Frontiers of Decision Theory

Siwen (Addison) Pan

School of Economics and Finance (Albany)
Massey University

This dissertation is submitted for the degree of
Doctor of Philosophy in Economics

Massey University August 2016

I would like to dedicate this thesis to:

My parents;

My main Ph.D. Supervisor Dr. Simona Fabrizi; and

Researchers who, with their seminal studies on decision theory,

have made it possible for me to base my own research on theirs.

I also wish to dedicate this thesis to:

All scholars who will find it useful when pursuing their research.

It is for those people that this thesis has been written.

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the text and acknowledgements.

Siwen (Addison) Pan
August 2016

Acknowledgements

I am grateful to Patrick Girard and Gabriele Gratton for their comments on the early stage idea upon which this dissertation is based, as well as to my supervisors, Simona Fabrizi, Thomas Pfeiffer and Matthew Ryan for their continuous support, encouragement and inputs. I also thank both John Hillas and Steffen Lippert for many useful conversations, comments and suggestions, at various stages of this project. I would like to extend my thank you to the participants in the 2nd Massey Business School PhD Symposium (Manawatu Campus), held on 27-28 August 2015; in the 34th Australasian Economic Theory Workshop, held at Monash University on 11-12 February 2016; in the 12th Conference on Logic and the Foundations of Game and Decision Theory, held at Maastricht University on 20-22 July 2016; in the Economics Seminars at Massey University (Albany) in November 2015 and at the Karlsruhe Institute of Technology in May and June 2016. I am also particularly grateful to Simon Grant, for very valuable suggestions on my dissertation work and about possible directions for future research, which he gave as my discussant during the 28th PhD Conference in Economics and Business, held on 11-13 November 2015 at the University of Queensland. Comments from Bernie Grofman, Jean-Francois Laslier and Charles Plott in the 7th Summer Workshop of Centre for Mathematical Social Science, held on 18-20 February 2016 at the University of Auckland are also gratefully acknowledged. I also wish to thank Satosi Nakada, Javier Rivas and Takashi Ui for comments and discussions on chapter 3 of this dissertation during the 5th World Congress of the Game Theory Society, held at Maastricht University on 24-28 July 2016; Michele Lombardi, Marcus Pivato and Naoki Watanabe for helpful comments and suggestions on chapters 1 and 2 of this dissertation during the 12th European Meeting on Game Theory, held at Syddansk University on 11-13 July 2016. I am also grateful for the wonderful hospitality received when visiting the School of Economics at the University of New South Wales in 2014; and when visiting the Institute of Economics at the Karlsruhe Institute of Technology in 2016. Support by the 2015 Massey University Research Fund (MURF) to Fabrizi and Pan for the research assistance by William Zhi and for the running of the experiments is gratefully acknowledged. I would like to extend my gratitude to Steffen Lippert and Valery Pavlov for their inputs and support to first obtain ethical clearance to conduct human-subject laboratory experiments from the Human Ethics Committee at the

University of Auckland; and then to provide access to DECIDE (Laboratory for Business Decision Making), based at the University of Auckland, as well as for helping with the smooth running of the experiments, which were conducted in October 2015, and which are part of this thesis. I am also grateful to my PhD Examiners, Ben Greiner, Clemens Puppe and Oscar Lau for valuable comments and important suggestions for possible future research. Finally, I would like to thank my friends, Edmund Lou and Sherry Wu, for their encouragement and support provided through insightful discussions.

Abstract

The well-known jury paradox – the more demanding the hurdle for conviction is, the more likely it is that a jury will convict an innocent defendant – heavily relies on Bayesian updating. However, with ambiguous information (e.g., a forensic test with accuracy of 60%, or more), standard Bayesian updating becomes invalid, challenging the existence of this paradox. By developing novel theoretical models and by testing their predictions in laboratory settings, this thesis advances our understanding of how individuals process more realistically imprecise measures of information reliability and how this impacts on information aggregation for the group decision-making. Hence, our findings inform the institutional design of collective deliberation, from small to large group decision-making.

Preface

"...two factors [are] commonly used to determine a choice situation, the relative desirability of the possible pay-offs and the relative likelihood of the events affecting them, but in a third dimension of the problem of choice: the nature of one's information concerning the relative likelihood of events. What is at issue might be called the ambiguity of this information, a quality depending on the amount, type, reliability and 'unanimity' of information, and giving rise to one's degree of confidence in an estimate of relative likelihoods."

Daniel Ellsberg, *The Quarterly Journal of Economics*, 1961, 75(4), p. 657-659.

Much real world negotiation and decision-making takes place in small groups. Members within groups are either chosen by public voting or by authorised nomination. They gather together to deliver their opinions or cast votes for determining an authoritative decision. For example, congressmen are chosen by national voting to join Congress to decide whether certain policies are allowed to take place or whether outdated codes need to be abolished. Similarly, corporation board members need to deliberate during seasonal meetings regarding business strategies, innovation projects, and so forth. Also, medical teams, judicial trials (up to the Court of Appeal level), policy offices, and any expert teams are typical sources of authoritative decisions ultimately impacting either treated patients, defendants, suspected offenders, and/or the general public. Hence, small-group deliberations often determine final outcomes (or consequences) mattering to – and affecting to various degrees – a multitude of agents, from single individuals, households, businesses, and organisations, to communities and the entire society.

The underlying mandate of those small groups and their related – implicit or explicit – obligation, when e.g., acting either as expert teams or committees within institutions, is to reach 'the' right/optimal/best decision for 'the' given case at hand. However, a final recommendation/decision by a group of individuals will likely be affected by the processes, the regulations as well as by the voting rules under which their deliberation(s) occurs. In turns, the processes, the regulations, the voting rules, combined with the quality of available information to the decision-partakers, all contribute to whether the outcome of a collective deliberation, beyond being in line with the 'declared' goal set out to be achieved by the deliberation group in the first instance, ultimately best reflects and is consistent with the 'true' nature of the case at hand.

Despite the fact that for any given selected process and voting rule combination it is always possible to characterise/conjecture which decision(s) could be reached by a small decision-group under specified preferences and information structure (whether the information is common knowledge to all parties involved or not), it is far from obvious to anticipate what the possible decision(s) would be in the presence of information which is inherently *ambiguous*.

Ambiguity exists not only in the inability to assign well-defined, numerical probabilities to specific events, as in Ellsberg (1961). It is also embedded in the language, the signals, and even in many among well-recognised social norms used by agents to communicate with one another within any given decision-making context. The language used in the regula-

tions may be interpreted differently by different individuals. The probability of reaching the expected payoff of implementing a new government policy might be difficult to assess. Similarly, decisions whether to grant a patent to an invention by a given patent office, which rely to some extent on the assessment of the quality/novelty of the idea/innovation within a given jurisdiction, could appear to be inconsistent when one patent office grants a patent on an innovation and another denies it. Yet, not all information related to an innovation can be perfectly and unambiguously codified, possibly explaining the apparent contradiction. Another example of when ambiguity matters, is within jury trials. It is not difficult to conceive that jurors often need to form a verdict, based on the evidence submitted to a court of law, the accuracy of which cannot though be assessed perfectly. In other words, jurors are required to cast their votes in favour or against (acquit or convict) a defendant, despite the potential source of ambiguity in the quality of the information provided to them. Thus, even if the very same set of information, ambiguous information, is given to all individuals within agencies/committees/jury in charge of making a decision, that information may still be responsible for generating differing priors among those individuals. And, if more individuals need to agree on the votes they cast, in favour or against, a given choice at hand, the presence of ambiguity could alter the way consensus will be reached, and, potentially, the outcome of such consensus, as opposed to predictions under canonical Bayesian settings. Ambiguity might lead to misunderstanding, sub-optimal choice, and ambiguity-avoiding strategies.

Other Motivating Examples

Below we provide a list of cases, to name but a few examples of other small group decision-making situations where a decision has to be reached for a binary choice under information ambiguity.

The Court of Criminal Appeals This Court responds to defendants who require a review of any adjudications made by the lower court during the original trials. When cases are to be reviewed, lawyers prepare material based on all relevant past cases, including the decisions as they were reached in these cases, and present them to the appeals court, consisting of three or five judges in total. There is no hearing or debating process during the appeal: the judges only read the briefs and the legal documents of the trial court and decide whether to dismiss an appeal. The appeal will be rejected whenever the majority of judges agree with the trial court, and vice versa. Although it does not require an unanimous agreement among the judges, the fact is that different judges might respond differently to similar cases given in the briefs. And, thus, they will hold different opinions regarding the decision of the lower court, which might explain why they eventually fail to reach an agreement even when faced

with the same materials.

The Surgical Team Assume a surgical team, consisting of one chief surgeon and a few attending surgeons, has to come up with a solution about how to treat a cancer patient. The surgeons could either opt to operate on the patient to try to remove the lesion, or conduct chemotherapy and hope the cancer cells will shrink. The surgical team has to reach a decision on the treatment before taking any further actions. Although surgeons are able to come up with the probability of success if surgery is chosen, based on a large database of similar cases, this ability alone does not guarantee that all surgeons will ex ante all agree on the same treatment decision, as not all past cases are exactly the same as the case at hand, due to the uniqueness of the human body, as well as the personal history and idiosyncratic characteristics of the patient[1]. Thus, surgeons may fail to unanimously agree on a particular course of action (decision regarding the treatment) at the state of the surgical consultation, due to different priors/beliefs about which treatment has better odds of success, if undertaken, for this particular patient. Based on those priors, they will most likely have to mediate their positions, to reach an agreement (whether an unanimous one or not) and to be able to treat the patient accordingly.

Organ Allocation Two heart failure patients are waiting for a heart transplant. Whether one of them or the other makes it to the top of the transplant list will determine whether a donor heart will go to the patient who needs it the most and can make the most out of the transplant. When there is a donor heart, the organ allocation center will have to decide to whom they will allocate this heart. Suppose there is a small medical team within the organ allocation center which has to analyse these two patients' cases and vote for who gets first on that list. Whenever an unanimous vote is reached, it decides the receiver of the donor heart. Although there are strict rules for evaluating who should be the receiver, there is still some chance that the two patients' medical conditions are extremely similar, and, thus, there is no obvious way of choosing whom the heart should go to. For example, these two candidate patients for a transplant have the same physical tissue and blood type matching, severity of the disease, recovery potential, etc. It could even be the case that twins are waiting for the heart, and, unfortunately, there is only one suitable donor organ available. Then, it will be

[1]The patients might have different ages or weights; their tumours might be of different sizes, or located in different organs, benign or malignant. Based on past cases, a surgeon clearly knows whether this type of surgery has succeeded or failed in the past. However, each patient treated before is different from the present one. If the surgeon thinks this way, the adequate past cases seem impractical to him and he will end up with nothing plausible to which to resort (Gilboa and Schmeidler, 2010). That is why the patient will be asked to sign an 'informed consent' form to capture his understanding of all the potential risks that might happen in the treatment, including death.

a very hard decision of who gets the heart and who does not.[2] Having to come up with a choice, agreed by the team may be affected by each expert's belief regarding who is more likely to react positively to a transplant, and the final allocation may be affected by possible differing priors about those chances of successful transplant (such decision needs to be a swift one, as the clock is ticking, determining the chances of any transplant to succeed at all).

The Innovation Funding Programme Suppose that the government offers a fund which is only sufficient to promote one innovation project. All major universities have the opportunity to submit their projects to the funding committee. The funding committee has to come up with only one recipient from among all the candidates and their research projects. The judgement standard includes the novelty, promise, feasibility of reproduction, and the potential social contributions of the research project, as well as the project proposer's academic background, publishing record, his/her network of the relevant experts in the field, their co-authors' backgrounds, research reputation, etc. After a few rounds of pre-selections, only two final projects remain in the final round of assessment. In order to minimise the potential dissent in the final decision, the fund will only be given to the project receiving some degree of consensus from the members of the committee. However, members of the committee are likely to each have their own idiosyncratic prior as to the merits of each project, based on their own subjective assessment of its chances of success, say, up to the commercialisation (an innovation may function, technically, but not be successful in the final market, for example, due to how the market receives it – e.g., consumer taste for something really new cannot be anticipated for certain, as there are no other innovations in use which resemble any of those proposed new ones). Obviously, if the committee cannot find one project that all or at least most of its members agree on, the fund will be lost, putting pressure on the committee to find the best possible agreeable allocation of those funds, obeying the idea of the government to promote the most promising innovation project. The intuition is that in this case there is not enough statistical evidence about the distribution of 'good projects' versus 'bad projects' in the economy such that all members of the committee will be able to necessarily all share the same belief about the exact chances of each project submitted to their attention to be of either type. The members' differing priors are likely to impact on the final selection of the recipient of the fund.

[2] It could also happen for the parents of the twins to decide which kid they want to save by agreeing to the heart transplant.

In the above examples, medical surgeons, organ allocation officers, and funding committee members are the decision makers who are faced with some sort of ambiguous information within each small group. Such small group decision-making could also extend to legislatures, expert panels[3] and other judicial bodies. Decision makers are asked to make a choice between possible alternatives. Provided with the very same pieces of information, such as the medical data of the patients, surgical history, project proposal, and the merit of a project, decision makers will then generate their own ideas/opinions or beliefs independently. According to Ellsberg (1961), ambiguity stems from such information, which exposes decision makers to a potential dissent from their initial positions/judgements/beliefs. Decision makers need to be aware of the fact that someone will have to vote against their received information/signal, 'aligning their minds' to eliminate any dissent to reach a decision (whether unanimously or not) and that the quality of the information received matters in determining how such alignment may be reached.

In the remainder of this thesis, we take the jury trial as the leading example, as the metaphor for other small group decision-making examples, to study whether an ambiguous information structure could affect collective deliberation processes, and if so how, in order to gain a better understanding of the effects of different institutions on collective decision outcomes.

To advance our understanding of how ambiguity can play a role in a jury trial setting, we embed identical, but – at least partially – ambiguous information into the canonical jury decision-making model of Feddersen and Pesendorfer (1998). Our main goal is to study the effects of introducing different forms of ambiguity on the probabilities of convicting the innocent (type I error) and acquitting the guilty (type II error), compared to the canonical jury trial case.

To that end, in chapter 1 we begin by exploring a model in which jurors may distrust the precision of the information given to them, leading to jurors adopting potentially differing priors and altering the formation of their posteriors, used when casting votes to convict or

[3]The case of a legislature shares some similarity with the jury trial, in which there exists a default option, in the case a consensus fails to be reached, which is the acquittal. Although legislatures might still be making a binary choice, in general, they are choosing between whether to dismiss a proposal or accept it, a 'Yes' or 'No' question. That is the same for the expert panel, if they are considering whether to adopt a new technology. The main difference between them is that the final choice indicates different results. When an unanimous decision is not accepted, the status quo remains instead (similarly to the jury trial); then, the decision of 'rejecting the legislative proposal' is the same as not voting. A binary choice in these cases is not to choose one option out of two; it becomes whether to maintain the status quo or not. However, in the surgery case, either the patient receives the surgery or he/she will have chemotherapy, neither of which are the status quo.

to acquit a defendant. Within this model, we can summarise the following findings. As the size of the jury grows sufficiently large, when voters share the same 'trusting' level of belief, voting according to their private signals leads to a smaller probability of convicting an innocent defendant. This suggests that if there were ways of framing all voters to believe that the quality of the private information is the highest among alternative ones provided to them, and that belief is wrong (jurors trust the precision to be higher than its 'true' underlying level), type I errors would be reduced, if not even eliminated. Therefore, asymptotically, being trusting of the information received or framing the information to induce more trust in it, makes the unanimity voting rule less unappealing. However, for a small jury size, distrusting the information provided would be best to reduce type I errors and to improve the performance of the unanimity rule.

In chapter 2, we report results from an array of experiments designed to capture the collective voting behaviour under the two-point non-common prior model introduced in chapter 1 and to contrast them against results of canonical collective voting behaviour models. Our aim is to investigate the collective decision-making outcomes under different voting rules when the quality of the private information given to voters when casting their votes is unmeasurable, triggering voters to adopt potentially differing beliefs about it. The results of these experiments validate the theoretical predictions of voting under the two-point non-common prior model, suggesting the importance of the quality of the information structure in determining the collective deliberation outcomes. These results help establish when, in the finite case, the unanimity voting rule can outperform majority voting rule if voters adopts two-point non-common priors.

In chapter 3 we generalise the jury voting model of Feddersen and Pesendorfer (1998) by embedding ambiguity into the private signal structure and considering voters who, being ambiguity averse, adopt a Maxmin approach to form subjective beliefs. The Maxmin Expected Utility Theorem (MMEU) of Gilboa and Schmeidler (1989) helps capture the voter's attitude towards ambiguity to analyse how this impacts the collective voting outcomes under both the majority rule and the unanimity rule. According to MMEU, voters assign their priors in an act-contingent manner, that is, ambiguity averse voters assign the prior, which gives them the best among the worst expected utility levels when evaluating alternatives choices (in this context, voting choices, namely whether to vote to convict or to acquit). Within this framework we prove the existence of an informative voting equilibrium and of strategic voting equilibria. Moreover, we find that if ambiguity exists in the precision of the private information, it is easier to sustain informative voting as an equilibrium strategy, that is, there

exists a larger set of reasonable doubt levels for the unanimity voting rule to prevail as an equilibrium of the voting game. This is an important result as voting informatively, especially under unanimity helps maintain the efficiency of information aggregation.

Our theoretical and experimental results call into question preconceived results about the performance of different institutional designs and voting rules for collective deliberation under differing information structures. When the objective probability of the information is imprecisely measured, that is when the common-prior assumption is relaxed, novel results arise which deserve further exploration, challenging our views about the virtues of adopting, say, majority voting, as opposed to unanimity voting, to avoid the bad outcome of exacerbating the odds of convicting an innocent defendant (jury paradox).

Table of contents

List of figures

List of tables