

Incorporating Preference Information into Formal Models
of Transitive Proxy Voting

MSc Thesis (*Afstudeerscriptie*)

written by

Jacqueline Harding

(born November 29, 1994 in London, United Kingdom)

under the supervision of **Ulle Endriss**, and submitted to the
Examinations Board in partial fulfillment of the requirements for the
degree of

MSc in Logic

at the *Universiteit van Amsterdam*.

Date of the public defense: **Members of the Thesis Committee:**
September 26, 2019

Ulle Endriss
Davide Grossi
Ronald de Haan
Ekaterina Shutova



INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

This thesis is concerned with giving a computational social choice-theoretic model of transitive proxy voting.

Transitive proxy voting (or ‘liquid democracy’) is a novel form of collective decision making. It is often introduced as an attractive hybrid of direct and representative democracy. Recently, it has been used by the German branch of the Pirate Party to aid intra-party decisions (Litvinenko (2012)).

Although the ideas behind liquid democracy have garnered widespread support, there has been little rigorous examination of the arguments offered on its behalf. In particular, there have been relatively few attempts to model liquid democracy formally. A formal model has the potential to serve as a testing ground for the conceptual and empirical claims put forward by supporters (and, of course, detractors) of liquid democracy.

Computational social choice is an emerging field at the intersection of economics and computer science (Brandt et al. (2016)). There are a variety of methodologies and techniques employed within the field, but a common theme in the heterogeneous approaches is a formal perspective on collective decision making. As such, tools from computational social choice seem natural candidates for modelling liquid democracy.

In this thesis, I’ll propose a novel model of transitive proxy voting. My model is individuated by the fact it takes a richer formal perspective on proxy selection (the process by which a voter chooses a proxy). I argue that this allows it better to capture features relevant to claims made about transitive proxy voting.

After proposing the model, I’ll examine it from an axiomatic perspective. I’ll then look at problems of manipulation and control in a proxy vote setting, using the model I have introduced.

Acknowledgements

There is a large number of people I'd like to thank; no inferences can be drawn from absences of names in this section, since such omissions are doubtless unintentional.

With regards to the Master of Logic itself, I'd like to thank Maria Aloni and Tanja Kassenaar. I feel extremely fortunate to have had the opportunity to study such a wide range of subjects in the last two years, and I think much of that breadth of offering is owed to the hard work put in by (amongst many others) these two.

My principal source of funding during the Master of Logic was the Amsterdam Science Talent Scholarship. I'd like to thank all of the donors to this scholarship; without their generosity, I would not have been able to pursue a degree at the UvA. With regards to the administration of the scholarship, I'd also like to thank (amongst others) Marjan Kuiper, Amarins Renema and Gamze Ulker. I received some additional support from Trinity College, Cambridge, in the form of a travelling studentship; thanks to Jeremy Butterfield and Hamish Low for recommending me for this.

With regards to this thesis, I'd like to first and foremost thank my supervisor, Ulle Endriss. This thesis profited not only from Ulle's encyclopedic knowledge of the field of computational social choice but also his attention to detail. It's not hyperbole to say that the model I present here would have been littered with inconsistencies were it not for Ulle. I emerged from each of our meetings with a host of fruitful suggestions and crucial corrections. It's worth adding that Ulle has a quality which it's difficult to overstate the importance of in a supervisor, namely that of being utterly dependable; sometimes I felt like he'd replied to my emails before I'd even sent them.

Secondly, I'd like to thank Ronald de Haan. Without Ronald's course on computational complexity, I think it's likely that I might never have considered doing a thesis with a computational flavour, and the section on parameterized complexity stems directly from a project we conducted alongside the thesis. As a teacher, Ronald's enthusiasm for the subject is infectious; he excels in distilling useful intuitions from difficult concepts.

Thirdly, I'd like to thank my friend and colleague Rachael Colley. From the point where we both decided to write theses on this area, I've found Rachael's suggestions for background reading and potential research avenues

invaluable.

On a more personal note, I'd like to thank (amongst others) Mario Giulianelli, Kristoffer Kalavainen, Florian Mohnert, Mina Young Pedersen and Maike Statz. Their friendship has made the process of writing this thesis enjoyable.

Finally, thanks to my partner, Maria Creech. I'll always treasure the memory of our time in Amsterdam. I owe so much to her support over the last five years, and I'm excited to see what the future holds for us.

Contents

1	Introduction	6
1.1	Overview	6
1.2	Structure of the Thesis	7
1.3	Proxy Voting	8
1.4	Liquid Democracy	9
1.5	Related Work	10
1.5.1	Pairwise Delegations	10
1.5.2	Liquid Democracy with Interdependent Binary Issues	13
1.5.3	Liquid Democracy and the Rationality of Delegation	14
1.5.4	The Structure of the Delegation Graph	15
1.5.5	Preferences over Gurus	16
1.5.6	Preference over Delegates	18
1.5.7	Epistemic Justifications for Liquid Democracy	20
1.5.8	Flexible Representative Democracy	21
2	Proposed Model	23
2.1	Proxy Selection	23
2.2	Formal Background	25
2.3	Social Choice/Welfare Functions	26
2.3.1	Properties of Social Choice Functions	27
2.4	Extending Classical Votes	29
2.4.1	Proxy Mechanisms	29
2.4.2	Proxy Votes	30
2.4.3	Agents' Preferences over the Outcomes of Proxy Votes	31
2.5	Discussion of the Model	32
2.5.1	What is a Proxy Mechanism?	33
2.5.2	Cycles	35
2.5.3	Networks	35
2.5.4	Existing Impossibility Results	37
2.6	Responses and Rejoinders	37
2.6.1	The Model Requires too much from Voters	37
2.6.2	The Model Renders Voter Behaviour Trivial	39
2.7	Concluding Remarks	40

3	Properties of Proxy Mechanisms and Proxy Votes: An Axiomatic Analysis	41
3.1	Proxy Mechanisms	41
3.1.1	Properties of Proxy Mechanisms	41
3.1.2	Characterising SUBSET	43
3.1.3	Discussion	46
3.2	Properties of Proxy Votes	47
3.2.1	Defining Some Properties	47
3.2.2	How do these relate to classical properties of social choice functions?	49
3.2.3	How do these relate to properties of proxy mechanisms?	50
3.2.4	A Proxy Vote Analogue of May's Theorem.	52
3.2.5	Proxy Vote Monotonicity: An Impossibility Result	53
3.2.6	Discussion	59
4	Manipulation and Candidate Control in Proxy Votes	61
4.1	Manipulation	61
4.1.1	Proxy Choice Manipulation	62
4.1.2	Preference Misrepresentation Manipulation	68
4.1.3	Discussion	70
4.2	Control	71
4.2.1	Generalising Candidate-based Control Problems	71
4.2.2	Complexity of Candidate Control in the Proxy Vote Setting	74
4.2.3	Parameterizing Problems with respect to the Number of Gurus: an FPT-Membership Result for PCCDC	74
4.2.4	Discussion	77
5	Conclusion and Future Work	78
6	References	81

Chapter 1

Introduction

1.1 Overview

This thesis is concerned with giving a computational social choice-theoretic model of transitive proxy voting.

Transitive proxy voting (or ‘liquid democracy’) is a novel form of collective decision making. It is often introduced as an attractive hybrid of direct and representative democracy. Recently, it has been used by the German branch of the Pirate Party to aid intra-party decisions (Litvinenko (2012)).

Although the ideas behind liquid democracy have garnered widespread support, there has been little rigorous examination of the arguments offered on its behalf. In particular, there have been relatively few attempts to model liquid democracy formally. A formal model has the potential to serve as a testing ground for the conceptual and empirical claims put forward by supporters (and, of course, detractors) of liquid democracy.

Computational social choice is an emerging field at the intersection of economics and computer science (Brandt et al. (2016)). There are a variety of methodologies and techniques employed within the field, but a common theme in the heterogeneous approaches is a formal perspective on collective decision making. As such, tools from computational social choice seem natural candidates for modelling liquid democracy.

In this thesis, I’ll propose a novel model of transitive proxy voting. My model is individuated by the fact it takes a richer formal perspective on proxy selection (the process by which a voter chooses a proxy). I argue that this allows it better to capture features relevant to claims made about transitive proxy voting.

After proposing the model, I’ll examine it from an axiomatic perspective. I’ll then look at problems of manipulation and control in a proxy vote setting, using the model I have introduced.

1.2 Structure of the Thesis

In this chapter (Chapter 1), I'll introduce transitive proxy voting, or 'liquid democracy'. I'll offer a brief survey of extant models of liquid democracy from within the computational social choice literature.

In Chapter 2, I'll outline the model of transitive proxy voting that I propose in this thesis. What distinguishes my model from existing models is its ability to include preference information in proxy selection. Formally, the model I present augments a classical election/vote (N, A, f) where

- $N = \{i, j, k, l, \dots\}$ is a set of voters, with $|N| = n$
- $A = \{a, b, c, d, \dots\}$ is a set of alternatives (or 'candidates'), with $|A| = m$
- f is a social choice function

with a novel function g , which I call a 'proxy mechanism'. So a 'proxy vote' (or 'proxy election') is a tuple (N, A, f, g) . This proxy mechanism g takes into account preference information supplied by voters and supplies each voter with a set of potential proxies. After sketching some reasons why the addition of a proxy mechanism has the potential to lead to interesting complications, I'll highlight some of the representational capacities of the model I propose.

In Chapter 3, I'll explore novel properties of proxy mechanisms. I'll characterize a natural proxy mechanism (the **SUBSET** mechanism) using some of these properties. We can also define properties of pairs (f, g) , where f is a social choice function and g is a proxy mechanism. I'll examine the interaction between these properties of pairs (f, g) , properties of proxy mechanisms g and classical properties of social choice functions f . I'll prove a proxy vote analogue of May's Theorem and present and prove an impossibility result in a proxy vote setting, showing that certain desirable properties of pairs (f, g) are incompatible with natural properties of their individual components f and g .

In Chapter 4, I'll examine manipulation and control in a proxy vote setting. I'll define a novel form of manipulation (which I call 'proxy choice manipulability') and explore connections between this form of manipulation and manipulation as it is classically understood (which I call 'Gibbard-Satterthwaite manipulability'). I'll also examine the effect on manipulation of domain restrictions (e.g. single peakedness) in a proxy vote setting, demonstrating that strategyproofness is strictly harder to come by in proxy elections. Finally, I'll extend classical candidate control problems into a proxy vote setting. I'll show that when we explore these problems from the perspective of computational complexity, hardness results carry over naturally into this setting. Using tools from parameterized complexity theory,

I'll present a novel choice of parameter for proxy votes, and present *FPT*-membership results using this parameterization. I'll finish by discussing the significance of these results.

I'll conclude the thesis by summarising the principal results I've obtained, and sketching several directions for future work using the model I've proposed.

1.3 Proxy Voting

In a standard vote, voters submit preferences over some set of alternatives.¹ In a proxy vote, voters can choose not to vote directly. Instead, they can delegate their vote to some other voter (who becomes their 'proxy'); this proxy will then cast a ballot on the original voter's behalf, as well as her own ballot.

One way to motivate proxy voting is to frame it as a hybrid between direct democracy and representative democracy. Direct democracy, where citizens vote directly on issues through frequent referenda, is seen as 'strongly democratic but highly impractical' (Green-Armytage (2015), p.190), whilst representative democracy, where citizens elect representatives to make decisions on their behalf, is 'practical but democratic to a lesser degree' (Green-Armytage (2015), p.190). If we view this trade-off between democratic representation and practicality as inherent to any real-world voting process, then it seems like we might want some happy medium, balancing pragmatic factors with the ability for a population to be represented.

Proxy voting purports to be just this. If people want their particular views to be represented in a vote, they can ensure this by voting directly. If they are undecided on an issue (or practical factors prevent them from becoming sufficiently informed, or even from casting their vote directly), they can choose to delegate their vote to someone they perceive as competent, or trustworthy.

As well as this ideological advantage, it has been proposed that proxy voting is accompanied by several practical benefits.

Increasing Voter Turnout. Depending on the situation where it is used, proxy voting may increase voter turnout. There are at least three arguments for this. Firstly, Miller (1969) argues that a major barrier to voters' participation in elections is simply the opportunity cost of voting directly; proxy voting has the potential to lower this cost. Secondly, both Miller (1969) and Alger (2006) identify apathy towards political representatives as a reason

¹Note that at this point I'm not assuming a particular formalisation of elections. For example, I will stay neutral here on what form ballots take. Later, though, I will assume that ballots are represented as linear (or sometimes partial) orders over the set of alternatives.

for poor voter turnout. If a voter can be represented by someone whom she trusts, they argue, she will be more likely to vote. The fact that proxy voting (at least as it is normally construed) allows voters to delegate their votes to any other voter makes it more likely that voters will be represented by someone they approve of. Finally, voters are often deterred from voting by the fact that they haven't made their mind up about all the alternatives being considered in the election (even if they have some sense of what they think). By choosing their proxy carefully, they can vote even if they haven't made their minds up fully.

Increasing Competence of Voters. One case in which a voter might delegate her vote is when she believes another voter to be more competent, or better informed than her. Assuming this perception of competence is truth-tracking, Green-Armytage (2015) argues that this implies that proxy voting leads to votes being cast by voters who are (on average) better informed. Given that at least one justification for increasing representation is epistemic, this is a point in proxy voting's favour.

Increasing Diversity of Viewpoints. Relatedly, Alger (2006) observes that proxy voting might be more likely to lead to a greater diversity in the viewpoints expressed by voters. In a representative democracy, only a very small proportion of the population is a potential representative; this means that such representatives tend to be pushed towards viewpoints with more broad appeal, with the consequence that some voters find their views unrepresented in the views of their representatives. By increasing the number of potential representatives, proxy voting could allow people to express more idiosyncratic viewpoints. Again, this appears to be favoured by an epistemic conception of democracy.

1.4 Liquid Democracy

In the previous section, I outlined proxy voting, and sketched some potential motivations for it. In this section, I will outline a specific form of proxy voting which has earned enough interest to be viewed as a separate form of voting in its own right: 'transitive proxy voting' or (the more catchy) 'liquid democracy'. As the name suggests, what distinguishes transitive proxy voting from a more vanilla form of proxy voting is the transitivity of proxy selection. Blum and Zuber (2016) characterise liquid democracy as the conjunction of four principles. Voters can:

- Directly vote on all policy issues (*direct democratic component*)
- Delegate their votes to a representative to vote on their behalf on (1) a singular policy issue, or (2) all policy issues

in one or more policy areas, or (3) all policy issues in all policy areas (*flexible delegation component*)

- Delegate those votes they have received via delegation to another representative (*meta-delegation component*)
- Terminate the delegation of their votes at any time (*instant recall component*)

(Blum and Zuber (2016), p.165)

It is the meta-delegation component (and, to a lesser extent, the instant recall component) which distinguishes transitive proxy voting from less flexible forms of proxy voting. Proponents of liquid democracy, such as Behrens (2017), argue that this flexibility accentuates the advantages of proxy voting. In the previous section, I emphasised a key strength of proxy voting, namely that it increases the number of potential representatives in an election. Meta-delegation, the ability for proxies to delegate their votes (and the votes that have been delegated to them), can only increase this number. The point is that one can be a representative without having made one's mind up on all the alternatives under consideration. So liquid democracy purports to deliver all of the benefits of proxy voting with some add-ons.

1.5 Related Work

The motivation for this thesis is as follows. In order to assess the arguments in favour of proxy voting (and transitive proxy voting specifically), it's helpful to have a formal model of decision making in a proxy voting setting. A formal model gives us a transparent testing ground for the conceptual and empirical arguments in the previous sections. Since we are dealing with collective decision making, it seems natural to turn to social choice theory when searching for a formal model. Furthermore, it's undeniable that the very notion of transitive proxy voting has an algorithmic whiff to it. With this in mind, in this section I'll examine some existing attempts to model (transitive) proxy voting from within the computational social choice literature.

1.5.1 Pairwise Delegations

Brill and Talmon (2018) propose 'Pairwise Liquid Democracy' (PLD). The key assumption behind PLD, which also operates in the background of the model I propose in this thesis, is that we can view ordinal preference rankings as collections of pairwise comparisons (or 'edges') between alternatives. When voters are asked to provide, for example, linear orders \succ over some set of alternatives A , they are effectively choosing whether $a \succ b$ or $b \succ a$ for each $a, b \in A$ (with the requirement that the pairwise choices they make be transitive and anti-symmetric).

Recall that a factor motivating proxy voting (and transitive proxy voting specifically) is the fact that a voter can fail to have a fully formed opinion and yet still be able to express her partial opinion, through the selection of a suitable proxy. Brill and Talmon’s suggestion is that, in the ordinal preference setting, we can model this voter as having fixed some pairwise comparisons/edges but not others. The model I propose in the next chapter uses this idea too.

There is an important feature of Brill and Talmon’s model which I will not adopt in this thesis. In Brill and Talmon’s model, for each pair of alternatives (a, b) that a voter has not decided between, she chooses some delegate from amongst the other voters to decide on her behalf whether $a \succ b$ or $b \succ a$. Note that this implies that a voter could have a different delegate for each edge she is undecided on. A consequence of this is that the voter can end up submitting an intransitive preference order. Let the set of voters be N . Say voter $i \in N$ is undecided on three pairwise comparisons: (a, b) , (b, c) and (a, c) . She gives the decision between a and b to $j \in N \setminus \{i\}$, the decision between b and c to $k \in N \setminus \{i, j\}$ and the decision between a and c to $l \in N \setminus \{i, j, k\}$. Suppose now that j decides $a \succ b$, k decides $b \succ c$ but l decides $c \succ a$. Then the order i ends up submitting will be intransitive.²

Of course, since social choice/welfare functions operate on total preference profiles (lists of linear orders over the set of alternatives), we cannot use the outputs of these pairwise delegations as inputs to a social choice function. This means that we are left with three options:

- Provide a systematic way of moving from the outputs of delegations (lists of possibly intransitive orders over the set of alternatives) to preference profiles.
- Place restrictions on proxy selections to ensure that every output of a pairwise delegation is a preference profile (i.e. to ensure that intransitivity doesn’t occur at the level of individual voters).
- Modify the social choice function to accommodate intransitivity.

The remainder of Brill and Talmon (2018) is spent exploring these options, particularly the first and third.

With regards to the first option, they find that the most natural response (looking for the minimal number of delegations we can ignore to reach a preference profile), is computationally intractable (NP -complete) to solve.

With regards to the third option, they sketch an initial attempt at a voting rule in their setting, based on minimising the number of pairwise alternative swaps voters have to make to end up with consensus on a ranking

²Note that anti-symmetry will not fail on Brill and Talmon (2018)’s model, since, for each undecided edge (a, b) , a single delegate makes an *exclusive* decision between $a \succ b$ and $b \succ a$.

over alternatives (essentially a liquid democracy version of the Kemeny rule). I see at least two problems with this approach. Firstly, it seems likely that the winner determination problem for any such rule would be *NP*-complete. Secondly, it seems strange to me that in a distance-based approach we would treat the edges a voter has delegated and the edges she has decided herself as having the same weight in the distance calculation. To see this, note that under the rule Brill and Talmon propose we are just as likely to flip the pairwise comparison a voter has made in the order she ends up submitting as we are to flip the pairwise comparison one of her delegates has made. Intuitively, though, it seems like the voter would care much more about the pairwise comparison she herself as made. After all, isn't the purpose of liquid democracy to allow voters who submit partial votes to have their views represented?

I think it's worth considering the second option, namely restricting the delegations which are available to a voter. Brill and Talmon briefly note two ways to do this. Either a voter delegates all possible pairwise comparisons to a proxy, or she collects her pairwise comparisons into a weak order and delegates each indifference class to a proxy. With regards to the latter option, Brill and Talmon argue that it asks too much of voters. Brill and Talmon dismiss the former option as inflexible (after all, they are committed to a pairwise delegation system).

Given that voters will have at most one delegate in my model, I feel it's important that I challenge the idea that it is necessarily advantageous to allow voters to delegate to separate proxies on separate edges. As noted, the flexibility comes at a price (the failure of transitivity). Rather than dwell on the practical (and computational) issues with fixing intransitivity, though, I think there is also something conceptually suspect about allowing intransitive delegations. According to Brill and Talmon, delegation is done on the basis of the perception of competence; voter i delegates the decision (a, b) to voter j because she thinks j is more competent at making the decision than her. Similarly, she delegates the decision (b, c) to voter k because voter k is competent on this issue, and (a, c) to voter l because l is competent on this issue. But if we accept that it is irrational to hold intransitive preferences oneself, then it is not at all obvious to me that it is rational to accept an intransitive preference resulting from delegation. Surely the conclusion voter i ought to draw when her delegates present her with the cycle $a \succ b \succ c \succ a$ is that she was mistaken in her initial assessment of the competence of her delegates? If we understand competence in the common sense way, in terms of a propensity to make correct decisions, then (assuming that intransitivity is the incorrect decision) at least one of her delegates must be incompetent. If we think that voters ought to delegate to competent people, then it appears that we are condoning irrationality at a distance.

I take it that any response to this point needs to give a proper account

of the reasons why a voter chooses a delegate. My aim is not to deny that this can be done, but rather to cast doubt on the idea that the flexibility provided by pairwise (or issue-wise) delegation is inherently advantageous. In effect, my point is about model selection. *Ceteris paribus*, I don't think there's any reason to prefer a model which allows delegations which result in intransitive orderings, given that we require that preferences be transitive. Intransitivity shouldn't be thought of as merely a practical problem, but rather a philosophical one.

1.5.2 Liquid Democracy with Interdependent Binary Issues

The model proposed by Christoff and Grossi (2017) uses interdependent binary issues instead of ordinal preference rankings. Let A be the set of binary issues under consideration. For each binary issue $p \in A$, voters either submit their decision on the issue or choose some delegate to decide whether $p = 0$ or $p = 1$.

The parallels between this model and that of Brill and Talmon (2018) should be clear. Once we translate ordinal rankings into a binary aggregation setting, pairwise comparisons become binary issues and transitivity is just one possible rationality requirement on individual judgements.

With this in mind, it's clear that (a more general variant of) the same problem arises for Christoff and Grossi's model as arose for Brill and Talmon's model, namely that issue-wise delegations can result in a voter having a judgement which breaches some rationality requirement inherent to the aggregation setting. Christoff and Grossi suggest a novel solution to this problem, namely to think of delegation as a diachronic process. From this perspective, which they call the 'vote-copying' perspective, voters begin with some default view on issues (some default judgement). At each timestep, they delegate decisions on some issues to individual proxies. Rather than updating their judgements with the decisions made by their delegates, they first check to see whether such an update would be consistent with any rationality requirement in the aggregation setting. Only updates which are consistent with any rationality requirements are performed, and the process continues until convergence. (Christoff and Grossi characterise the conditions required for convergence.)

Again, I think any breach of individual rationality resulting from issue-wise delegation constitutes a conceptual reason to be suspicious of issue-wise delegation, for the same reasons as outlined in my discussion of pairwise delegation (in Brill and Talmon (2018)). I think an account of proxy selection is also needed to justify the 'vote-copying' interpretation of liquid democracy. For example, why should we assume that all votes are copied simultaneously?

Christoff and Grossi (2017) considers another problem, namely the fact that delegations can result in delegation cycles. If i delegates on issue p to

j , j delegates issue p to k and k delegates issue p to i , then it is unclear what the input to an aggregator should be. In other words, how should we interpret i , j and k 's views on p ?

Standardly, agents in a delegation cycle are simply assumed to abstain (here, to abstain on a single issue, since cycles occur at the level of individual issues). This seems unsatisfactory, particularly since a practical justification for proxy voting appeals to its ability to raise voter turnout. Christoff and Grossi propose that voters submit a default view on each issue. In the case that some set of agents is involved in a delegation cycle on an issue, we look at the majority default view on the issue among the agents involved in the cycle, and interpret each agent in the cycle as having that majority view at the point of aggregation. This seems to capture the idea of agents being represented by their delegates without penalizing them for being involved in a delegation cycle. Of course, it comes at a cost; the requirement that voters submit a default view goes against a motivation for proxy voting, namely that one can express one's opinion without putting in the effort to generate a fully formed vote. That said, it's hard to see how else we should deal with cases where voters don't have suitable options for delegates (and I will argue that cycles constitute only one such case). The model I propose in the next chapter will follow Christoff and Grossi in having agents submit default views on issues.

1.5.3 Liquid Democracy and the Rationality of Delegation

One complaint I have made about the models I've discussed so far has been that little emphasis has been placed on analysing the reasons why an individual voter chooses a proxy rather than casting her vote directly. There is often a tacit assumption that there is some effort involved in voting directly, or that the voter thinks there will be a gain of accuracy by delegating to a more competent representative, but the assumption is undeveloped, to the extent that it has no counterpart in the actual formalism of the models.

Bloembergen, Grossi, and Lackner (2019) attempt to fill this lacuna by focusing on the conditions according to which it is rational for an agent to delegate her vote rather than voting directly.

The model they consider is very simple; voters are choosing between two alternatives. For each agent, one alternative is better (i.e. agents can be divided into two groups depending on which alternative is better for them), but agents are not aware which alternative is better for them (or which alternative is best for the other agents). If an agent votes directly, then, there is a chance she will vote for the alternative which is worse for her. Bloembergen, Grossi, and Lackner assume that the probability q_i that an agent i votes with her interests when she votes directly is always $q_i \geq 0.5$. Bloembergen, Grossi, and Lackner call this probability q_i an agent's 'accuracy'; each agent's accuracy is assumed to be public knowledge.

Agents are arranged into a network structure. Each agent can choose to vote directly or to delegate her vote to one of her neighbours in the network (so delegations are transitive). If an agent i votes directly, she incurs a cost e_i (interpreted as the effort it took her to vote). If she delegates her vote, she incurs no such cost. The utility an agent gains from voting is proportional to the probability that the voter who ends up casting her vote - her ‘guru’ (note that if she votes directly she will be this voter) - votes for the alternative which is better for her.

Under restricted conditions (firstly, where agents have deterministic hidden interests, rather than some non-degenerate distribution over hidden interests; secondly, where each cost $e_i = 0$), Bloembergen, Grossi, and Lackner show the existence of Nash Equilibria in their model.

I don’t want to examine their results in detail, but I do want to highlight two features of their model which I will employ in mine.

Firstly, I will also use the idea that agents incur some cost to voting directly. It seems to me that it’s important to make explicit such a cost in any formal model of proxy voting, since it features so prominently in philosophical justifications of proxy voting. In my model, rather than a cardinal cost, this will be a means for agents to decide between outcomes they are indifferent between (they will prefer reaching the same outcome having invested less energy in voting; this will be made more formal in the next chapter).

Secondly, I will also use the core idea of Bloembergen, Grossi, and Lackner’s model, namely that what drives an agent to pick a particular delegate as her proxy depends in some way on that delegate’s views on the alternatives at hand. I think it’s crucial to emphasise that proxy selection must depend on some feature of the proxy being selected. Because I will (usually) be dealing with settings with more than two alternatives, this will need to be fleshed out in a different way from the accuracy metric of Bloembergen, Grossi, and Lackner (2019).

1.5.4 The Structure of the Delegation Graph

The models we have discussed so far typically divide transitive proxy voting into two stages. In the first stage, preferences over alternatives and delegates are elicited; these are then combined into a ‘delegation graph’ (a graph representing delegations between voters). In the second stage, this delegation graph (or the profile resulting from it) is used as an input to some sort of aggregator.

It’s possible to consider questions regarding the two stages independently. Gözl et al. (2018) focus on the formation of the delegation graph from information about voters’ delegation preferences.

Recall that a purported advantage of transitive proxy voting is that it concentrates power amongst more competent voters. Some have argued

that this concentration of power in the hands of a few ‘super-delegates’ is problematic.³ It risks putting electoral outcomes at the whims of a few individuals. To combat this, Gözl et al. consider the problem of assigning delegations so as to minimise the number of voters delegating their vote to a single proxy (in other words, to minimise the maximum voting weight of a voter who votes directly).

In the model they propose, each agent specifies a subset of the other voters who they would be happy to delegate their vote to. A central mechanism then forms the delegation graph from this information, attempting to restrict the concentration of voting power by minimising the maximum weight given to any delegate. They find that not only is solving this problem *NP*-hard, but also that approximations to the problem are *NP*-hard (even in the very restricted case where we assume that each voter picks at most two potential proxies).

Boldi et al. (2011) examine a similar issue from a different perspective. In their ‘viscous democracy’, they propose a delegation factor $\alpha \in (0, 1)$, representing the extent to which delegation preserves voting power (which influences the ‘viscosity’ of the system). The smaller α is, the more weight is lost every time a vote is transferred. Intuitively, then, fine-tuning α could affect the feasible length of delegation chains. They discuss the impact of the structure of an underlying social network on the number of possible winners.

1.5.5 Preferences over Gurus

In Escoffier, Gilbert, and Pass-Lanneau (2018, 2019), the authors investigate the stability of delegations. Their model is as follows. Let $N = \{i, j, k, l, \dots\}$ be a set of voters, with $|N| = n$. Each voter is choosing whether to cast her own vote, choose some other voter to be her proxy or abstain. Escoffier, Gilbert, and Pass-Lanneau (2018) arrange the voters in a social network (with the accompanying restriction that voters are only allowed to delegate to their neighbours), whilst Escoffier, Gilbert, and Pass-Lanneau (2018) assume the social network is complete (such that there is no restriction on who can delegate to whom).

We assume that each $i \in N$ has a preference ordering \succ_i over $N \cup \{0\}$; we interpret this as representing who i would prefer to end up casting her vote,⁴ with ‘0’ representing the possibility of abstention. For example, $i \succ_i j$ implies that i would rather cast her own vote than have j end up casting her vote, whilst $j \succ_i 0$ implies that i would rather have j end up casting her

³For example, Kling et al. (2015) conduct an empirical analysis of internal election data from Germany’s Pirate Party (which used a transitive proxy voting system), showing that power ended up concentrating amongst the most active users of the system.

⁴Note that this is not a preference relation over who is her immediate proxy, but rather a preference relation over who is her terminal proxy (or ‘guru’).

vote than she would abstain.

The authors fix a delegation function $d : N \rightarrow N \cup \{0\}$. For some $i \in N$, $d(i) = j$ signifies that i delegates her vote to j , $d(i) = i$ signifies that she casts her own vote and $d(i) = 0$ signifies that she abstains. A delegation function is ‘Nash-stable for agent i ’ if i prefers her guru (strictly speaking, she is not guaranteed to have a guru, since she could end up abstaining) when she delegates according to $d(i)$ over any guru that results from some other feasible delegation. A delegation function is ‘Nash-stable’ if it is Nash-stable for every $i \in N$.

It is easy to see that there are delegation functions which are Nash-stable (simply consider the case where every voter prefers to cast her own vote) and delegation functions which are not (consider a voter who hates abstaining who delegates to a voter who abstains). For a given list of preferences $(\succ_i)_{i \in N}$ and social network (N, R) (where R is a binary relation), Escoffier, Gilbert, and Pass-Lanneau (2018) investigate (amongst other optimisation problems) the problem of finding whether a Nash-stable delegation function exists. They show that the problem is NP -hard even when we assume that R is complete, or has a bounded out-degree. They also show that the problem is $W[1]$ -hard when we parameterize by the tree-width of the social network.

Escoffier, Gilbert, and Pass-Lanneau (2018) examine the effect of restricting the allowed list of preferences $(\succ_i)_{i \in N}$ on the optimisation problems considered in Escoffier, Gilbert, and Pass-Lanneau (2018) (assuming here that R is complete). Specifically, they consider the effect of requiring that $(\succ_i)_{i \in N}$ is single peaked along some dimension. They show that when $(\succ_i)_{i \in N}$ is single-peaked, there will always exist a Nash-stable delegation function, which can be found in polynomial time.

The idea of having voters rank other voters is one I will use in my model. Rather than have voters express preferences over gurus, though, I will have voters express preferences over immediate proxies.

What is missing from the models in Escoffier, Gilbert, and Pass-Lanneau (2018) and Escoffier, Gilbert, and Pass-Lanneau (2018) is the actual election in which voters are participating. That is, what is driving the preferences expressed by the voters? In lieu of such an account, it’s unclear what the consequences of their model are for transitive proxy voting. To make this more concrete: unless we have some formal account of what generates preferences over gurus (for example, a notion of competence built into the model, or an idea of a guru agreeing on a particular issue), it’s not clear to me that we should rush to accommodate such preferences. If a guru has been chosen by a proxy of a voter and the voter is dissatisfied with the guru, then does that not imply that the voter should be dissatisfied with her choice of proxy? Part of the problem relates to the interpretation of delegation functions; do they represent (more plausibly) the voter’s preference over her delegates, or do they represent some strategic attempt to end up with as competent a guru as possible? Similarly, it’s unclear how we should inter-

pret the importance of single-peakedness in this setting; can it be thought of as second-order agreement on voter competence? Some philosophical work is needed to tease out the significance of the formal results.

1.5.6 Preference over Delegates

In the previous section, I discussed a model where voters submitted preferences over their potential terminal delegates, or gurus. As I noted, it also seems natural to consider a model where agents can submit preference information over their immediate proxies (i.e. the people they can delegate to directly), since these are the delegations that agents can control themselves. Kotsialou and Riley (2018) propose a model in which there is a set of agents N and a set of alternatives A . Agents can either submit:

- (partial) preferences over the set of alternatives A
- (partial) preferences over the set of voters N
- no information

In the first case, voters are taken to cast their own vote. In the second case, voters are taken to delegate their vote (with the delegate still to be decided on by a central mechanism). In the third case, voters are taken to abstain from voting.

Given that each agent has been sorted into one of these three categories (a direct voter, a delegator or an abstainer), a ‘delegation graph function’ then produces a directed graph (N, R, w) , where

- N is the set of agents
- R is a binary relation over N . We have $(i, j) \in R$ iff j features in the partial ranking i submits over the set of voters.
- Suppose there is an edge $(i, j) \in R$. Then $w(i, j)$ labels the edge with the position that j features in the preference ranking that i submitted (we know that j is ranked at some point in the preference ranking, since we are assuming the edge (i, j) exists in the graph).

A ‘delegation rule function’ takes into account this graph and the partial list of partial preference information over alternatives submitted by the voters, and produces a single delegation (or abstention) for each voter, resulting in a preference profile. This can then be fed into one’s preferred social choice (welfare) function.⁵ Kotsialou and Riley define two different types of

⁵There’s a small technical problem here. Since Kotsialou and Riley allow that voters can submit partial preferences over the set of alternatives, and require that any voter who submits preferences over alternatives casts her vote directly, they need to give some details about how these partial preferences should be extended to linear orders. For now, I will just assume that every voter who submits preference information over the set of alternatives submits complete information.

delegation rule function.

Firstly, they define a ‘depth-first’ function, which they observe is the standard interpretation of transitive proxy voting. In a depth-first function, to find a guru for voter $i \in N$, we move along edges with weight 1 (i.e. we move to i ’s favourite proxy, then to her favourite proxy’s favourite proxy, etc) until we reach a voter who votes directly. This voter is then taken to be her guru.

Secondly, they define a novel delegation rule function, a ‘breadth-first’ function. In a breadth-first function, we look for the shortest path from i to a voter who votes directly. If there are multiple shortest paths, we pick the path with the smallest weight.

Kotsialou and Riley show that when we use a depth-first method of delegation, there are profiles where voters would prefer not to be the guru of some other voter (so a proxy vote analogue of participation fails). When we use a breadth-first delegation rule, though, they show that this situation does not arise.

The model of Kotsialou and Riley (2018) attempts to give a fuller picture of the liquid democratic process (in that it connects delegation to a standard vote, since voters can submit preference information over alternatives). As noted in a previous section, I will also have voters submit preference information over potential proxies.

One issue with the Kotsialou and Riley (2018) model is that the two sorts of preference information submitted by voters (over alternatives and over proxies) are treated independently by the central delegation mechanism(s). That is, there is no attempt to capture the intuitive idea that voters might prefer proxies who agree with their views. Voters are immediately categorised into direct voters or delegators, regardless of the actual content of the preferences they submit (the existence of a preference order of either sort is sufficient to determine this categorisation). Recall that an advantage of (transitive) proxy voting is that it allows voters to express preferences on some issues but not others. The model by Kotsialou and Riley is unable to represent this idea, since it takes an all-or-nothing approach to delegation.⁶

With this in mind, I think we should also question the decisiveness of the participation property the authors focus on (‘guru’ participation, where a voter would always like to end up casting some other voter’s vote). Suppose we accept that some delegates are able better to represent a voter’s views than others, and that the voter’s preferences over delegates tracks this property. Then there is another natural participation property we would want satisfied, namely that a voter would rather delegate her vote than abstain.

⁶One solution to this worry might be to incorporate ideas from the ‘statement voting’ of Zhang and Zhou (2017), which allows delegations with (e.g.) conditional structure. However, it’s not clear how to marry these sorts of ballots with Kotsialou and Riley’s model.

As it stands, the model from Kotsialou and Riley is unable to accommodate this, since any voter who has preferences over alternatives casts her vote directly (so every delegator has no preferences over alternatives, implying that she is entirely neutral between participating and abstaining). But say we could extend the model to incorporate this idea. Then it seems unlikely that a breadth-first delegation function will satisfy this sort of participation property, since it can lead to delegations that a delegator is very unhappy with. In other words, breadth-first delegation may make gurus happy, but it seems unlikely to make delegators happy, once we augment the model with a representation of delegator satisfaction.

1.5.7 Epistemic Justifications for Liquid Democracy

Thus far I've considered accounts of liquid democracy which consider transitive proxy voting as a system of aggregating voters' preferences. But there are also models which examine the claim that transitive proxy voting tends to lead to better outcomes (where 'better' is understood in an epistemic sense, as 'more correct').

Cohensius et al. (2017) considers a situation where a (possibly infinite) population N of voters is distributed on some interval $[a, b]$. They consider two scenarios. In the first, the ground truth is taken to be the median of the voters' positions. In the second, it is taken to be the mean of their positions.

The basic set-up for both scenarios is the same. Some distinguished finite $N' \subseteq N$ (with $|N'| = n'$) is selected, representing the set of voters who are allowed to cast their votes directly (the 'agents' in Cohensius et al.'s terminology). A vote consists of (possibly falsely) stating one's position on the interval. Cohensius et al. compare the situation where the other voters - the non-agents in $N \setminus N'$ - abstain to the situation where the other voters delegate their votes to the members of N' . In the model of Cohensius et al. (2017), each voter delegates her vote to the closest agent in N' who chooses to cast her vote directly. Note that this means that delegations will never be transitive (so we are in a vanilla proxy voting situation). The authors find that proxy voting is always more accurate when the ground truth is the median position, and generally (through simulation) more accurate when the ground truth is the mean position.

Kang, Mackenzie, and Procaccia (2018) answer a similar question using a more familiar model. They assume that N voters are arranged in a social network. They are voting on a single binary issue, for which it is assumed there is a ground truth. Similarly to Bloembergen, Grossi, and Lackner (2019), each voter $i \in N$ has a competence level p_i , interpreted as the probability she would vote correctly if she voted directly. The competence level of each voter is public information.

Voters can either vote directly or delegate their vote to a neighbour whose competence level is strictly higher than their own (note that this eliminates

delegation cycles).

Kang, Mackenzie, and Procaccia define a ‘delegation mechanism’ as a function which takes in the social network and list of competence levels and returns, for each voter $i \in N$, a probability distribution over the delegations available to i . Delegations carry weight according to this probability distribution. The collective decision is then made by the majority rule.

The authors focus on a special class of delegation mechanisms, which they call ‘local’ mechanisms. For a voter $i \in N$, a local delegation mechanism is blind to any information outside of i ’s neighbourhood in the social network (so in two distinct networks where i has exactly the same neighbours with the same competence levels, the delegation mechanism will output exactly the same probability distribution over those neighbours).

Ideally, we would like a local mechanism whose expected return is always better than the case where everyone votes directly. Unfortunately, Kang, Mackenzie, and Procaccia prove an impossibility result, namely that there is no local mechanism which is always no worse than voting directly and better in some cases as the number of voters increases. Their results essentially work by concentrating voting power in the hands of a few voters.

1.5.8 Flexible Representative Democracy

Abramowitz and Mattei (2019) propose ‘Flexible Representative Democracy’ (FDR). FDR is a model of vanilla proxy voting, since it doesn’t allow proxies to further delegate the votes they have been given. However, it is worth outlining FDR, since it informs a key motivation for the model I will present, namely that delegates ought to represent the views of the voters who delegate their votes to them.

In FDR, there is an electorate that is voting on a set of binary issues. Abramowitz and Mattei divide the electorate into two distinct categories, voters and delegates. This risks undermining an argument for the claim that proxy voting increases voter turnout, namely that voters are able to delegate their votes to individuals who they know and trust.

In Abramowitz and Mattei’s model, both voters and delegates have preferences over the set of binary issues. Voter preferences are private, but delegate preferences are public (so every voter knows every delegate’s preferences and her own preferences, but no other voter’s preferences; we can assume some sort of election campaign has occurred). Voters then express preferences over delegates (Abramowitz and Mattei consider various ballots on which these preferences could be expressed, such as approval voting or standard ordinal preference voting). Crucially, a voter’s attitude towards a delegate is assumed to correspond with the degree of agreement between the voter’s preferences over the issues and the delegate’s preferences over the issues. So voters make delegation decisions according to how closely a delegate’s preferences match their own (with tie breaks being broken arbi-

trarily).

There are several important differences between Abramowitz and Mattei's model and mine. The fundamental setting is entirely different, to the extent that the primary question they consider (which set of delegates should we choose to cast votes, assuming we want to represent the views of the population?) wouldn't make sense in my setting. But the core idea - that a voter's choice of delegate should be in some way tied to correspondence between the voter's views and the delegate's - is at the very heart of the model I present. As I've emphasised throughout this chapter, it's an idea which is largely absent from the literature on transitive proxy voting.

Chapter 2

Proposed Model

In the previous chapter, I introduced transitive proxy voting and examined existing attempts to formalise it from within the computational social choice literature (broadly construed). During my discussion, I argued that we ought to assess these models by their ability to represent the features of transitive proxy voting that are central to its justification as a system of voting. Amongst others, these included:

- The ability for voters to express their views on some issues but not others, by choosing a suitable proxy.
- The empirical claim that voters will delegate to a proxy who they perceive as more competent.

I emphasised that very few of the existing models give an account of what it is for a voter to select a proxy. The result is that proxy selection is often treated by these models as a black box. This means that the models are often unable to represent the very features which make proxy voting attractive, since they depend on the notion of a proxy representing a voter. This is unfortunate, since a purpose of formalisation is giving us a more rigorous framework in which to test philosophical arguments regarding transitive proxy voting.

The model I present in this chapter is motivated by this deficiency. I start with the platitude that voters choose proxies to represent them, and attempt to formalise this intuition, whilst allowing sufficient flexibility to represent different features relevant to liquid democracy.

2.1 Proxy Selection

It's not my aim in this thesis to give a full account of the factors that go into a voter's choice of delegate. Indeed, it seems likely that such an account would have to contain so many features as to render it unsuitable for the sort

of abstraction required for a social choice-theoretic model. But I do think that a model of transitive proxy voting should try to take proxy selection into account. With this in mind, I begin with the following two intuitive principles.

- (1) There is a large range of factors which informs a voter's choice of delegate (for example, a perception of competence, charisma, intelligence, honesty, etc).
- (2) Voters pick proxies who (they think) will represent their interests.

I take it that (1) and (2) are plausible starting points for an account of proxy selection.

It is helpful to illustrate the relationship between (1) and (2) with an example. Suppose I were asked to give my preference regarding the nature of Britain's future relationship with the European Union, in the form of an ordinal ranking over the available options. There are a multitude of options at hand, including:

- (a) remaining in the EU
- (b) leaving the EU without a deal
- (c) leaving the EU with a customs union
- (d) leaving the EU with a backstop

and a variety of others, with varying degrees of specificity and complexity. Suppose (accurately!) that I am not sufficiently well informed about these options to submit a linear order over them. I know that I prefer remaining in the EU to the other options, but I am unsure about how to compare the various intermediate levels of integration at hand. If I am given the option of choosing a trusted delegate to submit an opinion on my behalf, I will opt for this option.

Suppose that my friend Alice is exceptionally well informed about the intricacies of the EU. She is a lawyer specialising in European law and regularly meets with industry experts on Brexit-related matters. *Ceteris paribus*, then, she would be an excellent candidate to be my delegate. She manifests various qualities which are relevant to my choice of delegate.

Suppose now that I learn that Alice prefers leaving the EU without a deal to remaining in the EU (so $b \succ a$, according to Alice). Recall that I am sure that I prefer remaining in the EU to leaving without a deal (so I think that $a \succ b$). The fact that Alice prefers a no-deal Brexit to remaining in the EU doesn't make me think that she's any less informed, or trustworthy, and so on, but it is sufficient to ensure that I won't pick her as my proxy. Since she disagrees with me so strongly on the issues on which I have made

up my mind, I don't think she will represent my interests if she votes on my behalf.

I use this thought experiment to show that a model of proxy selection should have at least two interacting components. First and foremost, voters will only consider delegates who represent their interests (this is (2)). That said, it is futile to attempt to place restrictions on their choice amongst potential delegates who they feel represent their interests, since many factors are relevant to this decision (this is (1)).

Of course, when we formalise this idea, we will need to flesh out what it means for a delegate to 'represent a voter's interests'. We will also need to formalise the notion of a voter choosing a delegate without making explicit the criteria behind the choice.

In the setting I use, a set of voters N will submit ordinal preferences over both a set of alternatives A and the other voters (i.e. over their potential proxies). The former preference can be partial, meaning it can omit certain pairwise comparisons between alternatives. The latter preference must be a linear (total) order. Based on the preferences submitted over alternatives, a central mechanism decides for a given voter which of the other voters represent her interests sufficiently. This subset of voters is called the voter's 'permitted proxies'. Based on the preferences she submitted over the other voters, one of these permitted proxies is then chosen as her delegate.

2.2 Formal Background

For a finite set X , let $\mathcal{P}(X)$ denote the set of all binary relations on X which are irreflexive, anti-symmetric and transitive.

I will call $P \in \mathcal{P}(X)$ a 'partial order', to emphasise that P need not be total. Technically, of course, the relation usually called a 'partial order' is reflexive rather than irreflexive. The reader should be mindful of this terminological idiosyncrasy, but it makes no substantive difference to the content of the thesis.

Following Brill and Talmon (2018), it will be helpful to think of a partial order as a set of strict pairwise comparisons. This affects the notation I use. Suppose $X = \{a, b, c\}$. Then, using my terminology, the following are all examples of partial orders on X :

- $P = \emptyset$
- $P' = \{a \succ b\}$
- $Q = \{a \succ b, a \succ c\}$

but the following would not be a partial order, since it is not closed under transitivity:

- $Q' = \{a \succ b, b \succ c\}$

I will also speak of specific pairwise comparisons (or ‘edges’) being members of a partial order. For example, I will say that P' contains the edge $a \succ b$. Formally, I will write that $a \succ b \in P'$ (or, equivalently, that $\{a \succ b\} \subseteq P'$), but $a \succ b \notin P$ ($\{a \succ b\} \not\subseteq P$). I will also write $|P|$ to express the number of pairwise comparisons P contains. For example, $|P| = 0$, $|P'| = 1$ and $|Q| = 2$.

Let $\mathcal{L}(X)$ denote the set of all binary relations on X which are irreflexive, anti-symmetric, transitive and also complete. Then I call $L \in \mathcal{L}(X)$ a ‘linear order’. Note that, by definition, $\mathcal{L}(X) \subseteq \mathcal{P}(X)$.

Throughout the thesis, I will speak of profiles of partial (linear) orders. We can think of a profile of partial orders as a list of partial orders, one for each agent. So if $N = \{1, \dots, n\}$ is the set of agents, and A is the set of alternatives, then $\mathbf{P} = (P_1, \dots, P_n) \in \mathcal{P}(A)^n$ is a list of partial orders (note I use the bold type face for the list of orders, and the normal type face for the partial orders themselves). By P_i , I designate the partial order submitted by agent i .

Fix some $\mathbf{P} = (P_1, \dots, P_n)$. Then, as is standard, we can also write $\mathbf{P} = (P_i, P_{-i})$ or $\mathbf{P} = (P_{i,j}, P_{-i,j})$, for some $i, j \in N$. I will write (P'_i, P_{-i}) to designate the profile that is an ‘ i -variant’ of (P_i, P_{-i}) (that is, the profile where at most agent i changes the order she submits, from P_i to P'_i). The same notational conventions apply to profiles of linear orders.

2.3 Social Choice/Welfare Functions

Let (N, A) be defined as follows:

- $N = \{i, j, k, l, \dots\}$ is a set of voters, with $|N| = n$. It will also sometimes be convenient to write $N = \{1, \dots, n\}$.
- $A = \{a, b, c, d, \dots\}$ is a set of candidates, with $|A| = m$.

Recall that $\mathcal{L}(A)$ is the set of all linear orders over A . Note that $\mathcal{P}(A)$ denotes the powerset of A ; this should not be confused with $\mathcal{P}(A)$, the set of partial orders over A .

There are two types of social choice functions (W. Zwicker and Herve Moulin (2016)).¹

Definition 2.1 (Irresolute Social Choice Function). An Irresolute Social Choice Function

$$f : \mathcal{L}(A)^n \rightarrow \mathcal{P}(A) \setminus \emptyset$$

aggregates agents’ total preferences over A into a set of winners of the election.

¹Of course, it is possible to see resoluteness as a property of irresolute social choice functions, and a resolute social choice function as a special sort of irresolute social choice function.

Definition 2.2 (Resolute Social Choice Function). A Resolute Social Choice Function

$$f : \mathcal{L}(A)^n \rightarrow A$$

aggregates agents' total preferences over A into a single winner of the election.

In this thesis, I will largely (out of convenience) be concerned with resolute social choice functions (so the reader can assume that the functions I consider have some sort of tie-breaking system built in). When I use the phrase 'social choice function', I intend to refer to a resolute social choice function. Occasionally, though, it will be important to emphasise that a particular result does not depend on the resoluteness of the underlying social choice function. I will make this clear when appropriate.

Definition 2.3 (Social Welfare Function). A Social Welfare Function

$$f : \mathcal{L}(A)^n \rightarrow \mathcal{L}(A)$$

aggregates agents' total preferences over A into a single linear order (interpreted as the preference of the group).

The overwhelming majority of the results I present in this thesis relate to social choice functions. The model I present, though, can also be used with social welfare functions. This is a potential avenue for future work.

Definition 2.4 (Election/Vote). A Classical Election, or Vote, is a triple (N, A, f) , where

- $N = \{i, j, k, l, \dots\}$ is a set of voters, with $|N| = n$. It will also sometimes be convenient to write $N = \{1, \dots, n\}$.
- $A = \{a, b, c, d, \dots\}$ is a set of candidates, with $|A| = m$.
- f is a social choice (welfare) function.

Each voter i (or 'agent') submits a linear order L_i over the set of alternatives A , generating a profile $\mathbf{L} = (L_1, \dots, L_n) \in \mathcal{L}(A)^n$. The outcome of the election is given by $f(\mathbf{L})$.

2.3.1 Properties of Social Choice Functions

There are various familiar properties of social choice functions which will be relevant during this thesis (W. Zwicker and Herve Moulin (2016)).

Definition 2.5 (Anonymity). A social choice function f is *anonymous* if, for any bijection $\psi : N \rightarrow N$ and profile $\mathbf{L} = (L_1, \dots, L_n) \in \mathcal{L}(A)^n$, we have that

$$f(L_1, \dots, L_n) = f(L_{\psi(1)}, \dots, L_{\psi(n)})$$

If a social choice function is anonymous, we can permute the names of the agents, and it is guaranteed not to change the result of the election.

Let $\psi : A \rightarrow A$ be a bijection. Let $P \in \mathcal{P}(A)$. By $\psi(P)$, I denote the alternative-wise application of the bijection. So if $P = \{a \succ b\}$, $\psi(a) = b$ and $\psi(b) = a$, then $\psi(P) = \{b \succ a\}$.

Definition 2.6 (Neutrality). A social choice function f is *neutral* if, for any bijection $\psi : A \rightarrow A$ and profile $\mathbf{L} = (L_1, \dots, L_n) \in \mathcal{L}(A)^n$, we have that

$$\psi(f(L_1, \dots, L_n)) = f(\psi(L_1), \dots, \psi(L_n))$$

If a social choice function is neutral and we permute the names of the alternatives, then we can simply calculate the winner of the new election by permuting the name of the previous winner.

Definition 2.7. (Weak Monotonicity) A social choice function f is *weakly monotonic* if the following holds for every $\mathbf{L} \in \mathcal{L}(A)^n$.

Suppose $f(\mathbf{L}) = a$, for some $a \in A$. Let $\mathbf{L}' = (L'_i, L_{-i})$ be an i -variant of \mathbf{L} , where

$$L'_i = L_i \setminus \{b \succ a\} \cup \{a \succ b\}$$

for some $b \in A$ (in other words, voter i moves alternative a up at most one place in her ordering). Then we have that $f(\mathbf{L}') = a$.

Definition 2.8. (Unanimity) A social choice function f is *unanimous* if the following holds for every $\mathbf{L} \in \mathcal{L}(A)^n$, $a \in A$.

Fix $a \in A$. Suppose that for every $i \in N$, for every $b \in A \setminus \{a\}$, we have $a \succ b \in L_i$ (in other words, every voter's favourite alternative is a). Then we must have $f(\mathbf{L}) = a$.

Definition 2.9. (Pareto Efficiency) A social choice function f is *Pareto efficient* if the following holds for every $\mathbf{L} \in \mathcal{L}(A)^n$, $a \in A$.

Suppose that there is some $b \in A \setminus \{a\}$ such that for every $i \in N$ we have $b \succ a \in L_i$ (in other words, there is an alternative, b , that every voter prefers to a). Then we must have $f(\mathbf{L}) \neq a$.

I will also include here a novel property of social choice functions, which I will make use of in Chapters 3 and 4.

Let $\mathbf{L}+$ be the profile we get when we augment \mathbf{L} with $|A|!$ new voters, one holding each possible ranking in $\mathcal{L}(A)$ (if f is anonymous, then we can think of $\mathbf{L}+$ as $\mathbf{L} \cup \mathcal{L}(A)$; otherwise, we must assume some ordering on the rankings in $\mathcal{L}(A)$).

Definition 2.10. (Uniform Voter Addition Invariance) Then f is *Uniform Voter Addition Invariant* (UVAI) iff $f(\mathbf{L}+) = f(\mathbf{L})$ for every $\mathbf{L} \in \mathcal{L}(A)^n$.

Uniform Voter Addition Invariance (UVAI) says that we can add a new set of voters to our existing voters, one holding each possible linear order over the set of alternatives, without changing the result of the election.

2.4 Extending Classical Votes

In this section, I will present the model I will use for the remainder of the thesis. My model extends a classical vote with a proxy mechanism, g .

2.4.1 Proxy Mechanisms

Recall that $\mathcal{P}(A)$ denotes the set of all partial orders over A . Recall that $\mathcal{P}(N)$ designates the powerset of N .

Definition 2.11 (Proxy Mechanism). A function

$$g : \mathcal{P}(A)^n \times N \rightarrow \mathcal{P}(N)$$

is a proxy mechanism iff, for every $\mathbf{P} = (P_1, \dots, P_n) \in \mathcal{P}(A)^n$, for every $i \in N$:

1. If $P_i = \emptyset$, then $g(\mathbf{P}, i) = N \setminus \{i\}$.
2. If $P_i \in \mathcal{L}(A)$, then $g(\mathbf{P}, i) = \{i\}$.
3. If $P_i \notin \mathcal{L}(A)$, then $i \notin g(\mathbf{P}, i)$.

Intuitively, a proxy mechanism takes in a profile of partial orders and assigns to each voter a set of permitted proxies, the voters who they are allowed to choose as their delegate. Recall that the idea is that this set of permitted proxies constitutes the delegates who could represent the voter's interests. Let's turn to the individual clauses in the definition.

Firstly, if agent i submits an empty order, we require (in 1.) that she can choose any other agent as her proxy (every other agent is in her set of permitted proxies). This is because she has no preferences over the alternatives, implying that there is no way for a potential delegate to fail to represent her interests.

Similarly, if agent i submits a linear order, we require (in 2.) that she casts her own vote (she is the only agent in her set of permitted proxies). The motivation for this is simple; if she has already made her mind up about the alternatives, there is no need for her to delegate her vote to another agent.

Finally, if agent i submits a partial order which is not a linear order, then she is not allowed (by 3.) to cast her own vote (she does not appear in her set of permitted proxies). This is because the aggregation function (a social choice or welfare function) takes profiles of linear orders as an input; the model I propose modifies the method of collecting preferences, not the method of aggregating preferences.

2.4.2 Proxy Votes

We are now ready to define proxy votes. Recall that a classical vote was a triple (N, A, f) . A proxy vote adds a proxy mechanism into the mix, and requires agents to submit linear orders over the set of potential proxies, as well as a default vote and a partial order over the set of alternatives.

Definition 2.12 (Proxy Vote). A proxy vote is a tuple

$$(N, A, f, g)$$

where:

- $N = \{i, j, k, l, \dots\}$ is a set of voters, with $|N| = n$. It will also sometimes be convenient to write $N = \{1, \dots, n\}$.
- $A = \{a, b, c, d, \dots\}$ is a set of candidates, with $|A| = m$.
- f is a social choice (welfare) function.
- g is a proxy mechanism.

An agent $i \in N$ submits a triple (P_i, S_i, D_i) , where:

- $P_i \in \mathcal{P}(A)$ is a partial order over the alternatives. So the model allows agents to have made their mind up about some pairwise comparisons but not others.
- $S_i \in \mathcal{L}(N)$ is a linear order over the voters. Intuitively, this order corresponds to a ranking over potential proxies (capturing all the reasons that i might have to prefer a delegate as her proxy independently of the delegate's ability to represent her).
- $D_i \in \mathcal{L}(A)$ is a linear order over the set of alternatives, with $P_i \subseteq D_i$. D_i is a 'default vote'. In the situation where i has no permitted proxies (so $g(\mathbf{P}, i) = \emptyset$), i is required to vote directly, submitting this default vote.

When each agent submits a triple, we have a *proxy vote profile* $(\mathbf{P}, \mathbf{S}, \mathbf{D})$, where

- \mathbf{P} is a (partial) *preference profile*.
- \mathbf{S} is a *proxy choice profile*.
- \mathbf{D} is a *default vote profile*.

Each voter i then receives $g(\mathbf{P}, i)$, a set of permitted proxies, given the preference profile.

If $g(\mathbf{P}, i) = \emptyset$, then i must submit her default vote $D_i \in \mathcal{L}(A)$.

If $g(\mathbf{P}, i) \neq \emptyset$, then i must pick some $j \in g_i(\mathbf{P})$ to cast her vote on her behalf. Let $N' \subseteq N$. Then by $S_i|_{N'}$ I denote the restriction of S_i to N' . Agent i will pick the potential proxy who is ranked highest when we consider $S_i|_{g(\mathbf{P}, i)}$ (in other words, the most preferred delegate from amongst her permitted proxies). Suppose that this is j . Then I will abuse notation by writing that $S_i|_{g(\mathbf{P}, i)} = j$. For the sake of convenience, I will write $S_i|_{\{i\}} = i$ and $S_i|_{\emptyset} = i$, since i casts her own vote if $g(\mathbf{P}, i) = \{i\}$ or if $g(\mathbf{P}, i) = \emptyset$.

So, given a voting profile $\mathbf{P} = (P_1, \dots, P_n)$ and proxy choice profile $\mathbf{S} = (S_1, \dots, S_n)$, each $i \in N$ has a proxy. So we have a *delegation graph* (N, R) where iRj iff

$$j = S_i|_{g(\mathbf{P}, i)}$$

Note that, where it does not have a negative impact on accuracy, I will speak of ‘ i choosing j to be her proxy’ as expressing this formal condition.

Let R^* be the transitive closure of R . For each i , let

$$\Pi_i = \{j \in N \mid iR^*j \text{ and } jRj\}$$

If Π_i is non-empty, it is easy to see that it will be a singleton $\{\pi_i\}$. Call π_i voter i 's *guru*. Note that if i casts her own vote, then we have $\pi_i = i$ (so i will be her own guru). We can then define a guru voting profile

$$\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}} = (P_{\pi_1, S_1, D_1}, \dots, P_{\pi_n, S_n, D_n})$$

Where P_{π_i, S_i, D_i} is the preference order submitted by voter i 's guru, generated according to $(\mathbf{P}, \mathbf{S}, \mathbf{D})$.

I use the notation $\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}}$ to emphasise that this profile results from the proxy vote profile $(\mathbf{P}, \mathbf{S}, \mathbf{D})$. The use of π is supposed to remind the reader that the votes are actually submitted by the gurus π_1, \dots, π_n .

Note that, by construction, $P_{\pi_i, S_i, D_i} \in \mathcal{L}(A)$, for every $i \in N$, since each guru must cast her own vote. So we can use $\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}}$ as the input to a social choice (welfare) function. The *outcome* of the proxy vote is given by $f(\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}})$.

2.4.3 Agents' Preferences over the Outcomes of Proxy Votes

In Chapter 4, we will explore manipulation in a proxy vote setting. To do this, we need to define what it means for an agent to prefer one outcome of a proxy vote over another. Suppose that f is a resolute social choice function. Let $(\mathbf{P}, \mathbf{S}, \mathbf{D})$ be a proxy vote profile. Then let $\mathbf{P}' = (P'_i, P_{-i})$, $\mathbf{S}' = (S'_i, S_{-i})$ and $\mathbf{D}' = (D'_i, D_{-i})$ (so \mathbf{P}' is an i -variant of \mathbf{P} , \mathbf{S}' is an

i -variant of S and D' is an i -variant of D). We will say that i prefers $f(\mathbf{P}'_{\pi',S',D'})$ to $f(\mathbf{P}_{\pi,S,D})$ iff either

$$f(\mathbf{P}'_{\pi',S',D'}) \succ f(\mathbf{P}_{\pi,S,D}) \in P_i \quad (\text{A})$$

or (

$$f(\mathbf{P}_{\pi,S,D}) \succ f(\mathbf{P}'_{\pi',S'}) \notin P_i \quad (\text{B1})$$

and (

$$\pi_i = i \text{ and } \pi'_i \neq i \quad (\text{B2.1})$$

or

$$|P'_i| < |P_i| \text{ and } \pi'_i \neq i). \quad (\text{B2.2})$$

Recall that $|P'_i|$ is the number of pairwise comparisons contained within P'_i , and $|P_i|$ is the number of pairwise comparisons contained within P_i .

Formally, the relationship between the equations that expresses the condition is

$$A \vee (B1 \wedge (B2.1 \vee B2.2)).$$

The first condition ('Strict Preference') is expressed by equation A whilst the second condition ('Effort Preference') is jointly expressed by equations B1 and B2.1/B2.2.

Strict Preference says that an agent will prefer changing what she submits when she prefers the winner of the new proxy vote to the winner of the old proxy vote in her original partial order. Strict Preference is simply the standard formalisation of an agent's preference over outcomes of a vote, updated to take into account the proxy vote setting.

Effort Preference says that an agent will prefer changing what she submits when it does not make the outcome any worse (according to her original ordering), whilst allowing her to submit fewer pairwise comparisons in her new preference ordering. Note that she cares about the size of the vote she actually has to cast (so if she casts her default vote, that is taken to be the size of the vote she casts, rather than the size of the partial order she submitted). In other words, an agent is happy when she has to put in less effort to achieve a result which is no worse. Effort Preference is a novel condition, designed to capture the idea of 'effort' in a proxy voting setting. It is motivated by the assumption that each pairwise comparison takes some effort to decide on and submit in a vote. This is an assumption (familiar from the previous chapter) which features in prominent defences of proxy voting, so I will not question it here.

2.5 Discussion of the Model

In this section, I will attempt to flesh out the conceptual underpinnings of the model, as well as highlighting some of its representational power.

2.5.1 What is a Proxy Mechanism?

A natural question concerns the nature of the proxy mechanism g . Recall that a proxy mechanism takes into account the partial preferences over the set of alternatives submitted by the voters and assigns to each voter a set of permitted proxies, the voters who she is allowed to delegate her vote to. I have suggested that we should interpret this set of permitted proxies as the set of delegators who would represent the voter's interests (based on the preferences they have submitted). But what is the proxy mechanism itself? How does it decide which voters are capable of representing which others?

The first point to make is that there are a lot of possible proxy mechanisms, just as there are a lot of social choice functions. I postpone examining proxy mechanisms from an axiomatic perspective until the next chapter. But it is worth giving some examples of simple proxy mechanisms here.

Definition 2.13. (TRIV)

$$\text{TRIV}(\mathbf{P}, i) = \begin{cases} N \setminus \{i\} & \text{if } P_i = \emptyset \\ \{i\} & \text{if } P_i \in \mathcal{L}(A) \\ \emptyset & \text{otherwise} \end{cases}$$

Definition 2.14. (SUBSET)

$$\text{SUBSET}(\mathbf{P}, i) = \begin{cases} N \setminus \{i\} & \text{if } P_i = \emptyset \\ \{i\} & \text{if } P_i \in \mathcal{L}(A) \\ \{j \in N \setminus \{i\} \mid P_i \subseteq P_j\} & \text{otherwise} \end{cases}$$

Definition 2.15. (STRICT-SUBSET)

$$\text{STRICT-SUBSET}(\mathbf{P}, i) = \begin{cases} N \setminus \{i\} & \text{if } P_i = \emptyset \\ \{i\} & \text{if } P_i \in \mathcal{L}(A) \\ \{j \in N \setminus \{i\} \mid P_i \subset P_j\} & \text{otherwise} \end{cases}$$

Definition 2.16. (UNIV)

$$\text{UNIV}(\mathbf{P}, i) = \begin{cases} N \setminus \{i\} & \text{if } P_i = \emptyset \\ \{i\} & \text{if } P_i \in \mathcal{L}(A) \\ N \setminus \{i\} & \text{otherwise} \end{cases}$$

Definition 2.17. (DICTATOR) For each $i \in N$, fix some $j \in N \setminus \{i\}$ (to make this concrete, we could, for example, pick the lexicographically earliest voter in $N \setminus \{i\}$). Then

$$\text{DICTATOR}(\mathbf{P}, i) = \begin{cases} N \setminus \{i\} & \text{if } P_i = \emptyset \\ \{i\} & \text{if } P_i \in \mathcal{L}(A) \\ \{j\} & \text{otherwise} \end{cases}$$

Definition 2.18. (HYBRID) Fix some $j \in N$. Then

$$\text{HYBRID}(\mathbf{P}, i) = \begin{cases} N \setminus \{i\} & \text{if } P_i = \emptyset \\ \{i\} & \text{if } P_i \in \mathcal{L}(A) \\ N \setminus \{i\} & \text{if } i = j \text{ and } P_i \notin \mathcal{L}(A) \\ \emptyset & \text{otherwise} \end{cases}$$

For most of the remainder of this thesis, I'll focus on the **SUBSET** mechanism, as I think it is a natural interpretation of what it means for a delegate to represent a voter. In the next chapter, I'll also show that it is the unique mechanism satisfying certain desirable properties.

For now, though, note that the choice of proxy mechanism g has a large effect on the proxy vote setting.

If $g = \text{TRIV}$, then every agent will cast her own vote, unless she has no preferences at all over the alternatives (in which case she will delegate). So we are close to a classical vote; proxy voting plays little role here. In particular, the proxy choice profile \mathbf{S} is often fairly irrelevant to the outcome of the election, although the default vote profile \mathbf{D} can be highly relevant.

By contrast, if $g = \text{UNIV}$, then every agent who has not made her mind up fully can delegate her vote to any other agent, regardless of what she thinks on the issues she has made her mind up on. In effect, this is the formal set up of many of the accounts we discussed in the previous chapter; the strictly partial components of the preference profile \mathbf{P} are irrelevant to the outcome of the vote, as is the default vote profile \mathbf{D} . Instead, it is the proxy choice \mathbf{S} that plays a large role in determining the outcome of the election.

If $g = \text{DICTATOR}$, then each agent i has a unique dictator j ; when i submits some but not all pairwise comparisons, then she must delegate her vote to j . Similarly, if $g = \text{HYBRID}$, then the mechanism acts like **UNIV** for some distinguished $j \in N$, and acts like **TRIV** for every $i \in N \setminus \{j\}$. These are not intended as real suggestion for a proxy mechanisms, but should serve to indicate the sheer range of available options.

Having sketched some example of proxy mechanisms and their effects on the voting system, it is time to turn to the question at hand. What actually is a proxy mechanism? I see at least two interpretations of proxy mechanisms.

Firstly, there is a *descriptive* interpretation of proxy mechanisms. On this interpretation, a proxy mechanism describes the behaviour of voters (assuming they act in their own interest). Voters will only choose delegates who represent their interests, and the proxy mechanism makes this constraint explicit. Note that proxy mechanisms allow for different voters to have different interpretations of what it takes for a delegate to represent their interests (since they take the name of the voter as an input). Different proxy mechanisms correspond to different constraints on the judgement of

voters. I think the descriptive interpretation best corresponds to the account of proxy choice I offered at the beginning of this chapter (consider, for example, the Brexit thought experiment).

There is also a *prescriptive* interpretation of proxy mechanisms. On this interpretation, a proxy mechanism is something external to a voter. It could be some aspect of a centralised voting system, or some rule which a voter is required to obey.

When it comes to the classical picture of liquid democracy as a happy medium between direct and representative democracies, I think it is clear that we ought to prefer the descriptive interpretation of proxy mechanisms. When it comes to transitive proxy voting more broadly, though, I think the idea of a centralised proxy mechanism becomes less strange. Suppose, for example, that agents are autonomous software agents making decisions over a large number of alternatives. At this point, one might plausibly want to constrain the sort of delegations that are allowed, to have some means to predict the behaviour of the system.

2.5.2 Cycles

The reader will note that proxy votes are currently underspecified. I have not yet discussed what should happen in the case of delegation cycles. Since other authors have taken delegation cycles seriously in their formalisation of transitive proxy voting, something clearly needs to be said about them.

The first point to make is that the model I have defined leaves freedom to reduce the occurrence of delegation cycles. If we use the **STRICT-SUBSET** mechanism, for example, then delegation cycles can only involve agents who have decided on absolutely no pairwise comparisons. Depending on the setting where transitive proxy voting is used, one might not be too worried about failing to include such agents. So the choice of proxy mechanism can ameliorate the problem of delegation cycles.

More importantly, though, the model permits an easy resolution to delegation cycles. Following Christoff and Grossi (2017), I have had agents submit a default vote which extends their existing vote. One could simply specify that this default vote is submitted directly by any agent who features in a delegation cycle. Thus delegation cycles penalise the agent a little (in that they have to submit more edges), but do not prevent them from voting.

2.5.3 Networks

There is an aspect of transitive proxy voting explored by other authors which is not explicitly present in my model. Several authors arrange agents into a social network (N, T) , and require that agents can only delegate to their neighbours in the network.

Although I have not arranged agents in a network, the model I have presented can accommodate a social network in at least two ways. I'll outline these here, so as to give a sense of the expressive power of the model.

Firstly, we can build the social network into the proxy mechanism g . To give an example, consider NETWORK-UNIV.

Definition 2.19. (NETWORK-UNIV) Let (N, T) be a social network, where T is irreflexive.

$$\text{NETWORK-UNIV}(\mathbf{P}, i) = \begin{cases} N \setminus \{i\} & \text{if } P_i = \emptyset \\ \{i\} & \text{if } P_i \in \mathcal{L}(A) \\ T[i] & \text{otherwise} \end{cases}$$

where $T[i]$ denotes the neighbourhood of i according to the binary relation T . Since T is irreflexive, $i \notin T[i]$, and so NETWORK-UNIV is a well-defined proxy mechanism.

If we use NETWORK-UNIV, we are essentially saying that an agent is capable of being represented by all and only her neighbours in the network. This models the proxy vote setting considered by many authors.

It's clear that we could easily create a networked version of an existing proxy mechanism (for example, SUBSET) by intersecting an agent's permitted proxies with her neighbourhood in the network. So we can impose a network structure on delegations whilst retaining the flexibility of the proxy mechanism framework if we desire.

There's another way to incorporate a social network into my model, namely to place a requirement on \mathbf{S} , the proxy choice profile. Fix some social network (N, T) . For each $i \in N$, we could stipulate that

$$(j, k) \in S_i \text{ iff } (j \in T[i] \text{ and } k \notin T[i]).$$

What this says is that i will always delegate to one of her neighbours if they are in her set of permitted proxies, but is able to delegate further afield if none of her neighbours represents her sufficiently. We could refine this idea further, by having i 's neighbour's neighbours ranked behind her neighbours but ahead of the remaining agents, and so on.

In this way, the model can capture the idea that agents will delegate to their neighbours in the network, whilst taking a more flexible perspective on other permitted delegations. If we assume that delegation is done using some sort of technology, then it seems plausible that an agent might be able to reach beyond her direct neighbours when making a delegation. So this interpretation of a social network may actually have significant conceptual advantages, depending on the setting where proxy voting is used. To my knowledge, this is the first time this way of modelling delegations in a network has been suggested.

2.5.4 Existing Impossibility Results

The final point to make in this section is that a proxy vote (N, A, f, g) is a straightforward generalisation of a classical vote (N, A, f) . In the case where $P_i \in \mathcal{L}(A)$ for every $i \in N$, every agent casts her vote directly. So the additional elements in the proxy vote model (the proxy mechanism g , the proxy choice profile \mathbf{S} and the default vote profile \mathbf{D}) are irrelevant to the outcome of the election.

In particular, this implies that any impossibility result concerning a social choice (welfare) function f will carry over into my setting. So, for example, Arrow's Impossibility Theorem (Arrow (1950)) and the Gibbard-Satterthwaite Theorem (Gibbard (1973), Satterthwaite (1975)) still hold in this novel setting.

2.6 Responses and Rejoinders

In the remainder of this chapter, I will attempt to anticipate some responses to the model I have proposed, and provide some brief rejoinders to those responses.

2.6.1 The Model Requires too much from Voters

As it stands, each voter i submits a triple (P_i, S_i, D_i) , where $P_i \in \mathcal{P}(A)$, $S_i \in \mathcal{L}(N)$ and $D_i \in \mathcal{L}(A)$. This is a lot of information to ask for from voters, especially as one of the principle justifications of proxy voting is that it reduces the workload required to express an opinion. One might worry that this makes the model unworkable from the outset.

One rejoinder to this response is a companions-in-guilt argument. As I showed in the previous section, various authors have required that voters submit similar amounts of information when modelling proxy voting. So there's a tacit assumption in the literature that this doesn't doom the model to failure.

Of course, this sort of response will only go so far, since it doesn't address the actual objection at hand. Let us look at what we require of voters.

A key assumption driving proxy voting is that each pairwise comparison between alternatives requires some effort for a voter i to decide on and then submit. Since P_i does not have to include every pairwise comparison, it must be at most as burdensome for a voter as the vote required by a standard election. So I take it that it is not P_i that is the problem.

Does submitting D_i , the default preference, require the same amount of effort from a voter as submitting a linear order? Well, it depends how we interpret how a voter picks the default preference. If we suppose the voter puts a significant amount of energy into making each pairwise comparison, then there is no essential difference between the default preference and a

complete vote; in other words, the voter might as well have voted directly, and we should forget about proxy voting.

This seems implausible to me, though. The very point of a default preference is that it is little more than a placeholder, a device which serves some practical purpose but has little ideological significance. I think a more realistic perspective on the default preference is that each voter extends P_i at random, or chooses the lexicographically earliest extension of P_i , and so on. Given such a perspective, the fact that a voter must submit a default preference shouldn't be an issue; she still puts in less effort than she would to make up her mind on each issue. As for the cost required for a voter i to submit a default preference, note that she only incurs this cost when she has no permitted proxies (or, depending on one's preferred account of cycles, is involved in a delegation cycle). In such a case, we are essentially saying that her original preference P_i was sufficiently baroque that no other voter could represent her interests. Here, it seems natural that she might have to pay some extra cost to express her view.

That leaves us with S_i , then, the linear order over the voters. If $|N| = n$ is large, the cost of forming and submitting such an order could be extremely high. I take it that S_i is the element of the ballot for which the objection at hand has the most bite, then.

Here, it is helpful to delineate two different perspectives on the proxy vote model, a *synchronic* perspective and a *diachronic* perspective.

The synchronic perspective takes the formal requirements of the model literally. It assumes we are dealing with a static environment in which voters really do submit triples (P_i, S_i, D_i) , from which the outcome of the election is calculated deterministically.

The diachronic perspective, by contrast, assumes that the model is a static representation of a process which is inherently dynamic. The diachronic perspective could interpret the model as follows. A voter i submits P_i . She is then told $g(\mathbf{P}, i)$, her set of permitted proxies. If $g(\mathbf{P}, i) = \emptyset$, she submits $D_i \supset P_i$, her default preference. If $g(\mathbf{P}, i) \neq \emptyset$, she picks some $j \in g(\mathbf{P}, i)$ to be her proxy.

From the diachronic perspective understood as above, the objection we are considering doesn't have any teeth. Rather than S_i , the voter i is really only required to specify the name of a proxy.

It's clear that any real world version of transitive proxy voting will be situated within a dynamic environment. Indeed, dynamics are essential to Blum and Zuber (2016)'s characterisation of liquid democracy. There is little doubt that the diachronic perspective will be the correct one when we look to the actual situations which the model purports to describe.

Why do we need the synchronic perspective at all, then? Why not build some sort of dynamics into the model directly, such that the sort of worry expressed in this subsection doesn't apply? Essentially, this question relates to a difficulty with any formal model of a highly complex system. We want

the model to contain a sufficient level of detail to describe features of the system relevant to its analysis, and a sufficient level of abstraction to ensure that such analysis is tractable. I have opted for a static model, trusting that it is sufficiently expressive to ask non-trivial questions about the setting it represents.

In Chapter 4, we'll see a drawback of the model's staticity when describing strategic behaviour amongst agents. It may be that the proper response to this drawback is to add dynamics to the model. For the remainder of this thesis, though, I'll use the static formal model with the diachronic perspective in the conceptual background.

2.6.2 The Model Renders Voter Behaviour Trivial

A second worry also relates to a general issue with modelling. When we abstract away information from the transitive proxy voting setting (as we must when we attempt to formalise it), there is a risk that we lose the ability to describe the behaviour of voters. We have formal features such as proxy choice profiles \mathcal{S} , but do these features actually latch onto their intended interpretation, or are they just useless formal appendages of the model?

One way to express this worry is to look at the behaviour of agents within the model. If we find that agents act differently from how they would in the situations the model purports to describe, then this is evidence of some deficiency in the model. More concretely, if we find that there is some dominant strategy in the formal setting (even in the absence of information about the preferences of other voters) that does not exist in the real-world setting, then this poses a problem for the model. I will spend the remainder of this subsection showing that there is no obvious dominant strategy. I consider two candidates for a dominant strategy of this sort.

Dominant Strategy 1: Will Agents Always Submit Linear Orders?

The first worry is that no agent would prefer to delegate her vote in the setting I have described. Instead, every agent will simply vote directly. It should be seen from the way I defined the notion of a voter's preference over outcomes of the election that this isn't the case, but I'll describe a toy example which makes this clearer.

Suppose f is unanimous (I make no assumptions about g). Let $N = \{1, 2\}$ and let $A = \{a, b\}$. Suppose that:

$$P_i = \{a \succ b\}, \forall i \in N$$

As it stands, a is the winner of the election, since f is unanimous. Consider agent 1. Currently $P_1 \in \mathcal{L}(A)$. But note that 1 would prefer to submit $P'_1 = \emptyset$. In such a situation, $g((P'_1, P_2), 1) = \{2\}$, and so $\pi_1 = 2$. So a would

remain the winner of the election, since

$$P_{\pi, S, D} = (P'_1, P_2)_{\pi', S, D}$$

But note that $|P'_1| < |P_1|$. So 1 would prefer to submit P'_1 than P_1 . This demonstrates that there are situations where agents would prefer not to submit linear orders.

Dominant Strategy 2: Will Agents Always Submit Empty Orders?

The second worry is the opposite of the first. The worry is that based on the way we have defined the notion of an agent preferring an outcome of the election, agents would always prefer to submit empty orders (since they require no effort from the agent). Again, I'll give a toy example which shows this isn't the case.

Suppose f is the plurality rule (I make no assumptions about g). Let $N = \{1, 2, 3\}$ and let $A = \{a, b\}$. Suppose that:

$$\begin{aligned} P_1 &= \{a \succ b\} \\ P_2 &= \{a \succ b\} \\ P_3 &= \{b \succ a\} \end{aligned}$$

and that

$$S_1 = \{3 \succ 2 \succ 1\}$$

As it stands, a is the winner of the election, since f is the plurality rule. Consider agent 1. Examine what would happen if $P'_1 = \emptyset$. In such a situation, $g((P'_1, P_{-1}), 1) = \{2, 3\}$, and so $\pi_1 = 3$. But then b would become the winner of the election, since f is the plurality rule. But $a \succ b \in P_1$. So 1 would not prefer to submit P'_1 over P_1 . This demonstrates that there are situations where agents would prefer not to change their vote to an empty preference order.

Of course, these two dominant strategies do not exhaust the possible trivialities. But I hope they give some evidence to the reader that the model doesn't immediately collapse into triviality, so that we can proceed with examining more interesting properties of the model.

2.7 Concluding Remarks

In this chapter, I've defined the model which I will use in Chapter 3 and Chapter 4. In the next chapter, I'll examine proxy mechanisms and proxy votes from an axiomatic perspective. I'll characterise the SUBSET mechanism using some of these axioms, and prove a general impossibility result for natural properties of (social choice function, proxy mechanism) pairs (f, g) .

Chapter 3

Properties of Proxy Mechanisms and Proxy Votes: An Axiomatic Analysis

In the previous chapter, I outlined a novel model of transitive proxy voting. At the heart of the model was a proxy mechanism, g . In this chapter, I will examine proxy mechanisms from an axiomatic perspective. I will also examine properties of pairs (f, g) , where f is a social choice function and g is a proxy mechanism.

3.1 Proxy Mechanisms

In this section, I'll define some natural properties which proxy mechanisms can satisfy. I'll then use some of these properties to characterise the **SUBSET** mechanism (defined in the previous chapter).

3.1.1 Properties of Proxy Mechanisms

Recall that a proxy mechanism is a function

$$g : \mathcal{P}(A)^n \times N \rightarrow \mathcal{P}(N).$$

Let us define some properties of proxy mechanisms as follows.

Let $\psi : N \rightarrow N$ be a bijection. Let $N' \subseteq N$. Then I write $\psi(N')$ to denote the image of N' under ψ . Let $\mathbf{P} \in \mathcal{P}(A)^n$ be a partial preference profile. Abusing notation, I write

$$\psi(\mathbf{P}) = \psi(P_1, \dots, P_n) = (P_{\psi(1)}, \dots, P_{\psi(n)})$$

Definition 3.1. (Proxy Mechanism Anonymity) A proxy mechanism g is *anonymous* iff for every preference profile $\mathbf{P} \in \mathcal{P}(A)^n$ and every bijection $\psi : N \rightarrow N$, we have that

$$\psi(g(\mathbf{P}, i)) = g(\psi(\mathbf{P}), \psi(i))$$

Proxy Mechanism Anonymity says that if we rename the agents, then a renamed agent's set of permitted proxies will just be the original agent's set of permitted proxies renamed. In other words, the proxy mechanism is blind to the identity of the individual agents.

Let $\psi : A \rightarrow A$ be a bijection. Let $P \in \mathcal{P}(A)$. By $\psi(P)$, I denote the alternative-wise application of the bijection. So if $P = \{a \succ b\}$, $\psi(a) = b$ and $\psi(b) = a$, then $\psi(P) = \{b \succ a\}$. Let $\mathbf{P} \in \mathcal{P}(A)^n$ be a partial preference profile. Abusing notation, I write

$$\psi(\mathbf{P}) = \psi(P_1, \dots, P_n) = (\psi(P_1), \dots, \psi(P_n))$$

Definition 3.2. (Proxy Mechanism Neutrality) A proxy mechanism g is *neutral* iff for every preference profile $\mathbf{P} \in \mathcal{P}(A)^n$ and every bijection $\psi : A \rightarrow A$, we have that

$$g(\mathbf{P}, i) = g(\psi(\mathbf{P}), i)$$

Proxy Mechanism Neutrality says that we can rename the alternatives without affecting each agent's set of permitted proxies.

Definition 3.3. (Proxy Availability (PA)) g satisfies *PA* iff for every $Q \in \mathcal{P}(A)$, for every $i \in N$, there is some $\mathbf{P} \in \mathcal{P}(A)^n$ such that $P_i = Q$ and

$$g(\mathbf{P}, i) \neq \emptyset$$

Proxy Availability (PA) says that every voter should be able to find potential proxies for their votes, regardless of what views they hold, in at least some profile. In other words, every voter is capable of being represented, regardless of her views.

Definition 3.4. (Independence of Irrelevant Proxies (IIP)) g satisfies *IIP* iff for every $\mathbf{P}, \mathbf{P}' \in \mathcal{P}(A)^n$, for every $i, j \in N$, if $P_i = P'_i$ and $P_j = P'_j$, then

$$j \in (\mathbf{P}, i) \text{ iff } j \in g(\mathbf{P}', i)$$

Independence of Irrelevant Proxies (IIP) says that whether j is a permitted proxy for i should depend only on i 's and j 's preferences, not on the preferences of the other agents.

Definition 3.5. (Zero Regret (ZR)) g satisfies *ZR* iff there is no triple $(\mathbf{P}, \mathbf{S}, \mathbf{D})$ (where $\mathbf{P} \in \mathcal{P}(A)^n$, $\mathbf{S} \in \mathcal{L}(N)^n$ and $\mathbf{D} \in \mathcal{L}(A)^n$) such that, for some $i \in N$:

$$P_i \not\subseteq P_{\pi_i, S_i, D_i}$$

Zero Regret (ZR) says that a proxy mechanism guarantees that every agent's vote ends up being cast by someone who agrees with them completely (i.e. that they have no regrets about the vote submitted by their guru).

Before defining the next condition on proxy mechanisms, it will be useful to define a couple of terms.

Definition 3.6. (Agreement and Disagreement) Let $P_i, Q_i \in \mathcal{P}(A)$. Then

$$\text{Agree}(P_i, Q_i) = \{a \succ_{P_i} b \mid a \succ_{Q_i} b\}$$

and

$$\text{Disagree}(P_i, Q_i) = \{a \succ_{P_i} b \mid b \succ_{Q_i} a\}$$

So $\text{Agree}(P_i, Q_i)$ returns the set of pairwise comparisons which P_i and Q_i agree on, and $\text{Disagree}(P_i, Q_i)$ returns the set of pairwise comparisons which P_i and Q_i disagree on. Note that if $P_i, Q_i \in \mathcal{L}(A)$ (i.e. if they are linear orders), we have

$$|\text{Agree}(P_i, Q_i)| + |\text{Disagree}(P_i, Q_i)| = \frac{1}{2}m(m-1)$$

since $\frac{1}{2}m(m-1)$ is the total number of pairwise comparisons that need to be made to form a total/linear order over $|A| = m$ alternatives.

With Definition 3.6 in mind, we can define another condition on proxy mechanisms.

Definition 3.7. (Preference Monotonicity (PM)) g satisfies *PM* iff the following condition holds for every $\mathbf{P} \in \mathcal{P}(A)^n$, for every $i \in N$. Suppose $j \in g(\mathbf{P}, i)$ and $j \neq i$. Then for every $k \in N \setminus \{i\}$, if

$$\text{Agree}(P_i, P_j) \subseteq \text{Agree}(P_i, P_k)$$

and

$$\text{Disagree}(P_i, P_k) \subseteq \text{Disagree}(P_i, P_j)$$

then $k \in g(\mathbf{P}, i)$.

Preference Monotonicity (PM) says that if j is a permitted proxy for i and k agrees with i on at least the same things as j whilst disagreeing with i on at most the same things as j , then k should also be a permitted proxy for i .

3.1.2 Characterising SUBSET

Recall that SUBSET was defined as follows:

Definition 3.8. (SUBSET)

$$\text{SUBSET}(\mathbf{P}, i) = \begin{cases} N \setminus \{i\} & \text{if } P_i = \emptyset \\ \{i\} & \text{if } P_i \in \mathcal{L}(A) \\ \{j \in N \setminus \{i\} \mid P_i \subseteq P_j\} & \text{otherwise} \end{cases}$$

Theorem 3.9. SUBSET is the unique proxy mechanism satisfying *Proxy Availability, Independence of Irrelevant Proxies, Zero Regret* and *Preference Monotonicity*.

Proof. Clearly, SUBSET satisfies *PA, IIP* and *ZR*. To see that SUBSET satisfies *PM*, suppose that $j \in \text{SUBSET}(\mathbf{P}, i)$ and $j \neq i$, for some $\mathbf{P} \in \mathcal{P}(A)^n$, $i, j \in N$. So $P_i \subseteq P_j$. Suppose that there is $k \in N$ such that $\text{Agree}(P_i, P_j) \subseteq \text{Agree}(P_i, P_k)$ and $\text{Disagree}(P_i, P_k) \subseteq \text{Disagree}(P_i, P_j)$. Then we must have $P_i \subseteq P_k$, since $\text{Agree}(P_i, P_j) = P_i$, since $P_i \subseteq P_j$. So $k \in \text{SUBSET}(\mathbf{P}, i)$, as required.

For the other direction (i.e. to show uniqueness), I prove the contrapositive. Suppose $g \neq \text{SUBSET}$ is a proxy mechanism, and suppose g satisfies *PA, IIP* and *PM*. I will show that g does not satisfy *ZR*.

It will help to prove the following intermediate claim.

Lemma 3.10. Let $\mathbf{P} \in \mathcal{P}(A)^n$ and $i, j \in N$, such that $P_i \notin \mathcal{L}(A)$. Then if $P_i \subseteq P_j$, and g satisfies *PA, IIP* and *PM*, we have $j \in g(\mathbf{P}, i)$.

Proof. Since g satisfies *PA*, there must be some $\mathbf{P}' \in \mathcal{P}(A)^n$ such that $P'_i = P_i$ and $g(\mathbf{P}', i) \neq \emptyset$. Suppose $k \in g(\mathbf{P}', i)$, for some $k \in N$. Then we can construct a new profile \mathbf{P}'' where

$$\begin{aligned} P''_i &= P'_i = P_i \\ P''_j &= P_j \\ P''_k &= P'_k \end{aligned}$$

By *IIP*, we must have $k \in g(\mathbf{P}'', i)$, since $P''_i = P'_i$ and $P''_k = P'_k$. But then by *PM*, we must have $j \in g(\mathbf{P}'', i)$, since

$$P''_i = P_i \subseteq P_j = P''_j$$

implying that j must agree at least as much with i as k in profile \mathbf{P}'' . But then by another application of *IIP*, we must have $j \in g(\mathbf{P}, i)$, since $P''_i = P_i$ and $P''_j = P_j$. \square

We are now ready to prove the uniqueness of SUBSET. Since $g \neq \text{SUBSET}$, there must be some $\mathbf{P} \in \mathcal{P}(A)^n$ and $i, j \in N$ with $P_i \notin \mathcal{L}(A)$ such that either

$$P_i \not\subseteq P_j \text{ and } j \in g(\mathbf{P}, i)$$

or

$$P_i \subseteq P_j \text{ and } j \notin g(\mathbf{P}, i)$$

But note that Lemma 3.10 rules out this latter case. So we only need to consider the case where

$$P_i \not\subseteq P_j \text{ and } j \in g(\mathbf{P}, i)$$

Since $P_i \not\subseteq P_j$, there must be some $a \succ b \in P_i$ such that $a \succ b \notin P_j$.

But now consider a profile \mathbf{P}' where, for some $k \in N$:

$$\begin{aligned} P'_i &= P_i \\ P'_j &= P_j \\ P'_j \cup \{b \succ a\} &\subseteq P'_k, \text{ and } P'_k \in \mathcal{L}(A) \end{aligned}$$

Note that this profile is well defined; since $a \succ b \notin P_j = P'_j$, we must have that $P'_j \cup \{b \succ a\}$ is still anti-symmetric, and thus can be extended to a linear order P'_k .

Since $P'_i = P_i$ and $P'_j = P_j$, we must have $j \in g(\mathbf{P}', i)$, by *IIP*. But then we must have $k \in g(\mathbf{P}', j)$ by Lemma 3.10, since $P'_j \subseteq P'_k$. So then if i picks j as her proxy and j picks k as her proxy, then k will be i 's guru. But $P'_i \not\subseteq P'_k$, since $a \succ b \in P_i = P'_i$ and $a \succ b \notin P'_k$ (because $b \succ a \in P'_k$). So g is not *ZR*. \square

Showing each condition is necessary

I have characterised **SUBSET** as the conjunction of four conditions: *Surjectivity*, *IIP*, *ZR* and *PM*. We can also show that each of these conditions is individually necessary for characterising **SUBSET**, by showing that the other three conditions are not jointly sufficient.

Proposition 3.11. *PA* is necessary for characterising **SUBSET**.

Proof. Recall that **TRIV** was defined as follows.

$$\text{TRIV}(\mathbf{P}, i) = \begin{cases} N \setminus \{i\} & \text{if } P_i = \emptyset \\ \{i\} & \text{if } P_i \in \mathcal{L}(A) \\ \emptyset & \text{otherwise} \end{cases}$$

TRIV is a proxy mechanism which does not satisfy *PA*. But note that **TRIV** does satisfy *IIP*, *ZR* and *PM*. \square

Proposition 3.12. *ZR* is necessary for characterising **SUBSET**.

Proof. Recall that **UNIV** was defined as follows:

$$\text{UNIV}(\mathbf{P}, i) = \begin{cases} N \setminus \{i\} & \text{if } P_i = \emptyset \\ \{i\} & \text{if } P_i \in \mathcal{L}(A) \\ N \setminus \{i\} & \text{otherwise} \end{cases}$$

UNIV is a proxy mechanism which does not satisfy *ZR*. But note that **UNIV** does satisfy *PA*, *IIP* and *PM*. \square

Proposition 3.13. *PM* is necessary for characterising **SUBSET**.

Proof. Consider g defined as follows:

$$g(\mathbf{P}, i) = \begin{cases} N \setminus \{i\} & \text{if } P_i = \emptyset \\ \{i\} & \text{if } P_i \in \mathcal{L}(A) \\ \{j \in N \mid P_j \in \mathcal{L}(A) \text{ and } P_i \subset P_j\} & \text{otherwise} \end{cases}$$

g is a proxy mechanism which does not satisfy PM (just consider some $P_j \notin \mathcal{L}(A)$ such that $P_i \subseteq P_j$). But note that g does satisfy PA , IIP and ZR . \square

Proposition 3.14. IIP is necessary for characterising **SUBSET**.

Proof. Consider g defined as follows:

$$g(\mathbf{P}, i) = \begin{cases} N \setminus \{i\} & \text{if } P_i = \emptyset \\ \{i\} & \text{if } P_i \in \mathcal{L}(A) \\ \{j \in N \mid P_j \in \mathcal{L}(A)\} & \text{if } P_i \notin \mathcal{L}(A) \text{ and } P_i \subseteq P_j, \forall j \in N \setminus \{i\} \\ & \text{such that } P_j \in \mathcal{L}(A) \\ \emptyset & \text{otherwise} \end{cases}$$

g is a proxy mechanism which does not satisfy IIP . But note that g does satisfy PA , PM and ZR . \square

3.1.3 Discussion

It is worth briefly commenting on the properties used to characterise **SUBSET**.

I take it that *proxy availability* is an important feature of any proxy mechanism. The whole point of proxy voting is that voters are (at least in principle) capable of finding delegates to represent them, regardless of their views or their identity. So I will not challenge proxy availability here.

Independence of irrelevant proxies plays a large role in the characterisation of **SUBSET**. Here is one way of motivating independence of irrelevant proxies. In the previous chapter, I argued that we should think of a voter i 's set of permitted proxies as the set of delegates who are capable of representing i 's interests. Membership in this set is binary; a delegate is either capable of representing i 's interests or she is not. Using proxy mechanisms, we are interpreting this capacity in terms of a correspondence between the voter's views and the delegates'. Independence of irrelevant proxies essentially says that each delegate's capacity to represent a voter is independent of what the other delegates think about the issues. If a delegate is capable (or not) of representing a voter and some other voter changes her mind, then this should not affect the capacity, since the correspondence between the voter and delegate's views still exists (or not).

The interpretation of the previous paragraph also lends itself to a defence of *preference monotonicity*, which goes as follows. Suppose some delegate

is capable of representing some voter. This indicates that there is some correspondence between the voter's views and the delegate's. Suppose now that there is an even greater correspondence between some other delegate's and the voter's views. Then this other delegate is also capable of representing the voter.

We are left, then, with *zero regret*. I take it that zero regret is desirable from the perspective of an individual voter, but implausible if we take a proxy mechanism to describe all real world delegations. After all, it seems a consequence of the transitivity of transitive proxy voting that a voter's vote can end up going to someone who they may disagree with. Moreover, one can imagine situations where a voter might have no objection to this (e.g. since the pairwise comparison she cares about most is still represented, or since she has still been spared the effort of voting directly). So we might prefer a proxy mechanism to satisfy zero regret, without rejecting out of hand any proxy mechanism that fails to satisfy it. Future work could address this.

3.2 Properties of Proxy Votes

In the previous section, we identified some plausible properties of proxy mechanisms, and characterised a specific mechanism (the SUBSET mechanism) in terms of these properties.

In this section, we examine properties of proxy votes as a whole. Recall that proxy votes are tuples

$$(N, A, f, g).$$

It follows that the properties we identify will be properties of pairs (f, g) . So these properties will interact both with novel properties of proxy mechanisms g , and with familiar properties of social choice functions f .

3.2.1 Defining Some Properties

Proxy Vote Anonymity

Before defining a notion of anonymity in a proxy vote setting, it will be useful to define some notation. Suppose $\psi : N \rightarrow N$ is a bijection.

Suppose $\mathbf{P} \in \mathcal{P}(A)^n$ is a preference profile. Then it will be convenient to write

$$\psi(\mathbf{P}) = (P_{\psi(1)}, \dots, P_{\psi(n)})$$

to denote the agent-wise application of the bijection ψ . Likewise with a default vote profile $\mathbf{D} \in \mathcal{L}(A)^n$.

Similarly, suppose $\mathbf{S} \in \mathcal{L}(N)^n$ is a proxy choice profile. Then I write

$$\psi(\mathbf{S}) = (S_{\psi(1)}, \dots, S_{\psi(n)})$$

to denote the agent-wise application of the bijection ψ .

Definition 3.15. (Proxy Vote Anonymity) A pair (f, g) , where f is a social choice function and g is a proxy mechanism, satisfies *Proxy Vote Anonymity* iff for every $\mathbf{P} \in \mathcal{P}(A)^n$, for every $\mathbf{S} \in \mathcal{L}(N)^n$ for every $\mathbf{D} \in \mathcal{L}(A)^n$, and for every bijection $\psi : N \rightarrow N$:

$$f(\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}}) = f(\psi(\mathbf{P})_{\pi, \psi(\mathbf{S}), \psi(\mathbf{D})})$$

Proxy Vote Anonymity says that renaming the agents does not affect the result of the proxy vote.

Neutrality

Before defining a notion of neutrality in a proxy vote setting, it will be helpful to define some notation. Suppose $\psi : A \rightarrow A$ is a bijection. Let $P \in \mathcal{P}(A)$. Then I write $\psi(P)$ to denote the alternative-wise permutation of P . For example, if $P = \{a \succ b\}$, then $\psi(P) = \{\psi(a) \succ \psi(b)\}$.

Suppose $\mathbf{P} \in \mathcal{P}(A)^n$ is a (partial) preference profile. Then, abusing notation slightly, it will also be convenient to write

$$\psi(\mathbf{P}) = (\psi(P_1), \dots, \psi(P_n))$$

to denote the agent-wise application of the bijection ψ . Likewise with a default vote profile $\mathbf{D} \in \mathcal{L}(A)^n$.

Definition 3.16. (Proxy Vote Neutrality) A pair (f, g) , where f is a social choice function and g is a proxy mechanism, satisfies *Proxy Vote Neutrality* iff for every $\mathbf{P} \in \mathcal{P}(A)^n$, for every $\mathbf{S} \in \mathcal{L}(N)^n$, for every $\mathbf{D} \in \mathcal{L}(A)^n$ and for every bijection $\psi : A \rightarrow A$:

$$\psi(f(\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}})) = f(\psi(\mathbf{P})_{\pi, \mathbf{S}, \psi(\mathbf{D})})$$

Proxy Vote Neutrality says that renaming the alternatives just renames the outcome of the proxy vote.

Proxy Vote Monotonicity

In a proxy vote setting, voters submit partial orders over alternatives. This means there are two ways they can increase their support for an alternative $a \in A$. They can either add an edge $a \succ b$, or remove an edge $b \succ a$. In a classical vote setting (in which agents submit linear orders over alternatives), these two notions coincide, since one cannot add an edge $a \succ b$ without removing an edge $b \succ a$ (and *vice versa*).

With this in mind, we can distinguish between two notions of monotonicity in a proxy vote setting: ‘addition monotonicity’ and ‘deletion monotonicity’.

Definition 3.17. (Proxy Vote Addition Monotonicity (PVAM)) A pair (f, g) , where f is a social choice function and g is a proxy mechanism, satisfies *PVAM* iff the following holds for every $\mathbf{P} \in \mathcal{P}(A)^n$, every $\mathbf{S} \in \mathcal{L}(N)^n$ and every $\mathbf{D} \in \mathcal{L}(A)^n$.

Suppose $f(\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}}) = a$ for some $a \in A$. Consider an i-variant of \mathbf{P} , $\mathbf{P}' = (P'_i, P_{-i})$, where $P'_i = P_i \cup \{a \succ b\}$, for some $b \in A$. Then $f(\mathbf{P}'_{\pi, \mathbf{S}, \mathbf{D}}) = a$.

Proxy Vote Addition Monotonicity (PVAM) says that if the winner under some proxy vote profile $(\mathbf{P}, \mathbf{S}, \mathbf{D})$ is a , and we modify \mathbf{P} by having some agent add a pairwise comparison to favour a , then the winner should remain a .

Definition 3.18. (Proxy Vote Deletion Monotonicity (PVDM)) A pair (f, g) , where f is a social choice function and g is a proxy mechanism, satisfies *PVDM* iff the following holds for every $\mathbf{P} \in \mathcal{P}(A)^n$, every $\mathbf{S} \in \mathcal{L}(N)^n$ and every $\mathbf{D} \in \mathcal{L}(A)^n$.

Suppose $f(\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}}) = a$, for some $a \in A$. Consider an i-variant of \mathbf{P} , $\mathbf{P}' = (P'_i, P_{-i})$, where $P'_i = P_i \setminus \{b \succ a\}$, for some $b \in A$. Then $f(\mathbf{P}'_{\pi, \mathbf{S}, \mathbf{D}}) = a$.

Proxy Vote Deletion Monotonicity (PVDM) says that if the winner under some proxy vote profile $(\mathbf{P}, \mathbf{S}, \mathbf{D})$ is a , and we modify \mathbf{P} by having some agent delete a pairwise comparison which favours some other alternative over a , then the winner should remain a .

3.2.2 How do these relate to classical properties of social choice functions?

Having defined these properties of pairs (f, g) , it's interesting to explore how they relate to properties of the individual components f and g . In this subsection, I'll focus on properties of f .

Proposition 3.19. If (f, g) satisfies *proxy vote anonymity*, then f satisfies *anonymity*.

Proof. By contraposition. Consider the case where every agent submits a linear order. \square

Proposition 3.20. If (f, g) satisfies *proxy vote neutrality*, then f satisfies *neutrality*.

Proof. By contraposition. Consider the case where every agent submits a linear order. \square

Recall that a social choice function f is *weakly monotonic* if the following holds for every $\mathbf{L} \in \mathcal{L}(A)^n$.

Suppose $f(\mathbf{L}) = a$, for some $a \in A$. Let $\mathbf{L}' = (L'_i, L_{-i})$ be an i -variant of \mathbf{L} , where

$$L'_i = L_i \setminus \{b \succ a\} \cup \{a \succ b\}$$

for some $b \in A$ (in other words, voter i moves alternative a up at most one place in her ordering). Then we have that $f(\mathbf{L}') = a$.

Proposition 3.21. If (f, g) satisfies *proxy vote addition monotonicity* and *proxy vote deletion monotonicity*, then f satisfies *weak monotonicity*.

Proof. By contraposition. Suppose f is not weakly monotonic. We will show that either (f, g) fails to satisfy PVAM or (f, g) fails to satisfy PVDM.

Since f is not weakly monotonic, there must be $\mathbf{P}, \mathbf{P}' \in \mathcal{L}(A)^n$, where $\mathbf{P}' = (P'_i, P_{-i})$ and

$$P'_i = P_i \setminus \{b \succ a\} \cup \{a \succ b\}$$

such that $f(\mathbf{P}) = a$ and $f(\mathbf{P}') \neq a$.

Define $P''_i = P_i - \{b \succ a\}$. By definition, $P_i = P''_i \cup \{a \succ b\}$. Define $\mathbf{P}'' = (P''_i, P_{-i})$. Fix some arbitrary proxy choice profile \mathbf{S} and default vote profile \mathbf{D} . We know $f(\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}}) = a$, by assumption. We also know $f(\mathbf{P}'_{\pi, \mathbf{S}, \mathbf{D}}) \neq a$, by assumption. If $f(\mathbf{P}''_{\pi, \mathbf{S}, \mathbf{D}}) = a$, then (f, g) fails to satisfy PVAM, by definition (since adding the edge $a \succ b$ changes the winner from a). If $f(\mathbf{P}''_{\pi, \mathbf{S}, \mathbf{D}}) \neq a$, then (f, g) fails to satisfy PVDM, by definition (since removing the edge $b \succ a$ has changed the winner from a). \square

3.2.3 How do these relate to properties of proxy mechanisms?

We can also explore the interaction between properties of f and g , and properties of the pair (f, g) .

Lemma 3.22. If f is anonymous and g is proxy mechanism anonymous, then (f, g) is proxy vote anonymous.

Proof. Let a proxy vote profile $(\mathbf{P}, \mathbf{S}, \mathbf{D})$ be arbitrary. Pick some bijection $\psi : N \rightarrow N$. Then we must have

$$\begin{aligned} f(\psi(\mathbf{P})_{\pi, \psi(\mathbf{S}), \psi(\mathbf{D})}) &= f(\psi(\mathbf{P})_{\pi, \mathbf{S}, \mathbf{D}}) && \text{(since } g \text{ is anonymous)} \\ &= f(\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}}) && \text{(since } f \text{ is anonymous)} \end{aligned}$$

\square

Lemma 3.23. If f is neutral and g is proxy mechanism neutral, then (f, g) is proxy vote neutral.

Proof. Let a proxy vote profile $(\mathbf{P}, \mathbf{S}, \mathbf{D})$ be arbitrary. Pick some bijection $\psi : A \rightarrow A$. Then we must have

$$\begin{aligned} f(\psi(\mathbf{P})_{\pi, \mathbf{S}; \psi(\mathbf{D})}) &= f(\psi(\mathbf{P}_{\pi, \mathbf{S}; \mathbf{D}})) && \text{(since } g \text{ is neutral)} \\ &= \psi(f(\mathbf{P}_{\pi, \mathbf{S}; \mathbf{D}})) && \text{(since } f \text{ is neutral)} \end{aligned}$$

□

Note that, in general, the other direction of Theorems 3.22 and 3.23 won't hold.

Recall that Uniform Voter Addition Invariance was defined as follows. Let $\mathbf{L} \in \mathcal{L}(A)^n$ be arbitrary. Let $\mathbf{L}+$ be the profile we get when we augment \mathbf{L} with $|A|!$ new voters, one holding each possible ranking in $\mathcal{L}(A)$ (if f is anonymous, then we can think of $\mathbf{L}+$ as $\mathbf{L} \cup \mathcal{L}(A)$; otherwise, we must assume some ordering on the rankings in $\mathcal{L}(A)$). Then f is *Uniform Voter Addition Invariant* (UVAI) iff $f(\mathbf{L}+) = f(\mathbf{L})$.

Lemma 3.24. If (f, g) satisfies *proxy vote addition monotonicity*, g satisfies *preference monotonicity* and f satisfies *uniform voter addition invariance*, then f satisfies *weak monotonicity*

Proof. By contraposition. Suppose f does not satisfy weak monotonicity but does satisfy uniform voter addition invariance. Suppose also that g satisfies preference monotonicity. We will show that (f, g) does not satisfy proxy vote addition monotonicity.

Since f does not satisfy weak monotonicity, we know that for some $\mathbf{L} \in \mathcal{L}(A)^n$ such that $f(\mathbf{L}) = a$, there is an i -variant $\mathbf{L}' = (L'_i, L_{-i})$ with

$$L'_i = L_i \setminus \{b \succ a\} \cup \{a \succ b\}$$

for some $b \in A$, and $f(\mathbf{L}') \neq a$.

Consider $\mathbf{L}+$ and $\mathbf{L}'+$, the uniform voter augmented versions of, respectively, \mathbf{L} and \mathbf{L}' . Since f satisfies uniform voter addition invariance, we have that

$$f(\mathbf{L}+) = f(\mathbf{L}) = a$$

and

$$f(\mathbf{L}'+) = f(\mathbf{L}') \neq a$$

Consider the profile \mathbf{P} that is exactly like $\mathbf{L}+$ (and $\mathbf{L}'+$), but where

$$P_i = L_i \setminus \{b \succ a\} = L'_i \setminus \{a \succ b\}$$

There are two cases.

In the first case, $g(\mathbf{P}, i) = \emptyset$. In this case, i must submit some ranking P'_i with $P_i \subseteq P'_i$. Simply suppose that $P'_i = L_i$.

In the second case, $g(\mathbf{P}, i) \neq \emptyset$. Since $P_{-i} = L+_{-i}$, we know there must be some $j \in N+$ with $P_j = L_i$. Since $P_i \subseteq L_i$, we have $P_i \subseteq P_j$. So we must have that $j \in g(\mathbf{P}, i)$, since g is preference monotonic. Simply suppose that i picks j as her proxy.

In either case, we end up with i 's guru submitting the vote L_i (note that i is her own guru in the first case). So we have that

$$\mathbf{P}_{\pi, S, D} = \mathbf{L}+$$

It follows that

$$f(\mathbf{P}_{\pi, S, D}) = f(\mathbf{L}+) = a$$

Consider now that profile $\mathbf{L}'+$. Note that it is exactly like \mathbf{P} except for the fact that

$$L'+_i = L_i = P_i \cup \{a \succ b\}$$

As noted above, we have that $f(\mathbf{L}'+) \neq a$. It follows that (f, g) violates proxy vote addition monotonicity. \square

3.2.4 A Proxy Vote Analogue of May's Theorem.

May's Theorem (May (1952)) is a well known result. When $|A| = 2$ and $|N|$ is odd, May shows that we can characterise the majority rule as the unique rule satisfying anonymity, neutrality and weak monotonicity. Here, I show that we can use the proxy vote analogues of these properties to achieve the same characterisation result.

Theorem 3.25. Suppose $|A| = 2$ and $|N|$ is odd. Then a pair (f, g) satisfies

- Proxy Vote Anonymity
- Proxy Vote Neutrality
- Proxy Vote Addition Monotonicity (PVAM), and
- Proxy Vote Deletion Monotonicity (PVDM)

iff f is the majority rule.

Proof. The left to right direction follows from Propositions 3.19, 3.20, 3.21 and May's Theorem.¹

For the other direction, suppose that f is the majority rule. So f is anonymous and neutral. Since $|A| = 2$, every proxy mechanism g will be anonymous and neutral (this is trivial to verify). By Lemmas 3.22 and 3.23, this implies that (f, g) is anonymous and neutral.

¹Note that we only need the requirement that $|N|$ is odd to ensure that the majority rule is resolute. We can drop this requirement if we replace weak monotonicity by its irresolute counterpart, positive responsiveness, and modify the definitions of anonymity and neutrality to accommodate irresoluteness.

It remains only to show that (f, g) satisfies PVAM and PVDM. I will write $A = \{a, b\}$. Suppose that for some proxy vote profile $(\mathbf{P}, \mathbf{S}, \mathbf{D})$, we have

$$f(\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}}) = a$$

Fix some $i \in N$. Then there are two cases to consider.

To see that (f, g) satisfies PVAM, suppose that $P_i = \emptyset$, then consider the case where $P'_i = \{a \succ b\}$. So i casts her own vote, meaning $P'_{\pi_i, S_i} = \{a \succ b\}$. Note that if j picked i as her proxy when $P_i = \emptyset$, then j must still pick i as her proxy (since this implies $P_j = \emptyset$, since $|A| = 2$). And we know that P_{π_i, S_i} was either $\{b \succ a\}$ or $\{a \succ b\}$.

To see that (f, g) satisfies PVDM, suppose that $P_i = \{b \succ a\}$, then consider the case where $P'_i = \emptyset$. So i delegates her vote, meaning P'_{π_i, S_i} is either $\{a \succ b\}$ or $\{b \succ a\}$. Note that if j didn't pick i as her proxy when $P_i = \{b \succ a\}$, then j must still not pick i as her proxy (since this implies either that $P_j \in \mathcal{L}(A)$, or that $P_j = \emptyset$ and j prefers some other voter in $N \setminus \{i\}$). And we know that P_{π_i, S_i} was $\{b \succ a\}$, since i cast her own vote.

In either case, changing from P_i to P'_i can only decrease the number of $\{b \succ a\}$ edges submitted in the profile $\mathbf{P}'_{\pi, \mathbf{S}, \mathbf{D}}$ from in the profile $\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}}$ and increase the number of $\{a \succ b\}$ edges submitted in the profile $\mathbf{P}'_{\pi, \mathbf{S}, \mathbf{D}}$ from in the profile $\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}}$. Since f is weakly monotonic, this implies that $f(\mathbf{P}'_{\pi, \mathbf{S}, \mathbf{D}}) = a$. So (f, g) satisfies PVAM and PVDM. \square

The reader might find it strange that Theorem 3.25 characterises a social choice function f in terms of properties of pairs (f, g) without making reference to a proxy mechanism g .

Of course, this is due to the fact that $|A| = 2$. When there are two alternatives, there is only a single (trivial) proxy mechanism. Voters can only submit linear orders or empty orders. If they submit an empty order, they can delegate their vote to any other agent. If they submit a linear order, they cast their own vote. So delegation plays a rather limited role in these proxy votes. And, as it happens, the sort of delegation that occurs does not violate PVAM and PVDM, assuming f is the majority rule.

3.2.5 Proxy Vote Monotonicity: An Impossibility Result

The reader will notice that we haven't yet explored the relationship between the proxy vote monotonicity properties and the proxy mechanism monotonicity property ('preference monotonicity').

In fact, we can show that, given some plausible restrictions on f and g , the monotonicity property of g is incompatible with monotonicity properties of the pair (f, g) . We have the following impossibility result.

Theorem 3.26. Suppose $|A| = 3$. Then, for every sufficiently large odd $|N|$, there is no pair (f, g) , where f is a social choice function and g is a proxy mechanism, such that:

- (f, g) satisfies *proxy vote addition monotonicity (PVAM)* and *proxy vote deletion monotonicity (PVDM)*
- f satisfies *anonymity* and *neutrality*
- g satisfies *preference monotonicity (PM)* and *independence of irrelevant proxies (IIP)*

Proof. Note that, since (f, g) satisfies PVAM and PVDM, we can assume without loss of generality that f is weakly monotonic (from Proposition 3.21). I proceed by means of three lemmas.

Lemma 3.27. Suppose that for some $i \in N$, for some $a, b \in A$, for some $P_{-i} \in \mathcal{P}(A)^{|N \setminus \{i\}|}$, if

$$P'_i = \{a \succ b\}$$

we have

$$g((P'_i, P_{-i}), i) \neq N \setminus \{i\}.$$

Then we can construct profiles $(\mathbf{Q}, \mathbf{S}, \mathbf{D})$ and $((Q'_i, Q_{-i}), \mathbf{S}, \mathbf{D})$, where

$$Q_i = \emptyset$$

and

$$Q'_i = \{a \succ b\}$$

such that

$$Q_{\pi_i, S_i, D_i} = \{b \succ a \succ c\}$$

and

$$Q'_{\pi'_i, S_i, D_i} = \{c \succ a \succ b\}$$

regardless of the behaviour of the set $N \setminus \{i, \pi_i, \pi'_i\}$ (that is, regardless of $(P_{-i, \pi_i, \pi'_i}, S_{-i, \pi_i, \pi'_i}, D_{-i, \pi_i, \pi'_i})$). In other words, we can construct a profile where the vote cast by i 's guru changes from $\{b \succ a \succ c\}$ to $\{c \succ a \succ b\}$ and we are free to specify the votes submitted by any voter who is not i 's guru.

Proof. Let $\mathbf{P} \in \mathcal{P}(A)^n$ be such that for some $i \in N$, for some $a, b \in A$ if

$$P'_i = \{a \succ b\}$$

we have, writing $\mathbf{P}' = (P'_i, P_{-i})$,

$$g(\mathbf{P}', i) \neq N \setminus \{i\}.$$

We construct the profiles $(\mathbf{Q}, \mathbf{S}, \mathbf{D})$ and $(\mathbf{Q}' = (Q'_i, Q_{-i}), \mathbf{S})$ as follows.

Let $Q'_i = \emptyset$ and $Q'_i = \{a \succ b\}$. We know that there must be some $k \in N \setminus \{i\}$ such that $k \notin g((Q'_i, P_{-i}), i)$. So let $Q_k = P_k$. Take some $j \in N \setminus \{i, k\}$, and let $Q_j = \{b \succ a \succ c\}$.

Note that since g satisfies IIP, we must have $k \notin g((Q'_i, Q_{j,k}, P_{-i,j,k}), i)$. Since g satisfies PM, this implies that $j \notin g((Q'_i, Q_{j,k}, P_{-i,j,k}), i)$, since otherwise we would have to have $k \in g((Q'_i, Q_{j,k}, P_{-i,j,k}), i)$ (because it would be impossible for Q_k to conflict with Q'_i more that Q_j does). Suppose that $S_i|_N = j$. Note that this implies that i picks j as her proxy in any profile of the form $((Q_i, Q_{-i}), \mathbf{S})$, since we know that $g(Q_i, Q_{-i}) = N \setminus \{i\}$ by definition (because g is a proxy mechanism).

Choose some $l \in N \setminus \{i, j\}$, and suppose that $S_i|_{N \setminus \{j\}} = l$ (i.e. that i picks l as her proxy if she is not allowed to pick j). Set $Q_l = \{c \succ a \succ b\}$. There are two cases.

For the first case, suppose $g((Q'_i, Q_{j,l}, P_{-i,j,l}), i) \neq \emptyset$. Then we must have $l \in g((Q'_i, Q_{j,l}, P_{-i,j,l}), i)$, since g satisfies PM. Then i delegates her vote to j (who is actually her guru, since $Q_j \in \mathcal{L}(A)$) in the profile $((Q_{i,j,l}, P_{-i,j,l}), \mathbf{S})$, and delegates her vote to l (who is actually her guru, since $Q_l \in \mathcal{L}(A)$) in the profile $((Q'_i, Q_{j,l}, P_{-i,j,l}), \mathbf{S})$.

For the second case, suppose $g((Q'_i, Q_{j,l}, P_{-i,j,l}), i) = \emptyset$. Then we can simply let $D_i = \{c \succ a \succ b\}$. Then i delegates her vote to j (who is actually her guru, since $Q_j \in \mathcal{L}(A)$) in the profile $((Q_{i,j,l}, P_{-i,j,l}), \mathbf{S})$, and casts her own vote in the profile $((Q'_i, Q_{j,l}, P_{-i,j,l}), \mathbf{S})$.

So, regardless of which case obtains, i 's vote is cast as $\{b \succ a \succ c\}$ in the profile $((Q_{i,j,l}, P_{-i,j,l}), \mathbf{S})$, and $\{c \succ a \succ b\}$ in the profile $((Q'_i, Q_{j,l}, P_{-i,j,l}), \mathbf{S})$, as desired. Note, crucially, that this is entirely independent of $(P_{-i,j,l}, S_{-i,j,l})$, since g is IIP and PM and we have fixed S_i . \square

Lemma 3.28. Suppose that for any $j \in N$, $a, b \in A$, $P_{-j} \in \mathcal{P}(A)^{|N \setminus \{j\}|}$, that if $P_j = \{a \succ b\}$, we have

$$g(\mathbf{P}) = N \setminus \{j\}$$

Suppose that for some $i \in N$, for some $a, b, c \in A$, for some $P_{-i} \in \mathcal{P}(A)^{|N \setminus \{i\}|}$, if

$$P'_i = \{a \succ b, c \succ b\}$$

we have

$$g((P'_i, P_{-i}), i) \neq N \setminus \{i\}.$$

Then we can construct profiles $(\mathbf{Q}, \mathbf{S}, \mathbf{D})$ and $((Q'_i, Q_{-i}), \mathbf{S}, \mathbf{D})$, where

$$Q_i = \{c \succ b\}$$

and

$$Q'_i = \{a \succ b, c \succ b\}$$

such that

$$Q_{\pi_i, S_i, D_i} = \{b \succ a \succ c\}$$

and

$$Q'_{\pi'_i, S_i, D_i} = \{c \succ a \succ b\}$$

regardless of the behaviour of the set $N \setminus \{i, \pi_i, \pi'_i\}$ (that is, regardless of $(P_{-i, \pi_i, \pi'_i}, S_{-i, \pi_i, \pi'_i}, D_{-i, \pi_i, \pi'_i})$). In other words, we can construct a profile where the vote cast by i 's guru changes from $\{b \succ a \succ c\}$ to $\{c \succ a \succ b\}$ and we are free to specify the votes submitted by any voter who is not i 's guru.

Proof. Suppose that for any $j \in N$, $a, b \in A$, $P_{-j} \in \mathcal{P}(A)^{|N \setminus \{j\}|}$, that if $P_j = \{a \succ b\}$, we have

$$g(\mathbf{P}) = N \setminus \{j\}.$$

Suppose that for some $i \in N$, for some $a, b, c \in A$, we have $P_{-i} \in \mathcal{P}(A)^{|N \setminus \{i\}|}$ such that if

$$P'_i = \{a \succ b, c \succ b\}$$

we have

$$g((P'_i, P_{-i}), i) \neq N \setminus \{i\}.$$

We construct \mathbf{Q} and $\mathbf{Q}' = (Q'_i, Q_{-i})$ as follows.

Let $Q_i = \{c \succ b\}$ and let $Q'_i = P'_i$. By assumption, we must have

$$g((Q_i, P_{-i}), i) = N \setminus \{i\}$$

We also know that

$$g((Q'_i, P_{-i}), i) \neq N \setminus \{i\}$$

So there must be $k \in g((Q_i, P_{-i}), i)$ such that $k \notin g((Q'_i, P_{-i}), i)$.

We can now essentially repeat the proof of Lemma 3.27. For $j \in N \setminus \{i, k\}$, set $Q_j = \{b \succ a \succ c\}$. For $l \in N \setminus \{i, j, k\}$, set $Q_l = \{c \succ a \succ b\}$. Set $S_i|_N = j$, $S_i|_{N \setminus \{j\}} = l$ and $D_i = \{c \succ a \succ b\}$ (this is allowed, since $Q'_i \subseteq \{c \succ a \succ b\}$). Then since $Q'_i \subseteq Q_l$, it is easy to see that $l \in g((Q'_i, Q_{j,l}, P_{-i,j,l}), i)$ if $g((Q'_i, Q_{j,l}, P_{-i,j,l}), i)$ is non-empty. Furthermore, by assumption, we know that it always holds that $j \in g((Q_i, Q_{j,l}, P_{-i,j,l}), i)$.

So, regardless of whether $g((Q'_i, Q_{j,l}, P_{-i,j,l}), i)$ is every non-empty, i 's vote is cast as $\{b \succ a \succ c\}$ in the profile $((Q_{i,j,l}, P_{-i,j,l}), \mathbf{S})$, and $\{c \succ a \succ b\}$ in the profile $((Q'_i, Q_{j,l}, P_{-i,j,l}), \mathbf{S})$, as desired. Note, crucially, that this is entirely independent of $(P_{-i,j,l}, S_{-i,j,l})$, since g is IIP and PM and we have fixed S_i . \square

Lemma 3.29. Suppose that for any $j \in N$, $a, b, c \in A$, $P_{-j} \in \mathcal{P}(A)^{|N \setminus \{j\}|}$, that if $P_j = \{a \succ c, b \succ c\}$, we have

$$g(\mathbf{P}) = N \setminus \{j\}$$

Then we can construct profiles $(\mathbf{Q}, \mathbf{S}, \mathbf{D})$ and $((Q'_i, Q_{-i}), \mathbf{S}, \mathbf{D})$, where

$$Q_i = \{b \succ a \succ c\}$$

and

$$Q'_i = \{a \succ c, b \succ c\}$$

(i.e. $Q'_i = Q_i - \{b \succ a\}$) such that

$$Q_{\pi_i, S_i, D_i} = \{b \succ a \succ c\}$$

and

$$Q'_{\pi'_i, S_i, D_i} = \{c \succ a \succ b\}$$

regardless of the behaviour of the set $N \setminus \{i, \pi_i, \pi'_i\}$ (that is, regardless of $(P_{-i, \pi_i, \pi'_i}, S_{-i, \pi_i, \pi'_i}, D_{-i, \pi_i, \pi'_i})$). In other words, we can construct a profile where the vote cast by i 's guru changes from $\{b \succ a \succ c\}$ to $\{c \succ a \succ b\}$ and we are free to specify the votes submitted by any voter who is not i 's guru.

Proof. Suppose that for any $j \in N$, $a, b, c \in A$, $P_{-j} \in \mathcal{P}(A)^{|N \setminus \{j\}|}$, that if $P_j = \{a \succ c, b \succ c\}$, we have

$$g(\mathbf{P}) = N \setminus \{j\}$$

Let $Q_i = \{b \succ a \succ c\}$ and $Q'_i = \{a \succ c, b \succ c\}$. Let $Q_j = \{c \succ a \succ b\}$ for some $j \in N \setminus \{i\}$. Let $S_i|_N = j$.

Since $Q_i \in \mathcal{L}(A)$, we know that i casts her own vote in any profile of the form $((Q_i, P_{-i}), (S_i, S_{-i}))$, regardless of P_{-i} and S_{-i} . So i will always cast the vote $b \succ a \succ c$ when she submits Q_i .

Since $Q'_i = \{a \succ c, b \succ c\}$, we must have (by assumption)

$$g((Q'_i, P_{-i}), i) = N \setminus \{i\}.$$

for any $P_{-i} \in \mathcal{P}(A)^{|N \setminus \{i\}|}$.

So, in particular

$$j \in g((Q'_i, Q_j, P_{-i, j}), i)$$

regardless of what $(P_{-i, j}, S_{-i, j})$ looks like. So when i submits (Q'_i, S_i) and j submits Q_j , i 's vote will always be case as $\{c \succ a \succ b\}$. \square

It might not be immediately clear how Lemmas 3.27, 3.28 and 3.29 will help us prove Theorem 3.26.

Note that in Lemmas 3.27 and 3.28 we construct profiles where a single voter adding an edge $\{a \succ b\}$ switches the vote that is cast on her behalf from $\{b \succ a \succ c\}$ to $\{c \succ a \succ b\}$. In Lemma 3.29, we construct profiles where a single voter removing an edge $\{b \succ a\}$ switches the vote that is cast on her behalf from $\{b \succ a \succ c\}$ to $\{c \succ a \succ b\}$.

Of course, each of these results is based on assumptions about the proxy mechanism g . The crucial point to observe is that the assumptions upon which these results rest collectively exhaust the available options for g . In Lemma 3.27, we assume that there is at least one case where a voter i submits a single edge $P_i = \{a \succ b\}$ and

$$g(\mathbf{P}, i) \neq N \setminus \{i\}$$

for some P_{-i} .

In Lemma 3.29, we assume that in every case where a voter i submits an edge of the form $P_i = a \succ c, b \succ c$, we have

$$g(\mathbf{P}, i) = N \setminus \{i\}$$

regardless of P_{-i} .

In Lemma 3.28, we assume that both of these assumptions are false. That is, we assume that in every case where a voter i submits an edge of the form $P_i = \{a \succ b\}$, we have

$$g(\mathbf{P}, i) = N \setminus \{i\}$$

regardless of P_{-i} , and we assume that there is at least one case where a voter i submits a single edge $P_i = a \succ c, b \succ c$ and

$$g(\mathbf{P}, i) \neq N \setminus \{i\}$$

for some P_{-i} .

What we have shown, then, is that given the conditions on g there must exist profiles where adding an edge $\{a \succ b\}$ or removing an edge $\{b \succ a\}$ switches a voter i 's final vote from $\{b \succ a \succ c\}$ to $\{c \succ a \succ b\}$. In these profiles, we require that the votes submitted by at most two other voters are fixed. In particular, we require that $Q_j = \{b \succ a \succ c\}$ and $Q_l = \{c \succ a \succ b\}$ for some $j, l \in N$. Crucially, though, we have shown that we are free to vary the votes of the other $|N \setminus \{i, j, l\}|$ voters as we wish (exploiting the fact that g is IIP) and that i 's final vote will still change in the constructed way.

We are now in a position to show that, given f satisfies anonymity, neutrality and weak monotonicity, this implies that (f, g) must fail to satisfy at least one of proxy vote addition monotonicity and proxy vote deletion monotonicity. Specifically, we need to fiddle with the votes submitted by the voters in the set $N \setminus \{i, j, l\}$ to construct profiles where a wins when i 's final vote is $\{b \succ a \succ c\}$, and c wins when i 's final vote is $\{c \succ a \succ b\}$.

But it's easy to see that this can be done, since f is anonymous, neutral and weakly monotonic. Just have one voter k vote for $\{b \succ c \succ a\}$, and have the other voters split their votes between $\{a \succ c \succ b\}$ and $\{c \succ a \succ b\}$, such that

$$\begin{aligned} |\{l' \in N \setminus \{i, j, k, l\} \mid P_{l'} = \{a \succ c \succ b\}\}| = \\ |\{l' \in N \setminus \{i, j, k, l\} \mid P_{l'} = \{c \succ a \succ b\}\}| + 1 \end{aligned}$$

Let us write $|\{l' \in N \setminus \{i, j, k, l\} \mid P_{l'} = \{c \succ a \succ b\}\}|$ as n' (note that $n' = \frac{1}{2}(|N \setminus \{i, j, k, l\}| - 1)$). Then when i 's guru votes for $\{b \succ a \succ c\}$ (i.e. i 's guru is j), there will be $n' + 1$ votes for $\{a \succ c \succ b\}$, $n' + 1$ votes for $\{c \succ a \succ b\}$, two votes for $\{b \succ a \succ c\}$ and one vote for $\{b \succ c \succ a\}$. When

i 's guru votes for $\{c \succ a \succ b\}$ (i.e. i 's guru is l), then there will be $n' + 1$ votes for $\{a \succ c \succ b\}$, $n' + 2$ votes for $\{c \succ a \succ b\}$, one vote for $\{b \succ a \succ c\}$ and one vote for $\{b \succ c \succ a\}$.

For sufficiently large $|N|$, the fact f is anonymous, neutral and weakly monotonic implies b cannot win in either case.² So either a or c must win in each case.

To see that a must win in the first case, suppose that c wins in the first case. Then suppose we change i 's vote from $\{b \succ a \succ c\}$ to $\{b \succ c \succ a\}$. By neutrality and anonymity, a would now have to win. So f would fail to be weakly monotonic.

To see that c must win in the second case, suppose that a wins in the second case. Then suppose we change i 's vote from $\{c \succ a \succ b\}$ to $\{a \succ c \succ b\}$. By neutrality and anonymity, c would now have to win. So f would fail to be weakly monotonic.

What we have seen, then, is that i will become the tiebreaker. If i 's guru votes for $\{b \succ a \succ c\}$, then a must win. If i 's guru votes for $\{c \succ a \succ b\}$, then c must win. So (f, g) must either fail to satisfy PVAM or fail to satisfy PVDM. \square

3.2.6 Discussion

What Theorem 3.26 shows is that the monotonicity properties of (f, g) are fundamentally incompatible with the monotonicity properties of g (assuming g is IIP). Essentially, this is because adding an edge $a \succ b$ (or deleting an edge $b \succ a$) can lead a voter to change guru, since proxy mechanisms are sensitive to changes in preferences over alternatives. This can mean that the voter's guru can end up submitting a vote which leads a to fail to win the election.

One might question Theorem 3.26 as follows. The result relies on the fact that f is anonymous and neutral, and we have been assuming that f is resolute. But it's well known that for certain values of $|N|$ and $|A|$ no resolute, anonymous and neutral social choice function exists (Herve Moulin (1983)). The worry is that we have just proven this result in a more complicated way.

There are two points to make here. The first is that Lemmas 3.27, 3.28 and 3.29 do not rely on any assumptions about f . So if f were irresolute and suitably non-trivial, and we were using some irresolute counterpart to the monotonicity properties at hand, then there's no reason to doubt that some version of Theorem 3.26 would still go through. Note that this would also allow us to remove the requirement that $|N|$ be odd.

The second (more important) point is that the proof applies to every $|N|$ of the form $|N| = 2n' + 5$ for some sufficiently large $n' \in \mathbb{N}$. In particular, for

²Note this is the only role the 'sufficiently large $|N|$ ' plays in the proof. It should be clear that this reasoning applies with even relatively small values of $|N|$.

example, the result will apply to every prime $|N| > 2$. Moulin shows, on the other hand, that anonymous, neutral and resolute social choice functions do exist when $|A|$ cannot be written as the sum of factors of $|N|$ greater than one. So the impossibility result here applies to pairs (N, A) for which there do exist anonymous, neutral and resolute social choice functions.

One final topic of discussion is the assumption that $|A| = 3$. Theorem 3.26 can be generalised to apply to arbitrary values of $|A| > 3$. To see this for $|A| = 4$, augment A with a dummy candidate d , and suppose that $a \succ d \in P_i$, for every $a \in A \setminus \{d\}$ and for every $i \in N$. Then the proof of Lemmas 3.27, 3.28 and 3.29 will be almost identical, since g is preference monotonic. We are guaranteed to end up with guru profiles where d is at the bottom of every linear order submitted. Since f is anonymous, neutral and weakly monotonic, we know that d cannot win the election. So a modified version of the second half of the proof will also go through.

Chapter 4

Manipulation and Candidate Control in Proxy Votes

In the previous chapter, I examined novel properties of proxy votes from an axiomatic perspective, presenting some results on their interaction with classical properties of social choice functions.

In this chapter, I want to give an example of the sort of analysis that can be done using my model of transitive proxy voting. Specifically, I will look at the topics of manipulation and control, which have hitherto received little attention in the literature on liquid democracy.

In the first section, I will define a novel form of manipulation (‘proxy choice manipulation’) and show that it occurs roughly as often as classical manipulation. I will then generalise classical manipulation to the proxy vote setting, and show that manipulation can occur strictly more often in proxy votes.

In the second section, I will discuss candidate control in proxy votes. After briefly generalising classical candidate control problems to the proxy vote setting, I will show that certain control immunity results fail when we allow proxy voting. I will then look at the (parameterized) complexity of candidate control. I will show that all hardness results carry over into the proxy vote setting. After defining a novel parameterization of a proxy vote, I will adapt an existing *FPT*-membership result to include a proxy vote control problem, and comment briefly on the significance of this result.

4.1 Manipulation

In a classical vote (N, A, f) , agents can manipulate by misrepresenting their preferences to achieve a better outcome. In the next subsection, I show that proxy voting increases the number of forms of manipulation.

4.1.1 Proxy Choice Manipulation

In a proxy vote (N, A, f, g) , there is an additional option for manipulation. Agents can manipulate by misrepresenting their choice of proxy (i.e. by picking one proxy over another for strategic reasons).

I call this sort of manipulation ‘proxy choice manipulation’. Note that in a proxy vote setting, manipulability is no longer a property of a social choice function f alone, but rather of a pair (f, g) .

Definition 4.1. (Proxy Choice Manipulation) A pair (f, g) is *proxy choice manipulable* (PC-manipulable) iff there exists $i \in N$, $\mathbf{P} \in \mathcal{P}(A)^n$, $\mathbf{S} \in \mathcal{L}(N)^n$, $\mathbf{D} \in \mathcal{L}(A)^n$ such that:

$$f(\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}}) \prec f(\mathbf{P}_{\pi, (\mathbf{S}'_i, \mathbf{S}_{-i}), \mathbf{D}}) \in P_i$$

for some $S_i, S'_i \in \mathcal{L}(N)$.

Intuitively, a pair (f, g) is PC-manipulable if there is a profile where an agent would prefer one of her potential proxies over another for purely strategic reasons.

A natural question to investigate is how PC-manipulability relates to the standard notion of manipulability, which I’ll call ‘Gibbard-Satterthwaite Manipulability’ (GS-manipulability). Recall that it is defined as follows.

Definition 4.2. (Gibbard-Satterthwaite Manipulation) A social choice function f is *Gibbard-Satterthwaite manipulable* (GS-manipulable) iff there exists $i \in N$, $\mathbf{P} \in \mathcal{L}(A)^n$ such that:

$$f(\mathbf{P}) \prec f((P'_i, P_{-i})) \in P_i$$

for some $P_i, P'_i \in \mathcal{L}(A)$.

One way of investigating the connection between PC-manipulability and GS-manipulability is to fix a particular proxy mechanism g .

In the previous chapter, I showed that SUBSET is the unique proxy mechanism satisfying zero regret, proxy availability, preference monotonicity and independence of irrelevant proxies. Since these properties are intuitively plausible requirements on g , let’s suppose that the proxy mechanism we are considering is SUBSET.

Theorem 4.3. If (f, SUBSET) is PC-manipulable for n voters and m alternatives, then f is GS-manipulable for n voters and m alternatives.

Proof. Suppose (f, SUBSET) is PC-manipulable for n voters and m alternatives. Then there is some preference profile \mathbf{P} , default profile \mathbf{D} , proxy choices of the voters S_{-i} and voters $i, j, k \in N$ where i strictly prefers the outcome of the vote when she picks k as her proxy to the outcome when she

picks j . So suppose $S_i|_N = j$ and $S'_i|_N = k$. Let $\mathbf{S} = (S_i, S_{-i})$ and define $\mathbf{S}' = (S'_i, S_{-i})$ (i.e. the profile that only differs by \mathbf{S} by changing S_i to S'_i).

If i herself has no proxies, then the proof is straightforward. Since

$$f(\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}}) \prec f(\mathbf{P}_{\pi, (S'_i, S_{-i}), \mathbf{D}}) \in P_i$$

we must have that

$$\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}} \neq \mathbf{P}_{\pi, (S'_i, S_{-i}), \mathbf{D}}$$

Since we are assuming that i is not a proxy, this implies that $f(\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}})$ and $f(\mathbf{P}_{\pi, (S'_i, S_{-i}), \mathbf{D}})$ differ only with regard to P_{π_i, S_i, D_i} . But then this implies that f is GS-manipulable (we need only consider a profile $\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}}$ where i switches from P_{π_i, S_i, D_i} to P'_{π_i, S'_i, D_i}).

If i herself has proxies, then the situation becomes more complicated.

Let

$$Proxy_i = \{j \in N \mid S_j|_{g(\mathbf{P}, j)} = i\}$$

be the set of voters selecting i as their proxy.

Since

$$f(\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}}) \succ_{P_i} f(\mathbf{P}_{\pi, (S'_i, S_{-i}), \mathbf{D}})$$

we must have that

$$\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}} \neq \mathbf{P}_{\pi, (S'_i, S_{-i}), \mathbf{D}}$$

It is easy to verify that this implies that $f(\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}})$ and $f(\mathbf{P}_{\pi, (S'_i, S_{-i}), \mathbf{D}})$ differ only with regard to P_{π_i} , for every $l \in Proxy_i$, and with regard to P_{π_i} (so for exactly $|Proxy_i| + 1$ voters).

Without loss of generality, suppose

$$f(\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}}) = b$$

and

$$f(\mathbf{P}_{\pi, (S'_i, S_{-i}), \mathbf{D}}) = a$$

By definition, this implies that $a \succ b \in P_i$. Since we are using the SUBSET mechanism, this implies that $a \succ b \in P_{\pi_i, S_i, D_i}$. Of course, this also implies that $a \succ b \in P_{\pi_l, S_l}$, for every $l \in Proxy_i$.

Suppose now that we move from the profile $\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}}$ towards the profile $\mathbf{P}_{\pi, (S'_i, S_{-i}), \mathbf{D}}$ by changing, for each $l \in Proxy_i$, P_{π_l, S_l} to P_{π_i, S'_i} , and finally by changing P_{π_i, S_i, D_i} to P'_{π_i, S'_i, D_i} .

We know that when we start, the social outcome is a . We know that when we have made all the changes, the social outcome is b . If the social outcome changes directly from a to b at some stage in the process, then we have a profile with respect to which f is GS-manipulable (since i would like to effect the change, and is capable of making it, since $P_{\pi_i} = P_{\pi_l}$, for every $l \in Proxy_i$). If the social outcome first changes to some $c \neq b \neq a$, then there are two cases. If $c \succ_{P_{\pi_i}} b$, then the same reasoning shows that f is

GS-manipulable. If $c \prec_{P_{\pi_i}} b$, then we can just carry on making the changes until the social outcome changes to b , then apply the same reasoning as above. It follows that f is GS-manipulable. \square

So we have shown that PC-manipulability implies GS-manipulability, assuming g is the SUBSET mechanism. To investigate the conditions under which GS-manipulability implies PC-manipulability, it is necessary to define an intermediate form of manipulability.

Definition 4.4. (IIA-Manipulability) A social choice function f is *IIA-manipulable* if there is some $\mathbf{L} \in \mathcal{L}(A)$ such that, for some $i \in N$, $L'_i \neq L_i$:

- $f(\mathbf{L}) = b$
- $f(L'_i, L_{-i}) = a$
- $a \succ b \in L_i$ and $a \succ b \in L'_i$

for some $a, b \in A$.


Intuitively, f is IIA-manipulable if an agent can reverse the social ranking of two alternatives whilst maintaining her personal ranking of the alternatives.¹ Clearly, f will be GS-manipulable if it is IIA-manipulable (since IIA-manipulability is just a special sort of GS-manipulability). We can also prove something about the converse.

Theorem 4.5. Suppose f is GS-manipulable and unanimous. Then f is IIA-manipulable.

Proof. I first prove the case with $N = \{i, j\}$ and $A = \{a, b, c\}$. Assume that f is GS-manipulable and unanimous. Assume, for reductio, that f is not IIA-manipulable.

Since f is GS-manipulable, there must be (without loss of generality) profiles $\mathbf{P} = (P_i, P_j)$ and $\mathbf{P}' = (P'_i, P_j)$ such that $a \succ b \in P_i$, $f(\mathbf{P}) = b$ and $f(\mathbf{P}') = a$.

So that means \mathbf{P} is partially described by one of these three cases:



¹IIA-manipulability is closely related to ‘one-way monotonicity’ (Sanver and W. S. Zwicker (2009)), which features in the preference reversal paradox (Peters (2017)). IIA-manipulability is a much weaker condition than one-way monotonicity, though. One-way monotonicity says that every example of GS-manipulability is an example of IIA-manipulability, whereas IIA-manipulability requires only that one example of GS-manipulability is an example of IIA-manipulability.

and \mathbf{P}' is partially described by one of these six cases:

(1)	<table style="border-collapse: collapse; margin: auto;"> <tr><td style="border: 1px solid black; padding: 2px;">i</td><td style="border: 1px solid black; padding: 2px;">j</td></tr> <tr><td style="border: 1px solid black; padding: 2px;">a</td><td style="border: 1px solid black; padding: 2px;"></td></tr> <tr><td style="border: 1px solid black; padding: 2px;">b</td><td style="border: 1px solid black; padding: 2px;"></td></tr> <tr><td style="border: 1px solid black; padding: 2px;">c</td><td style="border: 1px solid black; padding: 2px;"></td></tr> </table>	i	j	a		b		c		(2)	<table style="border-collapse: collapse; margin: auto;"> <tr><td style="border: 1px solid black; padding: 2px;">i</td><td style="border: 1px solid black; padding: 2px;">j</td></tr> <tr><td style="border: 1px solid black; padding: 2px;">a</td><td style="border: 1px solid black; padding: 2px;"></td></tr> <tr><td style="border: 1px solid black; padding: 2px;">c</td><td style="border: 1px solid black; padding: 2px;"></td></tr> <tr><td style="border: 1px solid black; padding: 2px;">b</td><td style="border: 1px solid black; padding: 2px;"></td></tr> </table>	i	j	a		c		b		(3)	<table style="border-collapse: collapse; margin: auto;"> <tr><td style="border: 1px solid black; padding: 2px;">i</td><td style="border: 1px solid black; padding: 2px;">j</td></tr> <tr><td style="border: 1px solid black; padding: 2px;">c</td><td style="border: 1px solid black; padding: 2px;"></td></tr> <tr><td style="border: 1px solid black; padding: 2px;">a</td><td style="border: 1px solid black; padding: 2px;"></td></tr> <tr><td style="border: 1px solid black; padding: 2px;">b</td><td style="border: 1px solid black; padding: 2px;"></td></tr> </table>	i	j	c		a		b	
i	j																												
a																													
b																													
c																													
i	j																												
a																													
c																													
b																													
i	j																												
c																													
a																													
b																													
(4)	<table style="border-collapse: collapse; margin: auto;"> <tr><td style="border: 1px solid black; padding: 2px;">i</td><td style="border: 1px solid black; padding: 2px;">j</td></tr> <tr><td style="border: 1px solid black; padding: 2px;">c</td><td style="border: 1px solid black; padding: 2px;"></td></tr> <tr><td style="border: 1px solid black; padding: 2px;">b</td><td style="border: 1px solid black; padding: 2px;"></td></tr> <tr><td style="border: 1px solid black; padding: 2px;">a</td><td style="border: 1px solid black; padding: 2px;"></td></tr> </table>	i	j	c		b		a		(5)	<table style="border-collapse: collapse; margin: auto;"> <tr><td style="border: 1px solid black; padding: 2px;">i</td><td style="border: 1px solid black; padding: 2px;">j</td></tr> <tr><td style="border: 1px solid black; padding: 2px;">b</td><td style="border: 1px solid black; padding: 2px;"></td></tr> <tr><td style="border: 1px solid black; padding: 2px;">a</td><td style="border: 1px solid black; padding: 2px;"></td></tr> <tr><td style="border: 1px solid black; padding: 2px;">c</td><td style="border: 1px solid black; padding: 2px;"></td></tr> </table>	i	j	b		a		c		(6)	<table style="border-collapse: collapse; margin: auto;"> <tr><td style="border: 1px solid black; padding: 2px;">i</td><td style="border: 1px solid black; padding: 2px;">j</td></tr> <tr><td style="border: 1px solid black; padding: 2px;">b</td><td style="border: 1px solid black; padding: 2px;"></td></tr> <tr><td style="border: 1px solid black; padding: 2px;">c</td><td style="border: 1px solid black; padding: 2px;"></td></tr> <tr><td style="border: 1px solid black; padding: 2px;">a</td><td style="border: 1px solid black; padding: 2px;"></td></tr> </table>	i	j	b		c		a	
i	j																												
c																													
b																													
a																													
i	j																												
b																													
a																													
c																													
i	j																												
b																													
c																													
a																													

Case 1: $\mathbf{P} = (1)$

By assumption, the winner in (1) is b . Since we are assuming that f is not IIA-manipulable, this implies that the winner in both (2) and (3) must not be a (since otherwise i could IIA-manipulate by switching from (1) to (2) or (3), respectively).

Since the winner in (1) is b and f is unanimous, this implies that j must not rank a as her top pick, since otherwise the winner would be a , by unanimity.

If j ranks b as her top pick, then b must win in profiles (5) and (6), by unanimity. This implies that $\mathbf{P}' = (4)$. So the winner in (4) is a . But then i could IIA-manipulate by switching from (e.g.) (4) to (5).

If j ranks c as her top pick, then c must win in profiles (3) and (4), by unanimity. So, by the reasoning above, a must win in either profile (5) or profile (6). If a wins in profile (6), then i can IIA-manipulate by moving from profile (6) to profile (4). If a wins in profile (5), then consider profile (2). We know the winner in profile (2) is not a . If the winner in profile (2) is b , then i could IIA-manipulate by switching from profile (2) to profile (3), since the winner in profile (3) is c . If the winner in profile (2) is c , then i could IIA-manipulate by switching from profile (2) to profile (5).

These sub-cases are exhaustive. It follows that $\mathbf{P} \neq (1)$.

Case 2: $\mathbf{P} = (2)$

By assumption, the winner in (2) is b . It follows that the winner in (1) and (3) cannot be a , since otherwise i could IIA-manipulate by switching from (2) to (1) or (3) respectively.

This also implies that j must not have a as her top pick, since otherwise the winner in (2) would be a , by unanimity.

If j has b as her top pick, then the winner in both (5) and (6) must be b , by unanimity. By our original assumption, this implies that the winner in (4) is a . But then i can IIA-manipulate by moving from (4) to (e.g.) (5).

If j has c as her top pick, then the winner in (4) must be c . But then i could IIA-manipulate by switching from (2) to (4), changing the outcome from b to c .

These sub-cases are exhaustive. It follows that $\mathbf{P} \neq (2)$.

Case 3: $\mathbf{P} = (3)$

By assumption, the winner in (3) is b . It follows that the winner in (1) and (2) cannot be a , since otherwise i could IIA-manipulate by switching from (3) to (1) or (2) respectively. This implies that j cannot have a as her top pick, since otherwise the winner in (1) and (2) would be a , by unanimity.

It also follows that j cannot have c as her top pick, since otherwise the winner in (3) would be c , by unanimity.

By the reasoning above, it follows that j must have b as her top pick. So b must be the winner in (5) and (6). So a must be the winner in (4), by our original assumption. But then i can IIA-manipulate by switching from (4) to (5) or (6).

These sub-cases are exhaustive. It follows that $\mathbf{P} \neq (3)$.

These cases are exhaustive. It follows that f must be IIA-manipulable. \square

We have the following bridging propositions, familiar from the literature on automating social choice theory.

Proposition 4.6. Suppose f is IIA-manipulable for n voters and m alternatives, with $n \geq 2$ and $m \geq 3$. Then f is IIA-manipulable for n voters and $m + 1$ alternatives.

Proof. See Tang and Lin (2008), Lemma 1. The proof works for IIA-manipulability. \square

Proposition 4.7. Suppose f is IIA-manipulable for n voters and m alternatives, with $n \geq 2$ and $m \geq 3$. Then f is IIA-manipulable for $n + 1$ voters and m alternatives.

Proof. See Tang and Lin (2008), Lemma 1. The proof works for IIA-manipulability. \square

Earlier, we saw that a social choice function f was GS-manipulable if the pair (f, SUBSET) was PC-manipulable. We are now ready to describe a converse relationship.

Recall that Uniform Voter Addition Invariance was defined as follows.

Let $\mathbf{L} \in \mathcal{L}(A)^n$ be arbitrary. Let $\mathbf{L}+$ be the profile we get when we augment \mathbf{L} with $|A|!$ new voters, one holding each possible ranking in $\mathcal{L}(A)$ (if f is anonymous, then we can think of $\mathbf{L}+$ as $\mathbf{L} \cup \mathcal{L}(A)$; otherwise, we must assume some ordering on the rankings in $\mathcal{L}(A)$). Then f is *Uniform Voter Addition Invariant* (UVAI) iff $f(\mathbf{L}+) = f(\mathbf{L})$.

Theorem 4.8. Let f be a social choice function. Suppose that f is:

- GS-manipulable over n voters and m alternatives.

- Unanimous
- Invariant to Uniform Voter Additions

Then (f, SUBSET) is PC-manipulable over $n + m!$ voters, and m alternatives.

Proof. From Theorem 4.5, we know that f is IIA-manipulable, for some profile $\mathbf{L} \in \mathcal{L}(A)^n$. So there must be some $i \in N$ such that, for some $L'_i \neq L_i$, we have

$$f(L'_i, L_{-i}) \succ_{L_i} f(\mathbf{L})$$

For the sake of readability, let $f(\mathbf{L}) = b$ and $f(L'_i, L_{-i}) = a$. Since f is IIA-manipulable, we can assume that both

$$a \succ_{L_i} b$$

and

$$a \succ_{L'_i} b$$

without loss of generality.

Let us now consider $\mathbf{L}+$ and $\mathbf{L}'+$, the uniform-voter augmentations of \mathbf{L} and (L'_i, L_{-i}) respectively. Since f is IUVA, it follows that

$$f(\mathbf{L}+) = b$$

and

$$f(\mathbf{L}'+) = a$$

Now define

$$P_i = a \succ b$$

Let (P_i, L_{+-i}) be the profile which is exactly like $\mathbf{L}+$, but where voter i submits the partial order P_i instead of her previous (linear) order L_i .

Since $\mathbf{L}+$ contains, for every linear order over A , at least one voter who submits that order, we must have voters j and k such that $L_{+j} = L_i$ and $L_{+k} = L'_i$.

Note, crucially, that both $P_i \subset L_{+j}$ and $P_i \subset L_{+k}$ (using the fact, observed above, that we are constructing this profile from an instance of IIA-manipulation). Since we are using the **SUBSET** proxy mechanism, this implies that both j and k are permitted proxies for i in the profile (P_i, L_{+-i}) .

If i picks j as her proxy, then the guru profile for (P_i, L_{+-i}) , written as $(P_i, L_{+-i})_{\pi, \mathbf{S}}$, is simply $\mathbf{L}+$. So we must have $f((P_i, L_{+-i})_{\pi, \mathbf{S}}) = b$.

If i picks k as her proxy, then the guru profile for (P_i, L_{+-i}) , written as $(P_i, L_{+-i})_{\pi, \mathbf{S}'}$, is simply $\mathbf{L}'+$. So we must have $f((P_i, L_{+-i})_{\pi, \mathbf{S}'}) = a$.

Since $a \succ_{P_i} b$, it follows that we have a situation where i would strictly prefer picking k over j as her proxy. So f is PC-manipulable on a profile of $n + m!$ voters. \square

4.1.2 Preference Misrepresentation Manipulation

PC-manipulation is a novel form of manipulation. It has no analogue in the classical setting. But it is also natural to generalise GS-manipulation in the proxy vote setting. I call this ‘Preference Misrepresentation Manipulation’ (PM-manipulation).

Definition 4.9. (Preference Misrepresentation Manipulation) A pair (f, g) is *preference misrepresentation manipulable* (PM-manipulable) iff there exists $i \in N$, $\mathbf{P} \in \mathcal{P}(A)^n$, $\mathbf{S} \in \mathcal{L}(N)^n$ such that:

$$f(\mathbf{P}_{\pi, \mathbf{S}, \mathbf{D}}) \prec f((P'_i, P_{-i})_{\pi, \mathbf{S}, \mathbf{D}}) \in P_i$$

for some $P_i, P'_i \in \mathcal{P}(A)$.

So PM-manipulability is just the generalisation of GS-manipulability to the proxy vote setting. It should therefore be no surprise that the following result holds.

Proposition 4.10. If f is GS-manipulable, then (f, g) is PM-manipulable, for any proxy mechanism g .

Proof. Trivial. We need only consider profiles where every agent votes directly. \square

Given such a large class of social choice functions is vulnerable to GS-manipulation, it is common to look for subdomains where the Gibbard-Satterthwaite theorem doesn’t hold. The most famous of these domain restrictions is single-peakedness (Black (1948)).

One might wonder whether these same domain restrictions also result in PM-strategyproofness. The following result shows that this does not hold. There are social choice functions f such that (f, SUBSET) is PM-manipulable on the domain of single-peaked preferences but f is not GS-manipulable on the domain of single-peaked preferences. So PM-strategyproofness is strictly more demanding a condition than GS-strategyproofness.

Theorem 4.11. When $|A| \geq 3$, there is no social choice function f which is non-dictatorial, surjective, and such that (f, SUBSET) is PM-strategyproof, even when we restrict the domain to include only single-peaked preference profiles.

Proof. We know that if a social choice function f is GS-manipulable, the pair (f, SUBSET) is PM-manipulable. By contraposition, it follows that if the pair (f, SUBSET) is PM-strategyproof, f is GS-strategyproof.

Moulin characterises the class of surjective, non-dictatorial and GS-strategyproof social choice functions on the domain of single-peaked preferences as the class of generalised median voter rules (H. Moulin (1980)).

To prove the theorem at hand, then, it suffices to show that for a generalised median voter rule f , (f, SUBSET) is PM-manipulable on the domain of single-peaked preferences when $|A| \geq 3$. For the sake of convenience, I focus on the case where $|A| = \{a, b, c\}$.

Let f be an arbitrary generalised median voter rule with $n-1$ phantoms. There are two cases.

Case 1: Every phantom has the same peak. Without loss of generality, suppose this is a . Consider the profile $\mathbf{P} = (P_1, \dots, P_n)$, where

$$\begin{aligned} P_j &= \{a \succ b\} & \forall j \in N \setminus \{i\} \\ P_i &= \{c \succ b \succ a\} \end{aligned}$$

Suppose also that $S_j|_{\emptyset} = \{a \succ c \succ b\}$ for every $j \in N \setminus \{i\}$. Note that (among other dimensions) \mathbf{P} is single-peaked along the dimension a, c, b .

As it stands, we have that $g(\mathbf{P}, j) = \emptyset$ for every $j \in N \setminus \{i\}$. It follows that each of these agents will cast her own vote as $\{a \succ b \succ c\}$. So the peak of the median voter will be a . So the winner will be a .

Now suppose i switches from $\{c \succ b \succ a\}$ to $\{c \succ a \succ b\}$. Note that the profile is still single-peaked along the dimension a, c, b . Now we have that $g(\mathbf{P}, j) = \{i\}$ for every $j \in N \setminus \{i\}$. It follows that each of the N voters has i as her guru. So the peak of the median voter will be c . So the winner will be c . Since $c \succ_{P_i} a$, it follows that i has an incentive to change her preference from $\{c \succ b \succ a\}$ to $\{c \succ a \succ b\}$.

Case 2: The phantoms have at least two distinct peaks. Suppose, without loss of generality, that at least one phantom has peak a and at least one phantom has peak b . We can then use exactly the same profiles as in the previous case to show that i can PM-manipulate. \square

It's worth briefly discussing the reasons why this result goes through. It may seem strange that a single agent can exert so much influence on the domain of single-peaked preferences, especially since Moulin's strategyproofness result applies to coalitional manipulation. The point is that in a proxy vote setting, an agent can acquire proxy votes from voters whose votes were previously cast on the other side of her 'peak' (because these voters are indifferent to which side of her peak their peak falls on). This marks the critical difference from a standard case of coalitional manipulation.

To make this clearer, it's worth identifying conditions under which the fact that (f, SUBSET) is PM-manipulable implies that f is GS-manipulable. It is easy to see that when the delegation graph is unaffected by a case of PM-manipulability, then we can construct a case of GS-manipulability.

Let $\mathbf{P} \in \mathcal{P}(A)^n$, let $\mathbf{S} \in \mathcal{L}(N)^n$, let $\mathbf{D} \in \mathcal{L}(A)^n$ and let $i \in N$. By $\text{Proxy}_{\mathbf{P}, \mathbf{S}, \mathbf{D}, i}$, I denote the set of agents whose vote 'flows through' i (i.e. the set of agents that are related to i in the delegation graph) under profile $(\mathbf{P}, \mathbf{S}, \mathbf{D})$.

Proposition 4.12. Suppose a pair (f, SUBSET) is PM-manipulable by agent i . So we have $(\mathbf{P}, \mathbf{S}, \mathbf{D})$ and $(\mathbf{P}', \mathbf{S}, \mathbf{D})$, where $\mathbf{P}' = (P'_i, P_{-i})$ and agent i manipulates by switching from P_i to P'_i . Suppose that

$$\text{Proxy}_{\mathbf{P}', \mathbf{S}, \mathbf{D}, i} = \text{Proxy}_{\mathbf{P}, \mathbf{S}, \mathbf{D}, i}$$

Then f is GS-manipulable.

Proof. Trivial. □

4.1.3 Discussion

In this section, I've investigated formal counterparts of classical notions of manipulation in a proxy vote setting. Manipulability in this setting becomes a property of pairs (f, g) . I've provided some initial exploration of the relationship between these novel forms of manipulation and classical manipulation.

PC-manipulation occurs when an agent chooses her proxies strategically. The fact that so many social choice functions are vulnerable to PC-manipulation (at least when paired with the SUBSET proxy mechanism) is not particularly surprising, but it does pose a problem for the diachronic interpretation of the model I've proposed in this thesis. Recall that in the diachronic interpretation, voters are presented with a set of their permitted proxies before they make their choice of proxy. Depending on what information is available to them, then, the results in this section show that they will often have the opportunity to act strategically. This worry has less force for the synchronic interpretation, since voters are taken to submit preferences and proxy choices simultaneously.

There are various avenues for future work to pursue. For example, one question concerns manipulation under partial information. Classically, this partial information is understood as partial knowledge of the preference profile submitted by agents (Reijngoud and Endriss (2012), Endriss et al. (2016)). In the proxy vote setting, there are more sources of information which could be restricted. For example, we could restrict information about the proxy choices of the other agents, or about their default values.

Another question concerns the structure of the delegation graph when PM-manipulation occurs. I have shown that if an agent i has an incentive to manipulate (f, SUBSET) when the set of voters choosing i as their proxy stays the same, then this implies that f is GS-manipulable. It would be interesting to explore the connection between PM-manipulation and GS-manipulation more. For example, it's easy to see that there are cases where an agent i has an incentive to manipulate, even when she reduces the number of voters choosing her as their proxy.

Relatedly, I've shown that single-peakedness is not sufficient for PM-strategyproofness. It would be interesting to look for non-trivial domain restrictions which do guarantee PM-strategyproofness.

4.2 Control

In this section, I present a straightforward generalisation of candidate control problems to the proxy vote setting. I show that hardness results carry over into the proxy vote setting, and adapt an existing *FPT*-membership result using a parameterization unique to the new setting.

4.2.1 Generalising Candidate-based Control Problems

Bartholdi, Tovey, and Trick (1992) introduce various control problems. I will focus on four candidate control problems here.

Constructive Control By Adding Candidates (CCAC)

Definition 4.13. (CCAC) In the classical CCAC problem, we have:

- An election (N, A, f) .
- B , a set of spoiler candidates.
- Some distinguished $a \in A$, interpreted as a preferred candidate.
- A bound $k \leq |N|$, the number of spoiler candidates we are allowed to add to the election.

The problem is to decide whether we can find some set of spoiler candidates $B' \subseteq B$ with $|B'| \leq k$ such that a wins the election $(N, A \cup B')$.

Constructive Control By Deleting Candidates (CCDC)

Definition 4.14. (CCDC) In the classical CCDC problem, we have:

- An election (N, A, f) .
- Some distinguished $a \in A$, interpreted as a preferred candidate.
- A bound $k \leq |N|$, the number of candidates we are allowed to remove from the election.

The problem is to decide whether we can find some set of $A' \subseteq A$ with $|A - A'| \leq k$ such that a wins the election (N, A') .

Destructive Control By Adding Candidates (DCAC)

Definition 4.15. (DCAC) In the classical DCAC problem, we have:

- An election (N, A, f) .
- B , a set of spoiler candidates.

- Some distinguished $a \in A$, interpreted as a disliked candidate.
- A bound $k \leq |N|$, the number of spoiler candidates we are allowed to add to the election.

The problem is to decide whether we can find some set of spoiler candidates $B' \subseteq B$ with $|B'| \leq k$ such that a does not win the election $(N, A \cup B')$.

Destructive Control By Deleting Candidates (DCDC)

Definition 4.16. (DCDC) In the classical DCDC problem, we have:

- An election (N, A, f) .
- Some distinguished $a \in A$, interpreted as a disliked candidate.
- A bound $k \leq |N|$, the number of candidates we are allowed to remove from the election.

The problem is to decide whether we can find some set of $A' \subseteq A \setminus \{a\}$ with $|A \setminus \{a\} - A'| \leq k$ such that a does not win the election $(N, A' \cup \{a\})$.

Note that each control problem is defined for a social choice function f . If, for a given control problem, there is some profile which yields a yes answer to the control problem, we say that f is *vulnerable* to control of the sort specified in the problem. Otherwise, we say f is *immune* to such control.

The proxy vote analogues of these classical control problems are defined by simply replacing a classical election (N, A, f) by a proxy vote (N, A, f, g) . I denote the proxy vote analogues of these classical control problems (respectively): PCCAC, PCCDC, PDCAC and PDCDC.

As is now familiar, when we move to a proxy vote setting, control problems are now defined relative to pairs (f, g) , where f is a social choice problem and g is a proxy mechanism. So we now speak of pairs (f, g) as being vulnerable or immune to control, rather than social choice functions alone.

As soon as we move to the proxy vote setting, differences can appear. Let *Condorcet* be some social choice function which outputs the unique Condorcet winner on profiles in which such a winner exists.

Proposition 4.17. (Bartholdi, Tovey, and Trick (1992)) *Condorcet* is immune to CCAC when we restrict the domain to profiles where there is a Condorcet winner.

Proposition 4.18. (Bartholdi, Tovey, and Trick (1992)) *Condorcet* is immune to DCDC when we restrict the domain to profiles where there is a Condorcet winner.

When we move from classical votes to proxy votes, these immunity results no longer hold, at least when we use the SUBSET proxy mechanism.

Proposition 4.19. (*Condorcet*, SUBSET) is vulnerable to PCCAC.

Proof. Suppose $N = \{1, 2, 3\}$, $A = \{a, b\}$ and $B = \{c, d\}$. I will define $\mathbf{P} = (P_1, P_2, P_3)$ as follows:

$$\begin{aligned} P_1 &= \{a \succ b \succ c \succ d\} \\ P_2 &= \{b \succ a \succ d \succ c\} \\ P_3 &= \{c \succ d\} \end{aligned}$$

In this election, 1 and 2 will cast their votes directly, whilst 3 will delegate. Suppose that

$$S_3 = \{2 \succ 1 \succ 3\}$$

When the set of candidates is A , 3 will delegate her vote to 2. So b will be the Condorcet winner. When the set of candidates is $A \cup B$, though, we have that

$$2 \notin \text{SUBSET}(\mathbf{P}, 3)$$

since $d \succ c \notin P_3$. So 3 will delegate her vote to 1, meaning a becomes the Condorcet winner. \square

Proposition 4.20. (*Condorcet*, SUBSET) is vulnerable to PDCDC.

Proof. Suppose $N = \{1, 2, 3\}$ and $A = \{a, b, c, d\}$. I will define $\mathbf{P} = (P_1, P_2, P_3)$ as follows:

$$\begin{aligned} P_1 &= \{a \succ b \succ c \succ d\} \\ P_2 &= \{b \succ a \succ d \succ c\} \\ P_3 &= \{c \succ d\} \end{aligned}$$

In this election, 1 and 2 will cast their votes directly, whilst 3 will delegate. Suppose that

$$S_3 = \{2 \succ 1 \succ 3\}$$

When the set of candidates is A , we have that

$$2 \notin \text{SUBSET}(\mathbf{P}, 3)$$

since $d \succ c \notin P_3$. So 3 will delegate her vote to 1, meaning a is the Condorcet winner.

When the set of candidates is $A \setminus \{c, d\}$, though, 3 will delegate her vote to 2. So a will no longer be the Condorcet winner. \square

The point is that, in a proxy vote setting, preferences over alternatives determine not only how voters vote directly, but also how they delegate their votes (since the proxy mechanism g takes into account the voters' preferences over alternatives). So a small change in the set of alternatives can result in a large change in the delegation graph.

4.2.2 Complexity of Candidate Control in the Proxy Vote Setting

Of course, most social choice functions are not immune to control. It's natural, then, to explore how difficult it is to control an election for a given social choice function. This difficulty is often cashed out in terms of computational complexity; intuitively, a social choice function is *resistant* to some control problem if there exists no efficient algorithm for deciding the problem. More formally, a social choice function is resistant to some control problem if the control problem is *NP*-hard (Faliszewski, Rothe, and Hervé Moulin (2016)).

Proposition 4.21. Suppose f is resistant to some control problem C . Then (f, g) is resistant to the proxy vote analogue of the control problem, PC , for any proxy mechanism g .

Proof. Consider some instance of the classical control problem. Note that it is also an instance of the proxy vote control problem (where everyone casts their own vote). It follows that the identity function constitutes a polynomial-time reduction from the classical control problem for f to the proxy vote control problem for (f, g) , regardless of how g is defined. So (f, g) must be resistant to the proxy vote analogue of the control problem, PC . \square

NP-hardness is a useful conceptual tool to analyse this sort of difficulty, but it is also blind to the structure of the input space. For example, the control problems we've outlined above have many different interacting features (to give a few examples: the number of voters, the number of candidates, and the number k of alternatives we are allowed to add/delete). When we move to a proxy voting setting, the number of such features only increases.

In the face of so much input structure, it's natural to wonder whether we can obtain a more fine-grained perspective on the difficulty of control. An attempt to do this formally comes from the field of parameterized complexity theory (Niedermeier (2006)).

Of course, the proof of Proposition 4.21 also shows that any parameterized hardness result will carry over to the proxy vote setting. So it's interesting to focus on membership results, using a parameterization specific to the proxy vote setting. In the following section, I'll present an *FPT*-membership result of this sort.

4.2.3 Parameterizing Problems with respect to the Number of Gurus: an *FPT*-Membership Result for PCCDC

Chen et al. (2015) show that if f is the plurality rule, CCDC is in *FPT* when we parameterize by $|N|$, the number of voters.

Theorem 4.22. (Chen et al. (2015)) CCDC using the plurality rule can be solved in $O(|A| \cdot |N| \cdot 2^{|N|})$ time (implying *FPT* membership when we parameterize by $|N|$, the number of voters in the election).

In this section, I'll prove a proxy vote analogue of this result, using a novel parameterization.

Let (N, A, f, g) be a proxy vote. Consider the control problem PCCDC, with k as the (maximum) number of candidates that can be deleted.

Define $LIN \subseteq N$ as follows:

$$LIN = \{i \in N \mid i \text{ submits a linear order over some } B \subseteq A \text{ alternatives} \\ \text{with } |B| = |A| - k\}$$

Define $SOLO \subseteq N$ as follows:

$$SOLO = \{i \in N \mid i \text{ has no proxies in the vote } (N, A)\}$$

Define $N' = LIN \cup SOLO$. Before presenting the main result, it is necessary to present some intermediate propositions. I write $\mathbf{P}|_B$ to denote the restriction of the preference profile \mathbf{P} to the candidate set B .

Proposition 4.23. Suppose $\text{SUBSET}(\mathbf{P}|_B, i) = \emptyset$ in some vote (N, B) , where $B \subseteq A$. Then $i \in SOLO$.

Proof. Since $\text{SUBSET}(\mathbf{P}|_B, i) = \emptyset$, we must have by definition that

$$\{j \in N \setminus \{i\} \mid P_i|_B \subseteq P_j|_B\} = \emptyset$$

It follows there must be some $a, b \in B$ such that $\{a \succ b\} \in P_i|_B$ and $\{a \succ b\} \notin P_j|_B$ for every $j \in N \setminus \{i\}$. But since $B \subseteq A$, it follows that $a, b \in A$. So then

$$\{j \in N \setminus \{i\} \mid P_i \subseteq P_j\} = \emptyset$$

It follows that $\text{SUBSET}(\mathbf{P}, i) = \emptyset$. So $i \in SOLO$. □

Proposition 4.24. Suppose $\text{SUBSET}(\mathbf{P}|_B, i) = \{i\}$ in some vote (N, B) , where $B \subseteq A$ with $|B| \geq |A| - k$. Then $i \in LIN$.

Proof. Follows directly from the definition of LIN , given that SUBSET is a proxy mechanism. □

From Propositions 4.23 and 4.24, we can deduce the following result.

Theorem 4.25. Suppose that $i \in N$ is a guru for some $j \in N$ (possibly with $i = j$) in the vote (N, B) , where $B \subseteq A$ with $|B| \geq |A| - k$. Then we must have that $i \in N'$.

Proof. Since i is a guru for some $j \in N$, we must have that i submits her own vote in the vote (N, B) . So we must either have that:

$$\text{SUBSET}(\mathbf{P}|_B, i) = \emptyset$$

or that

$$\text{SUBSET}(\mathbf{P}|_B, i) = \{i\}$$

From Proposition 4.23, the first case implies that $i \in \text{SOLO}$. From Proposition 4.24, the second case implies that $i \in \text{LIN}$. The two cases are exhaustive, so it follows that $i \in N'$. \square

We are now ready to present the main result in this section, namely that the ‘hardness’ in CCDC is localised to the size of N' when f is the plurality rule and g is the SUBSET mechanism. Our proof is very similar to the proof of Chen et al.’s Proposition 4.22.

Theorem 4.26. When f is the plurality rule and g is the SUBSET mechanism, PCCDC is in FPT with respect to $|N'| = n'$.

Proof. Standardly, it is assumed a polynomial length certificate for a control problem is simply a list of candidates which should be deleted (added). But it’s also possible to take a more voter-centric approach when specifying certificated for control problems. Instead of specifying candidates which should be deleted from the election as a whole, we could also specify for each voter the vote they end up submitting. The following proof of Theorem 4.26 exploits this fact.

We consider each of the $2^{n'}$ subsets of N' , one at a time. For each considered subset $H \subseteq N'$, we do the following. For each $i \in H$, we delete every candidate $b \in A$ such that $\{b \succ a\} \in P_i$ (recall that $a \in A$ was our preferred candidate). This means that there is no voter $i \in H$ who ranks a below some other alternative in the ranking she submits.

At this point, we establish who the current winner in this election. If it is some $c \in A \setminus \{a\}$, then we delete $c \in A$ and repeat the winner determination until a is the winner.

If at any stage in this process we have deleted more than k candidates, we reject. If the process terminates for some particular $H \subseteq N'$ and we have not deleted more than k candidates, then we accept.

Suppose our algorithm leads to an acceptance. Then it follows that there must exist a solution for this instance of the PCCDC control problem, since we will actually have found some particular solution.

Suppose there is a solution for the PCCDC control problem. Then it follows that there must be some $A' \subseteq A$ with $|A - A'| \leq k$ such that a wins the election (N, A') .

But then some subset $H \subseteq N$ of the gurus in the election (N, A') must rank a first in their ranking (since the gurus are the only voters who actually

end up casting votes). So, by Theorem 4.25, we must have that $H \subseteq N'$ ranks a first in their ranking. So the algorithm will find a solution when it considers this subset $H \subseteq N'$.

It remains only to note the running time of this algorithm. We consider $2^{|N'|=n'}$ subsets, and each subset can clearly be evaluated in polynomial time with respect to $|N| = n$ and $|A| = m$. It follows that PCCDC is in FPT with parameter $|N'|$, when f is the plurality rule and g is the SUBSET mechanism. \square

4.2.4 Discussion

Theorem 4.26 says that when very few voters have made their mind up about a significant number of the options, there exist relatively tractable algorithms to solve PCCDC for (*Plurality*, SUBSET).

Of course, in the case where everyone submits a linear order (or close to a linear order), there is no difference between Theorem 4.26 and Proposition 4.22. The difference between Theorem 4.26 and Proposition 4.22 is when there is a large number of voters of whom very few have made their minds up about (nearly) all of the alternatives. Note, though, that this is exactly the sort of scenario which is used to motivate transitive proxy voting. From the preliminary results in this section, then, it appears that there is strictly more potential for candidate control in a proxy vote setting.

Chapter 5

Conclusion and Future Work

In this thesis, I introduced a novel model of transitive proxy voting. I argued that extant models of transitive proxy voting paid insufficient attention to ‘proxy selection’, the process by which voters select delegates.

The model I proposed featured a two-dimensional analysis of proxy selection. First, a set of permitted proxies was formed for each voter by a ‘proxy mechanism’. The formation of these sets depended only on the preferences over alternatives submitted by the voters. Second, some delegate was chosen from within the set of permitted proxies according to a ranking submitted by the voter. I have deliberately stayed silent on the origins of this ranking, in recognition of the myriad of factors that can inform preferences over potential proxies.

After introducing the model, I explored some of its properties from an axiomatic perspective. The principal result in this section was an impossibility result. I showed that (given plausible assumptions) we cannot expect proxy votes to satisfy intuitively desirable monotonicity properties. I take it that this result does pose a challenge for proponents of liquid democracy. It serves to highlight the instability of any transitive proxy voting system. Small changes in preferences (either over alternatives or proxies) can lead to unexpected effects in the result of the vote as a whole. Future work could assess the implications of this result more thoroughly (for example, I have not shown that each assumption is individually necessary for the proof to go through).

In the final chapter of the thesis, I put the model to work in analysing manipulation and control in a proxy vote setting. I showed that not only do novel forms of manipulation arise in a proxy vote setting, but also that there are strictly more situations in which classical manipulation is available to agents in a proxy vote than in a classical vote. I also showed that certain candidate control immunity results fail when we allow proxy voting, and that there exist scenarios in which candidate control of proxy votes is strictly easier than candidate control of classical votes.

The results I've presented in this thesis do not begin to exhaust the available research directions using a computational social choice-theoretic model of transitive proxy voting. My aim has not been to provide a full coverage of topics relevant to transitive proxy voting, but rather to showcase interesting features of the model I've introduced. My hope is that the reader thinks my model sufficiently rich to enable non-trivial formal discussion of the merits of liquid democracy.

During individual chapters, I've given some sense of topic-specific future research paths. Here, I want briefly to sketch some areas which future work could engage with. A non-exhaustive list:

- In Chapter 2, I sketched two ways in which a social network could be accommodated within my model. It would be interesting to pursue this idea. For example, most properties of proxy mechanisms will need to be modified (e.g. anonymity, neutrality) when we build a social network into the proxy mechanism. Do the results I've presented in the rest of the thesis still hold when such modifications come into effect?
- Throughout the thesis, I've focused on elections using a social choice function. But, as I noted above, we could also use a social welfare function in the model.
- In Chapter 3, I gave a few examples of properties of proxy mechanisms, and characterised the SUBSET mechanism using these. But there are many interesting properties, and many plausible proxy mechanisms, which I have not discussed. It would be interesting to explore further the relationships between these properties and mechanisms. Are there specific voting scenarios in which we would prefer one proxy mechanism over another?
- Similarly, a motivation for transitive proxy voting is that it purportedly increases voter turnout. But I have not examined participation properties in this thesis. It would be interesting to extend the model to allow voters to abstain.
- In my model, voters submit partial orders over alternatives, linear orders over voters and default votes. This suggests a novel 'possible winner problem'. Given a partial preference profile, how many possible winners are there when we are free to fill out the proxy choice profile as we wish? How does the complexity of classical possible winner problems change in this setting? How does the choice of proxy mechanism affect this?
- In Chapter 4, I examined manipulation and control problems. But there are other forms of strategic behaviour that are relevant to as-

sessing liquid democracy. For example, it would be interesting to look at voter control problems and bribery problems.

To conclude, liquid democracy appears to be a promising means for collective decision-making, especially with the proliferation of technology in society. Before it achieves widespread adoption, though, it is necessary to scrutinise the claims offered by its proponents and detractors. One form of scrutiny is technical: how do formal models of liquid democracy behave? By writing this thesis, I hope to have contributed a formal tool for conducting such analysis.

Chapter 6

References

- Abramowitz, Ben and Nicholas Mattei (2019). “Flexible Representative Democracy: An Introduction with Binary Issues”. In: *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*. International Joint Conferences on Artificial Intelligence Organization, pp. 3–10. DOI: 10.24963/ijcai.2019/1. URL: <https://doi.org/10.24963/ijcai.2019/1>.
- Alger, Dan (2006). “Voting by proxy”. In: *Public Choice* 126 (1). URL: <https://doi.org/10.1007/s11127-006-3059-1>.
- Arrow, Kenneth (1950). “A Difficulty in the Concept of Social Welfare”. In: *Journal of Political Economy* 58. URL: <https://EconPapers.repec.org/RePEc:ucp:jpolec:v:58:y:1950:p:328>.
- Bartholdi, John J., Craig A. Tovey, and Michael A. Trick (1992). “How hard is it to control an election?” In: *Mathematical and Computer Modelling* 16.8, pp. 27–40. ISSN: 0895-7177. DOI: [https://doi.org/10.1016/0895-7177\(92\)90085-Y](https://doi.org/10.1016/0895-7177(92)90085-Y). URL: <http://www.sciencedirect.com/science/article/pii/089571779290085Y>.
- Behrens, Jan (2017). “The Origins of Liquid Democracy”. In: *The Liquid Democracy Journal* 5.
- Black, Duncan (1948). “On the Rationale of Group Decision-Making”. In: *Journal of Political Economy* 56, pp. 23–34.
- Bloembergen, Daan, Davide Grossi, and Martin Lackner (2019). “On Rational Delegations in Liquid Democracy”. In: *Association for the Advancement of Artificial Intelligence, AAAI 2019*.
- Blum, Christian and Christina Isabel Zuber (2016). “Liquid Democracy: Potentials, Problems, and Perspectives”. In: *Journal of Political Philosophy* 24.2, pp. 162–182. DOI: 10.1111/jopp.12065. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/jopp.12065>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/jopp.12065>.
- Boldi, Paolo, Francesco Bonchi, Carlos Castillo, and Sebastiano Vigna (2011). “Viscous Democracy for Social Networks”. In: *Commun. ACM* 54.6,

- pp. 129–137. ISSN: 0001-0782. DOI: 10.1145/1953122.1953154. URL: <http://doi.acm.org/10.1145/1953122.1953154>.
- Brandt, Felix, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (2016). *Handbook of Computational Social Choice*. 1st. New York, NY, USA: Cambridge University Press. ISBN: 1107060435, 9781107060432.
- Brill, Markus and Nimrod Talmon (2018). “Pairwise Liquid Democracy”. In: IJCAI’18, pp. 137–143. URL: <http://dl.acm.org/citation.cfm?id=3304415.3304436>.
- Chen, Jiehua, Piotr Faliszewski, Rolf Niedermeier, and Nimrod Talmon (2015). “Elections with Few Voters: Candidate Control Can Be Easy”. In: AAAI’15, pp. 2045–2051. URL: <http://dl.acm.org/citation.cfm?id=2886521.2886605>.
- Christoff, Zoé and Davide Grossi (2017). “Binary Voting with Delegable Proxy: An Analysis of Liquid Democracy”. In: pp. 134–150. DOI: 10.4204/EPTCS.251.10. URL: <https://doi.org/10.4204/EPTCS.251.10>.
- Cohensius, Gal, Shie Mannor, Reshef Meir, Eli Meir, and Ariel Orda (2017). “Proxy Voting for Better Outcomes”. In: AAMAS ’17, pp. 858–866. URL: <http://dl.acm.org/citation.cfm?id=3091125.3091247>.
- Endriss, Ulle, Svetlana Obraztsova, Maria Polukarov, and Jeffrey S. Rosenschein (2016). “Strategic Voting with Incomplete Information”. In: IJCAI’16, pp. 236–242. URL: <http://dl.acm.org/citation.cfm?id=3060621.3060655>.
- Escoffier, Bruno, Hugo Gilbert, and Adèle Pass-Lanneau (2018). “Iterative Delegations in Liquid Democracy with Restricted Preferences”. In: *MS (ArXiv)*.
- Escoffier, Bruno, Hugo Gilbert, and Adèle Pass-Lanneau (2018). “The Convergence of Iterative Delegations in Liquid Democracy”. In: *CoRR* abs/1809.04362. arXiv: 1809.04362. URL: <http://arxiv.org/abs/1809.04362>.
- Faliszewski, Piotr, Jörg Rothe, and Hervé Moulin (2016). “Control and Bribery in Voting”. In: *Handbook of Computational Social Choice*. Ed. by Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. Cambridge University Press, pp. 146–168. DOI: 10.1017/CB09781107446984.008.
- Gibbard, Allan (1973). “Manipulation of Voting Schemes: A General Result”. In: *Econometrica* 41.4, pp. 587–601. ISSN: 00129682, 14680262. URL: <http://www.jstor.org/stable/1914083>.
- Gölz, Paul, Anson Kahng, Simon Mackenzie, and Ariel D. Procaccia (2018). “The Fluid Mechanics of Liquid Democracy”. In: *Proceedings of WINE: Web and Internet Economics*.
- Green-Armytage, James (2015). “Direct voting and proxy voting”. In: *Constitutional Political Economy* 26.2, pp. 190–220. URL: <https://EconPapers.repec.org/RePEc:kap:copoec:v:26:y:2015:i:2:p:190-220>.

- Kang, Anson, Simon Mackenzie, and Ariel Procaccia (2018). “Liquid democracy: An Algorithmic Perspective”. In: *Proceedings of 32nd AAAI Conference on Artificial Intelligence*.
- Kling, Christoph Carl, Jérôme Kunegis, Heinrich Hartmann, Markus Strohmaier, and Steffen Staab (2015). “Voting behaviour and power in online democracy”. In: *Proceedings of the 9th International AAAI Conference on Web and Social Media (ICWSM)*.
- Kotsialou, Grammateia and Luke Riley (2018). “Incentivising Participation in Liquid Democracy with Breadth First Delegation”. In: *MS (ArXiv)*.
- Litvinenko, Anna (2012). “Social media and perspectives of liquid democracy: The example of political communication in the pirate party in Germany”. English. In: *Proceedings of the European Conference on e-Government, ECEG*, pp. 403–407. ISBN: 9781908272416.
- May, Kenneth O. (1952). “A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decision”. In: vol. 20. 4. [Wiley, Econometric Society], pp. 680–684. URL: <http://www.jstor.org/stable/1907651>.
- Miller, James C. (1969). “A program for direct and proxy voting in the legislative process”. In: *Public Choice* 7.1, pp. 107–113. ISSN: 1573-7101. DOI: 10.1007/BF01718736. URL: <https://doi.org/10.1007/BF01718736>.
- Moulin, H. (1980). “On strategy-proofness and single peakedness”. In: *Public Choice* 35.4, pp. 437–455. ISSN: 1573-7101. DOI: 10.1007/BF00128122. URL: <https://doi.org/10.1007/BF00128122>.
- Moulin, Herve (1983). *The strategy of social choice*. Advanced textbooks in economics. North-Holland Publishing Company. ISBN: 9780444863713. URL: <https://books.google.co.uk/books?id=iBzGAAAAIAAJ>.
- Niedermeier, Rolf (2006). *Invitation to Fixed-parameter Algorithms*. Oxford Lecture Series in Mathematics and its Applications. Oxford University Press.
- Peters, Dominik (2017). “Condorcet’s Principle and the Preference Reversal Paradox”. In: pp. 455–469. DOI: 10.4204/EPTCS.251.34. URL: <https://doi.org/10.4204/EPTCS.251.34>.
- Reijngoud, Annemieke and Ulle Endriss (2012). “Voter Response to Iterated Poll Information”. In: AAMAS ’12, pp. 635–644. URL: <http://dl.acm.org/citation.cfm?id=2343776.2343787>.
- Sanver, M. Remzi and William S. Zwicker (2009). “One-way monotonicity as a form of strategy-proofness”. In: *International Journal of Game Theory* 38.4, pp. 553–574. ISSN: 1432-1270. DOI: 10.1007/s00182-009-0170-9. URL: <https://doi.org/10.1007/s00182-009-0170-9>.
- Satterthwaite, Mark Allen (1975). “Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions”. In: *Journal of Economic Theory* 10.2, pp. 187–217. ISSN: 0022-0531. DOI: [https://doi.org/10.1016/0022-0531\(75](https://doi.org/10.1016/0022-0531(75)

- 90050-2. URL: <http://www.sciencedirect.com/science/article/pii/S0022053175900502>.
- Tang, Pingzhong and Fangzhen Lin (2008). “A computer-aided proof to Gibbard–Satterthwaite theorem”. In: *Technical report*. URL: http://iiis.tsinghua.edu.cn/~kenshin/GS_proof.
- Zhang, Bingsheng and Hong-sheng Zhou (2017). “Brief Announcement: Statement Voting and Liquid Democracy”. In: *Proceedings of the 36th ACM Symposium on Principles of Distributed Computing (PODC)*.
- Zwicker, William and Herve Moulin (2016). “Introduction to the Theory of Voting”. In: *Handbook of Computational Social Choice*. Ed. by Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D.Editors Procaccia. Cambridge University Press, pp. 23–56. DOI: 10.1017/CB09781107446984.003.