8. Camerer, C., 1992. Recent tests of generalizations of expected utility theory, in: *Utility Theories: Measurements and Applications*. Springer, pp. 207–251.
9. Cubitt, R., Sugden, R., 2001. On money pumps. *Games Econ. Behav.* 37, 121–160.
10. Edwards, W., 1955. The prediction of decisions among bets. *J. Exp. Psychol.* 50, 201.
11. Ellsberg, D., 1961. Risk, ambiguity and the Savage axioms. *Q. J. Econ.* 75, 643–669.
12. Fishburn, P., 1982. Nontransitive measurable utility. *J. Math. Psychol.* 26, 31–67.
13. Gilboa, I., Postlewaite, A., Schmeidler, D., 2012. Rationality of belief or: why savage's axioms are neither necessary nor sufficient for rationality. *Synthese* 187, 11–31.
14. Gonzalez, R., Wu, G., 1999. On the Shape of the Probability Weighting Function. *Cognit. Psychol.* 38, 129–166.
15. Hájek, A., 2008. Dutch book arguments. *Oxf. Handb. Ration. Soc. Choice* 173–195.
16. Higgins, E.T., 1996. Ideals, oughts, and regulatory focus. *Psychol. Action Link. Cogn. Motiv. Behav.* 91–114.
17. Higgins, E.T., 1997. Beyond pleasure and pain. *Am. Psychol.* 52, 1280.
18. Higgins, E.T., 1998. Promotion and prevention: Regulatory focus as a motivational principle. *Adv. Exp. Soc. Psychol.* 30, 1–46.
19. Howard, R., 1988. Decision analysis: practice and promise. *Manag. Sci.* 34, 679–695.
20. Howard, R., Matheson, J., 1984. Influence Diagrams, in: R.A. Howard and J.E. Matheson (Eds.), *The Principles and Applications of Decision Analysis*. Strategic Decisions Group, Menlo Park, CA.
21. Howard, R., Matheson, J., 2005. Influence diagrams. *Decis. Anal.* 2, 127–143.
22. Isen, A., Patrick, R., 1983. The effect of positive feelings on risk taking: When the chips are down. *Organ. Behav. Hum. Perform.* 31, 194–202.
23. Kahneman, D., Tversky, A., 1979. Prospect Theory: An analysis of decision under risk. *Econometrica* XLVII, 263–291.
24. Lerner, J., Keltner, D., 2001. Fear, anger, and risk. *J. Pers. Soc. Psychol.* 81, 146.
25. Loewenstein, G., Weber, E., Hsee, C., Welch, N., 2001. Risk as feelings. *Psychol. Bull.* 127, 267–286.
26. Loomes, G., Sugden, R., 1986. Disappointment and dynamic consistency in choice under uncertainty. *Rev. Econ. Stud.* 53, 271–282.
27. Machina, M., 1987. Choice Under Uncertainty: Problems Solved and Unsolved. *Econ. Perspect.* 1, 121–154.

28. Pratt, J., 1964. Risk aversion in the small and in the large. *Econometrica* 32, 122–136.
29. Quiggin, J., 1982. A theory of anticipated utility. *J. Econ. Behav. Organ.* 3, 323–343.
30. Rottenstreich, Y., Hsee, C., 2001. Money, kisses, and electric shocks: On the affective psychology of risk. *Psychol. Sci.* 12, 185–190.
31. Seo, M.-G., Goldfarb, B., Barrett, L., 2010. Affect and the framing effect within individuals over time: Risk taking in a dynamic investment simulation. *Acad. Manage. J.* 53, 411–431.
32. Sitkin, S., Pablo, A., 1992. Reconceptualizing the determinants of risk behavior. *Acad. Manage. Rev.* 17, 9–38.
33. Trepel, C., Fox, C., Poldrack, R., 2005. Prospect theory on the brain? Toward a cognitive neuroscience of decision under risk. *Cogn. Brain Res.* 23, 34–50.
34. Tsoukiàs, A., 2008. From decision theory to decision aiding methodology. *Eur. J. Oper. Res.* 187, 138–161.
35. Tversky, A., Fox, C., 1995. Weighing risk and uncertainty. *Psychol. Rev.* 102, 269–283.
36. Tversky, A., Kahneman, D., 1981. The framing of decisions and the psychology of choice. *Science* 211, 453.
37. Tversky, A., Kahneman, D., 1992. Advances in prospect theory: cumulative representation of uncertainty. *J. Risk Uncertain.* 5, 297.
38. Von Neumann, J., Morgenstern, O., 1944. *Theory of Games and Economic Behavior*. Princeton University Press.
39. Von Winterfeldt, D., Edwards, W., 1986. Theory of value and utility measurement, in: *Decision Analysis and Behavioral Research*. Cambridge University Press, Cambridge.
40. Wakker, P., 2010. *Prospect theory: For risk and ambiguity*. Cambridge University Press Cambridge.