Eliciting and Aggregating Information for Better Decision Making

by

Rupert Freeman

Department of Computer Science
Duke University

Date: _____
Approved:

_____
Vincent Conitzer, Supervisor

_____
Edith Elkind

_____
Kamesh Munagala

_____
Ronald Parr

_____
Aleksandar Sasa Pekec

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in the Department of Computer Science
in the Graduate School of Duke University
2018

ABSTRACT

Eliciting and Aggregating Information for Better Decision Making

by

Rupert Freeman

Department of Computer Science
Duke University

Date: _____
Approved:

_____
Vincent Conitzer, Supervisor

_____
Edith Elkind

_____
Kamesh Munagala

_____
Ronald Parr

_____
Aleksandar Sasa Pekec

An abstract of a dissertation submitted in partial fulfillment of the requirements for
the degree of Doctor of Philosophy in the Department of Computer Science
in the Graduate School of Duke University
2018

Abstract

In this thesis, we consider two classes of problems where algorithms are increasingly used to make, or assist in making, a wide range of decisions. The first class of problems we consider is the allocation of jointly owned resources among a group of agents, and the second is the elicitation and aggregation of probabilistic forecasts from agents regarding future events. Solutions to these problems must trade off between many competing objectives including economic efficiency, fairness between participants, and strategic concerns.

In the first part of the thesis, we consider shared resource allocation, where we relax two common assumptions in the fair divison literature. Firstly, we relax the assumption that goods are private, meaning that they must be allocated to only a single agent, and introduce a more general public decision making model. This allows us to incorporate ideas and techniques from fair division to define novel fairness notions in the public decisions setting. Second, we relax the assumption that decisions are made offline, and instead consider online decisions. In this setting, we are forced to make decisions based on limited information, while seeking to retain fairness and game-theoretic desiderata.

In the second part of the thesis, we consider the design of mechanisms for forecasting. We first consider a tradeoff between several desirable properties for wagering mechanisms, showing that the properties of Pareto efficiency, incentive compatibility, budget balance, and individual rationality are incompatible with one another.

We propose two compromise solutions by relaxing either Pareto efficiency or incentive compatibility. Next, we consider the design of decentralized prediction markets, which are defined by the lack of any single trusted authority. As a consequence, markets must be closed by popular vote amongst a group of anonymous, untrusted arbiters. We design a mechanism that incentivizes arbiters to truthfully report their information even when they have a (possibly conflicting) stake in the market themselves.

Contents

List of Tables

List of Figures

Acknowledgements

There are many people I have to thank for a fun and productive last five years; too many to thank by name here. But I will do my best. In approximate chronological order from first appearance in my life...

Thanks to my family and friends in New Zealand for their love and support. I'm lucky to have you all. In particular to Mum, Dad, and Issy, who have only ever encouraged me in my chosen pursuits.

For setting me down this path, and making me feel welcome in an academic community, thanks to some wonderful people at the University of Auckland: Steven Galbraith, Mark Wilson, and especially Arkadii Slinko.

My time at Duke has been better than I could have imagined. To the extent that this thesis has been successful, I have Vince Conitzer to thank. Vince has been a fantastic advisor. As well as a brilliant researcher, Vince has been an excellent mentor and constant source of advice. It says a lot that I have always felt like my success is Vince's higest priority, even when he has so many other things going on.

Thanks also to the other members of my committee – Kamesh Munagala, Ron Parr, Sasa Pekec, and Edith Elkind – for their help in many forms over the years. I (quite literally) couldn't have a Ph.D. without you! Thanks too to Markus Brill for acting as an informal co-advisor for two years, and Nisarg Shah for teaching me a lot about doing research over the years.

Of course, thank you to my collaborators without whom none of this work would

exist: Vince Conitzer, Haris Aziz, Markus Brill, Shaddin Dughmi, Edith Elkind, Sam Haney, Svante Janson, Andreas Krause, Martin Lackner, Sébastien Lahaie, Ben Lee, Yuqian Li, Debmalya Panigrahi, David Pennock, Nisarg Shah, Milind Tambe, Jenn Wortman Vaughan, Toby Walsh, Jens Witkowski, Haifeng Xu, and Seyed Majid Zahedi.

Thanks to my friends at Duke, and especially Nat, Sam, and Andrew, for making my time in Durham a lot of fun. And for teaching me the difference between beers and bears.

Thanks to the Duke CS staff who made me feel welcome right from Day 1. Special thanks to Pam for keeping me fed, Kathleen and Alison for keeping me reimbursed, and Marilyn for keeping my entire life in order. That you do the same for 100 other grad students defies belief.

I spent two wonderful and productive summers at Microsoft Research in New York City. Thanks to Dave Pennock, Jenn Wortman Vaughan, and Sébastien Lahaie for hosting me. I'm excited to be coming back for another year, and that's because of you guys.

A huge thanks to Facebook for the fellowship support for my final year!

Finally, thanks to Amy for all the sacrifices you've made for me. I'm glad I went to Campout in 2013.

1

Introduction

Increasingly, algorithmic platforms are being utilized to make decisions based on the private information of heterogeneous agents. On one hand, algorithmic solutions are available to guide a group of individuals towards a mutually beneficial outcome in situations where more ad-hoc approaches may be suboptimal. Examples include dividing a set of items from, say, an inheritance [95, 9], or deciding on a location for a group lunch [6, 8]. On the other hand, some algorithms capture information from individuals with the goal of informing the decision of some central agent, or principal. These might include using prediction markets [5, 7] to harness wisdom-of-crowd effects for evaluating policy decisions, revenue maximization on e-commerce sites like eBay or Amazon, or matching kidney donors to patients [144, 145].

These platforms leverage research in economics, theoretical computer science, and artificial intelligence to produce good solutions. The exact nature of a 'good' solution is not always clear, but there are certain guiding principles. Does the solution provide high utility to the participants (that is, is it economically efficient)? Is utility distributed fairly? Is the system user-friendly? Can a solution by computed in a reasonable amount of time? Do participants have an incentive to truthfully reveal

their information? As we will see, it is not usually possible to simultaneously resolve all these questions positively.

In this thesis, we consider algorithm design for two classes of problem. First, we consider the design of algorithms for allocating shared resources. As a canonical example of shared resource allocation, consider the allocation of computing resources. Limited resources, such as processor time and memory space, are continually allocated among competing applications in datacenters, shared clusters, and at the level of individual machines. The second class of problems concerns using individual information to probabilistically forecast future events. For instance, a company may want to estimate the probability that its new product will ship on time; one way to do this is to obtain estimates from the employees involved in the development of the product. Media outlets may wish to forecast the result of an election, or a government agency may wish to estimate the effect of a new policy on some social indicator.

We now describe each of these areas, and our contribution, in more detail.

Shared Resource Allocation. The sharing economy has already had significant impact on the way people access a myriad of services including transportation (Uber, Lyft) and accommodation (AirBnB, Couchsurfing). But existing models mostly rely on shared *usage*, where a single owner shares their item with many paying customers. However, at their heart, these models look a lot like our other economic systems, where individuals pay for a service. A single owner has full control over their item, and seeks to maximize profit with no regard to fairness. This is a familiar objective.

When not only the usage, but also the *ownership* of resources is shared, the objective changes from an economic design standpoint. Under this model, we need protocols that achieve efficient allocation while also being fair, by guaranteeing that benefits from sharing are distributed evenly among stakeholders. More pragmatically,

shared ownership requires clear rules and an unambiguous democratic process put in place, rather than a single owner having the final say. Thus, shared ownership models are naturally suited to algorithmic solutions, where the rules can be checked and agreed upon in advance, and consistently executed as required.

In modern computing, a shared ownership model of resource allocation is common. Resource requirements can fluctuate heavily over time, but adverse effects can be mitigated by pooling resources with other users. Examples include maintaining a shared cluster that individual researchers submit jobs to, allocating CPU time among competing applications, and storing data in a shared cache. However, even beyond computing, there are many areas where allocating shared resources is a key concern. Cars, for instance, are often jointly owned between members of one or more families (and, with the advent of self-driving cars and increased urbanization, this trend appears to be accelerating). As another example, public funds and property are resources jointly owned by taxpayers; deciding what to build on a piece of public land is a (shared) resource allocation problem in this sense.

The work in this part of the thesis bears similarity to the literature on fair division. Informally, fair division is the study of allocating scarce resources among competing agents. However, work in fair division generally[1] makes (at least) two assumptions that make it unsuitable for the resource sharing applications we have discussed.

1. **Private Goods.** Most work in fair division assumes that goods are private, meaning that at most one agent can derive postive utility per good. The classic cake-cutting problem lies in this setting, as does the problem of allocating indivisible goods among agents. However, in many resource sharing settings, we are required to make decisions where more than one agent can benefit simultaneously. Examples include storing data in a shared cache (many different

---

[1] There are, of course, exceptions that will be discussed as relevant.

	Private Goods	Public Goods
Offline	Standard Fair Division	Chapter 2
Online	Chapter 4	Chapter 3

FIGURE 1.1: Summary of problem settings for the first part of this thesis.

applications can access the cached data), building public infrastructure (many people use public roads, hospitals, parks, etc), and deciding who gets to use a jointly-owned vehicle (which can carry more than one person).

2. **Offline decisions.** Often, we have the power to re-allocate resources over time (e.g., you can take the car today, but I get it tomorrow). In these settings, in makes sense to consider the allocation problem as a repeated, or online, process. Doing so allows us to exploit heterogeneity in demand not only across goods, but also across time.

The key contribution of this part of the thesis is to propose solutions to resource allocation problems where these assumptions do not apply. Figure 1.1 provides an overview of the problem settings that we consider. Note that, throughout these chapters, we will not consider the use of money. Although money can make information elicitation easier, there are many settings where it is either disallowed for legal or moral reasons, or simply unsuitable. For example, when allocating use of a resource among friends over time, it is uncommon for people to exchange money regularly (of course, all the friends may have contributed money to acquire the resource originally, but this is a one-off payment). When voting on the provision of public funds or to elect a representative, we generally abide by the "one person, one vote" principle, rather than allowing vote-buying.[2]

---

[2] A line of work on *quadratic voting* is a notable exception to this rule [134, 161, 135], though even there it's not quite as simple as votes being sold for a single anonymous price.

Forecasting. Reasoning about future events is a central problem for making good decisions. One way to do this is to aggregate the beliefs of many other agents, harnessing the "wisdom of crowds" effect. The wisdom of crowds effect relies on two factors. First, that many agents provide an estimate, and second, that agents truthfully report their beliefs.[3] The work presented in this part of the thesis therefore aims to achieve these twin objectives.

The most common way to incentivize agents to put effort into gathering and (truthfully) reporting their information is to reward them with payments that depend on the quality of their forecasts. Several approaches have been designed for this purpose – in this thesis, we focus on *wagering mechanisms* and *prediction markets*, but other mechanisms include scoring rules and forecasting competitions.

In a wagering mechanism, forecasters submit both a probabilistic prediction regarding the outcome of some event, and an amount of money that they are prepared to wager. Once the outcome is realized, the mechanism redistributes the wagers among the agents depending on their wager and the quality of their prediction. A prediction market works similarly, but is more dynamic. Agents trade securities that pay off depending on the outcome of the event. At any point in time, securities corresponding to different outcomes can be bought or sold at a given price, which is interpreted as the market's perceived probability of that outcome being realized.

We address two key problems pertaining to the design of wagering mechanisms and prediction markets.

1. Existing incentive-compatible wagering mechanisms suffer from a *low stakes problem*: In most cases, participants lose far less money than their full wager,

---

[3] It can be sufficient for agents to misreport their beliefs in a predictable way, so that their true beliefs can be inferred from their reports. Among other problems, this is unreliable in practice as it may not be the case that all agents construct their reports in the predicted manner. Further, by the *revelation principle* [126], the existence of a non-truthful mechanism implies the existence of a truthful mechanism with the same properties.

even when they make a bad prediction. This introduces several problems. It disincentivizes participation, since agents stand to make only small profits and losses, and it can lead to distorted incentives, even with respect to the probabilistic predictions, if agents choose to artificially inflate their wager. We formalize the tradeoff between incentive compatibility and a Pareto optimality condition that captures the low-stakes problem, and propose mechanisms to address the problem.

2. Recently, there has been a wave of *decentralized prediction markets* [132, 2, 4, 3], that operate without a single trusted governing entity. There are many advantages to such a model, but one disadvantage is that markets must be closed according to consensus among a group of arbiters, rather than by some authority. Since these platforms are, by design, anonymous, it is impossible to prevent the arbiters from also participating in the market. Therefore, in these markets we not only have to consider the incentives of a trader, but also of an arbiter (who may also be a trader).

While the two parts of this thesis are presented as distinct components, forging a connection between them remains a compelling direction for future work. A particularly promising direction is to apply forecasting mechanisms to online resource allocation problems. The challenge in designing online algorithms for these problems, and online algorithms in general, is the need to commit to a decision without any knowledge of the future. However, if we can run a prediction mechanism, either built in to an online algorithm or running parallel to it, then we may be able to incentivize agents to predict their future reports, and adapt our online algorithm accordingly. In this way, we may be able to direct additional utility to agents who increase social welfare by accurately predicting their future reports.

1.1 Structure of this Thesis and Overview of Results

This thesis describes a selection of my work on shared resource allocation and fore-casting, with three chapters devoted to each topic.

1.1.1 Shared Resource Allocation

Each chapter in this part of the thesis corresponds to a different relaxation of the standard fair division setting, where private goods are allocated among agents in a one-shot (offline) manner.

Chapter 2: Public Resources, Offline Decisions. In this chapter, we consider a multiple issue, *public decision* model that generalizes the problem of allocating indivisible private goods. The primary question that we consider is the following: How can we best define notions of fairness in this public decision setting?

In the private goods setting, the dominant notion of fairness is that of *envy-freeness* [76], which says that no agent should prefer any other agent's allocation to their own, and its relaxations such as envy-freeness *up to one good (EF1)* [117]. Unfortunately, when resources are public, agents no longer receive a well-defined allocation, and the notion of envy is therefore not applicable.

In this chapter, we take a different route to defining fair allocations. Our starting point is the *proportionality* guarantee [153], which says that each agent should receive at least a $1/n$ fraction of the utility she would receive if she were a dictator capable of unilaterally making all decisions. Proportionality is weaker than envy-freeness, but for deterministic outcomes it is still impossible to guarantee for all instances.

We therefore define fairness guarantees that relax proportionality. We introduce the *Round Robin Share (RRS)* and the *Pessimistic Proportional Share (PPS)*, which provide utility guarantees to the agents based on the utilities that they receive according to simple protocols. In another direction, and inspired by the definition of

7

EF1, we define the *Proportionality up to One Issue (Prop1)* guarantee, which says that each agent should be able to guarantee themselves their proportional share if we give them the power to change the outcome on a single issue.

After defining these three fairness properties, we consider two mechanisms that are known to satisfy desirable properties in the private goods case. The first is the *Maximum Nash Welfare (MNW)* mechanism [124], which maximizes the product of the utilities of all the agents, and the second is the *leximin* mechanism, which maximizes the utility of the agent with the minimum utility. We show that MNW satisfies Prop1 but not RRS or PPS, while leximin satisfies RRS and PPS but not Prop1. An exciting question that our work leaves open is whether there exists a mechanism that satisfies all three properties, as well as the efficiency property of Pareto Optimality (PO).

Finally, we consider the complexity of computing outcomes that satisfy our fairness properties in conjunction with PO. We show that satisfying RRS or PPS and PO simultaneously is $\mathcal{NP}$-hard. However, for the special case of private goods division, we provide a polynomial-time algorithm to compute an allocation satisfying PPS and PO.

Chapter 3: Public Resources, Online Decisions. In this chapter, we consider the same public decision framework as in Chapter 2, but we allow issues to arrive online. Motivated by the performance of the MNW algorithm for the offline version of the problem, we examine algorithms for online Nash Welfare maximization.

We focus on two greedy algorithms. The first, GREEDY, makes decisions that maximize the cumulative Nash Welfare. The second, PF, chooses the alternative that maximizes the sum of percentage increases in cumulative utility over agents. Due to fundamental difficulties that arise from multiplying and dividing by zero, it is unclear how these algorithms should behave, or even be defined, when some agents

have not accumulated any utility at the present time.

We provide a framework that both unifies the GREEDY and PF algorithms and defines their output when some agents have zero accumulated utility. The main idea is to endow each agent with a random infinitesimal hallucinated utility that allows us to sidestep problems arising from zero-utility agents. Not only does this provide a way to choose a single alternative (by randomly sampling infinitesimal utilities), but we provide efficient algorithms for computing the full set of possibly chosen alternatives.

Our second contribution is to provide an axiomatic characterization of the PF algorithm in terms of four axioms: scale-freeness, plurality, separability into single-minded agents, and no zero-dominated alternatives. Finally, we compare the performance of both algorithms on a computer systems example and compare them to a state-of-the-art algorithm in terms of theoretical guarantees. Both algorithms outperform the state-of-the-art on our dataset, with GREEDY slightly outperforming PF.

Chapter 4: Private Resources, Online Decisions. In the previous two chapters, we have not considered the agents' (lack of) incentives to provide their true utilities to the mechanism. Indeed, when resources are public, any consideration of incentives is difficult because of the famous *free rider problem.* Informally, if it is impossible to prevent an agent from using a resource, then an agent that gets high utility from some resource can simply pretend not to like it, and be 'compensated' by the algorithm in some other way.

In this chapter, we consider the problem of designing incentive-compatible mechanisms for the repeated allocation of private goods, thus avoiding the free rider problem. In our setting, each agent contributes some number of identical units of a private good to a shared pool. At each of a finite number of rounds, this pool

of resources must be reallocated among the agents that contribute to it. At each round, each agent has a demand for resources that defines a piecewise linear utility function. For every resource received up to their demand, agents receive high (H) utility per resource. For every resource over and above their demand, agents receive low (L) utility per resource.

Our goal is to design online allocation mechanisms that satisfy strategy-proofness (no agent can improve their utility by lying about their demands) and sharing incentives (agents do better by contributing their resources to the common pool than by withholding them). There are two key ways in which our work diverges from previous work [96, 97, 99, 20, 14, 17, 16, 154, 146]. The first is that we require our economic properties to hold ex-post, meaning that participants can not regret their truth-telling or participation decisions even in hindsight. Other works only require that participants make their decision with distributional knowledge of other participants' utilities, and that truthful reporting is optimal in expectation with respect to these distributions. The second is that we allow $L > 0$, meaning that agents can still derive positive utility from extra resources by, say, running background tasks or over-consuming now to reduce demand later.[4]

We design two mechanisms. The *Flexible Lending* mechanism satisfies strategy-proofness and approximate sharing incentives, while providing close to optimal performance in practice, in the sense that resources are almost always allocated to agents that have high utility for them. For settings where exact sharing incentives is desirable, we also design the T-Period mechanism, which satisfies both strategy-proofness and sharing incentives. In doing so, however, it restricts the space of possible allocations, resulting in significantly lower performance than the flexible lending mechanism in practice.

---

[4] Some previous work allows for more general utility functions, but does not obtain ex-post axiomatic guarantees.

1.1.2 Forecasting

Chapters 5 and 6 investigate a tradeoff between economic efficiency and incentive compatibility in wagering mechanisms. Chapter 7 concerns the design of prediction markets that must be closed by popular vote, rather than by a single trusted entity.

Chapter 5: Incentive Compatible and Efficient Wagering: The Double Clinching Auction
This chapter is inspired by the observation that existing wagering mechanisms [112, 113, 56] ask agents to specify a maximum acceptable loss (their wager), but in most instances agents incur a loss that is only a small fraction of the specified amount. Some of the wager is therefore left unused, in the sense that the agent would prefer to use it to bet against another agent with a different probability estimate.

Our first contribution in this chapter is to formalize this intuition via the notion of Pareto optimality. We say that a wagering mechanism is Pareto optimal if, after all agents have made their reports to the mechanism, no group of agents would like to make a side bet amongst themselves in addition to the bets facilitated by the mechanism. Informally, this says that the mechanism is extracting all possible trade. Unfortunately, we show that Pareto optimality is incompatible with three other fundamental properties in wagering mechanism design: individual rationality (all agents should participate willingly), weak budget balance (the mechanism should not lose money), and weak incentive compatibility (agents should not profit from lying about their probability estimate).

Therefore, we seek to design a wagering mechanism that retains the three core properties, while coming 'close' to Pareto optimality in some sense. To that end, we design the *Double Clinching Auction (DCA)*. The DCA uses the observation that any wagering mechanism can be expressed in terms of the allocation of Arrow-Debreu securities. It then utilizes an auction format known as the *adaptive clinching auction*

to allocate securities to agents – conveniently, the adaptive clinching auction requires each agent to report a fixed valuation per item and a budget. The former can be inferred from an agent's probability estimate, and the latter corresponds exactly to their wager. The final innovation in the design of the DCA is to have the mechanism sell the 'right' number of securities to guarantee incentive compatibility. This turns out to be the point at which the mechanism exactly breaks even on the lowest priced security.

In experiments based on real contest data, we show that the DCA comes a lot closer to achieving Pareto optimality than existing incentive-compatible wagering mechanisms, in at least two senses. First, the total risk incurred by the agents is higher than for other mechanisms, and second, a large proportion of the agents have a worst case loss that is exactly equal to their wager.

Chapter 6: Efficient Wagering by Relaxing Incentive Compatibility In this chapter we consider the tradeoff between incentive compatibility and Pareto optimality from a different angle. Rather than insisting on incentive compatibility and sacrificing Pareto optimality, we relax incentive compatibility and retain Pareto optimality (as well as budget balance and individual rationality). To achieve this, we consider the *Parimutuel Consensus Mechanism (PCM)*, defined by Eisenberg and Gale [72].

We show that the PCM not only satisfies individual rationality, strict budget balance, and Pareto optimality, but also anonymity (the outcome does not depend on the agents' identities), sybilproofness (agents splitting or merging identities does not affect the outcome), and envy-freeness (no agent envies the bet defined by the mechanism for any other agent). Further, subject to a mild condition on the reports of the agents, the PCM is the only wagering mechanism that satisfies these six properties.

The rest of this chapter argues that, despite not satisfying incentive compatibility,

the PCM does retain some desirable incentive properties. First, it satisfies *incentive-compatibility in the large* [22], which says that incentives to misreport vanish as the number of agents grows large. Therefore, we would expect few opportunities to misreport in data with large numbers of agents, which is exactly what we observe in our real contest data. Not only that, but profitable misreports, when they exist, are not very profitable and arguably not worth the inherent risk that a misreporting agent faces due to their uncertainty regarding the reports of other agents. We also show that even on small instances, the PCM remains fairly robust to misreporting.

Chapter 7: Crowdsourced Outcome Determination in Decentralized Prediction Markets In this chapter we turn to prediction market design. Inspired by the rise of decentralized prediction market platforms, we design a prediction market mechanism that does not rely on the existence of a trusted center to close the market. Instead, the market must be closed by popular vote amongst a group of arbiters. The fundamental difficulty is that the true outcome cannot be verified, and there is no way to prevent some arbiters from also having a position in the market. This creates a strong incentive problem, where arbiters may want to close the market in the direction of the outcome that achieves the highest profit for them, rather than the true outcome.

To escape this intractability, we make a key assumption: that the total budget of any given trader is upper bounded by some constant B. This allows us to bound the total number of securities that a trader is able to buy, thus limiting the profit that they can achieve by changing the market outcome. However, arbiters still have non-zero incentive to manipulate the market outcome to match their position. To counter this, we incorporate a peer prediction mechanism to incentivize arbiters to truthfully report the outcome that they observe.

Our mechanism incorporates three key innovations. First, in the market trading phase, we incorporate a trading fee. Fees are common in real-world prediction mar-

kets but are usually a necessary inconvenience rather than a feature. Our trading fee ensures that the price of any security does not become too low, thus guaranteeing that a trader with a bounded budget can only buy a bounded number of securities (without this, the budget bound B would not be useful). Further, we use the funds raised from the trading fees to pay the arbiters (via the peer prediction mechanism) for closing the market.

Second, when eliciting votes from the arbiters, we pay them according to a modified version of the 1/prior peer prediction mechanism [162, 106, 107], that we term the *1/prior with midpoint* mechanism. This is a technical adaptation that reduces the effect of any asymmetry in the belief update model held by the arbiters.

Third, to determine the outcome of the market (and therefore the payoff for each security), we allow continuous outcomes rather than a binary yes/no. In particular, we declare the market outcome to be the *fraction* of arbiters that report the event to have occurred. One advantage of this is that each market has a well-defined outcome. Even if the event to be predicted is ambiguous or unclear, traders are explicitly predicting the behavior of the arbiters relative to the question. But the key advantage is that the continuous payoff structure prevents any single arbiter from having too much effect on the outcome. Each arbiter can modify the payoff of a security by only $1/m$, where m is the number of arbiters, rather than (potentially) being able to swing the outcome from 0 to 1, or vice versa.

1.2 Bibliographic Notes

All work presented in this thesis is based on work published with co-authors. I am the primary author, or one of two joint primary authors, on each paper.

Chapter 2 is based on joint work with Vincent Conitzer and Nisarg Shah [61]. Chapter 3 is based on joint work with Vincent Conitzer and Seyed Majid Zahedi [83]. Chapter 4 is based on joint work with Vincent Conitzer, Benjamin Lee, and Seyed

Majid Zahedi [84].[5]   Chapter 5 is based on joint work with David Pennock and Jennifer Wortman Vaughan [82].   Chapter 6 is based on joint work with David Pennock [77]. Chapter 7 is based on joint work with Sébastien Lahaie and David Pennock [81].

1.2.1 Omitted Work

Work completed during my Ph.D. studies that has been excluded from this thesis includes:

- Work on voting: runoff scoring rules [78], tiebreaking [79], multi-winner approval voting [25, 26, 43] and societal tradeoffs [59, 60].

- Work on game theory: possible/necessary equilibrium actions [41] and signaling in Bayesian Stackelberg games [167].

- Design of forecasting competitions [165].

- Design of false-name-proof recommendation systems [42].

- Price of stability in network design games [80].

---

[5] This work also appears in Zahedi's Ph.D. dissertation.

2

Public Goods, Offline Allocation

2.1 Introduction

In this chapter, we study a model of decision making that generalizes classic fair division. In the long history of fair division - dating back to at least the work of Steinhaus [153], most work focuses on the fair division of *private goods*, in which a set of m items must be divided among a set of n agents. Agents express their preferences by specifying their value for each good, and our goal is to find a division of the goods that is fair to all agents.

One particularly appealing notion of fairness is *envy-freeness* [76], which says that no agent should want to switch her set of items with that of another agent. This is a natural and strong notion of fairness that has long been the subject of fair division research [156, 141, 110, 100, 37, 38, 46]. It actually implies many other fairness notions such as *proportionality* [153] — each agent should get at least a $1/n$ fraction of her value for the entire set of goods — and *envy-freeness up to one good (EF1)* [117] — no agent should envy another agent after removing at most one good from the latter agent's bundle. Unfortunately, envy-freeness cannot always be

guaranteed, and therefore its relaxations have been focused on [117, 44, 138, 49].

Division of private goods, however, is not the only application in which we may desire a fair outcome. Often, we may need to make decisions where every alternative gives positive utility to many agents, rather than to just one agent as in the case of private goods. For instance, consider a couple, Alice and Bob, deciding where to go to dinner. Alice likes Italian food the most, but does not like Indian, whereas Bob prefers Indian food but does not like Italian. When there is only a single decision to make, we are simply in a classic bargaining game where agents must attempt to arrive at a mutually agreeable solution. Nash [124] proposed maximizing the product of agents' utilities (the *Nash welfare*) as an elegant solution that uniquely satisfies several appealing properties. But no matter how we arrive at a decision – and there is a myriad of work in computational social choice [39] discussing how exactly we should do so – some tradeoff must necessarily be made, and we may not be able to make everyone happy.

However, if we have several public decisions to make, maybe we can reach a compromise by making sure that all agents are happy with at least some of the decisions. For example, if Alice and Bob are to follow their dinner with a movie, then maybe Bob will be willing to eat Italian food for dinner if he gets to pick his favorite movie, and maybe Alice will agree to this compromise.

Note that this setting generalizes the classic private goods setting, because in this special case we can view each public decision as the allocation of a single good. While envy is a compelling notion in the private goods setting, it makes less sense for public decisions. In our example, irrespective of where Alice and Bob go for dinner, because they are eating the same food, it is not clear what it would mean for Alice to envy Bob. If she could somehow trade places with Bob, she would still be sitting at the other end of the dinner table, eating the same food, and not be any better off. Thankfully, proportionality still has a sensible interpretation: Each agent should

17

get at least a $1/n$ fraction of the utility she would get if her most desired alternative was chosen for each decision. Unfortunately, as with envy-freeness, proportionality cannot always be guaranteed. Therefore in this work we consider relaxations of proportionality in order to arrive at fairness notions that can be guaranteed.

2.1.1 Our Results

Formally, a public decision making problem consists of m issues, where each issue has several associated alternatives. Each of n agents has a utility for each alternative of each issue. Making a decision on an issue amounts to choosing one of the alternatives associated with the issue, and choosing an overall outcome requires making a decision on each issue simultaneously. The utility to a agent for an outcome is the sum of her utilities for the alternatives chosen for different issues. This is a very simple setting, but one in which the problem of fairness is already non-trivial.

We propose relaxations of proportionality in two directions. The first, *proportionality up to one issue (Prop1)*, is similar in spirit to EF1, stating that a agent should be able to get her proportional share if she gets to change the outcome of a single issue in her favor. The second direction is based on the guarantees provided by the round robin mechanism. This mechanism first orders the set of agents, and then repeatedly goes through the ordering, allowing each agent to make her favorite decision on any single issue, until decisions are made on all the issues. Our first relaxation in this direction, the *round robin share (RRS)*, guarantees each agent the utility that she would have received under the round robin mechanism if she were the last agent in the ordering. Note that the round robin mechanism lets each agent make decisions on roughly the same number of issues. A further relaxation in this direction, the *pessimistic proportional share (PPS)*, guarantees each agent the utility that she would get if her favorite alternatives were chosen for (approximately) a $1/n$ fraction of the issues, where these issues are chosen adversarially.

Table 2.1: Axioms satisfied or approximated by the mechanisms we consider. The MNW solution is split into private goods and general decisions because we obtain significantly stronger results for private goods. Results for the leximin mechanism and the round robin method apply equally to private goods and public decisions. The approximation results are lower bounds; we omit the upper bounds from the table for simplicity.

	PO	PPS	RRS	Prop1
MNW, Private goods	✓	✓ (Th. 10)	$\frac{1}{2}$ (Th. 10)	✓ (Th. 4)
MNW, Public decisions	✓ (Th. 4)	$\frac{1}{n}$ (Th. 8)	$\frac{1}{n}$ (Th. 8)	✓ (Th. 4)
Leximin Mechanism	✓	✓ (Th. 2)	✓ (Th. 2)	$\frac{1}{2}$ (Th. 2)
Round Robin Method	×	✓ (Th. 1)	✓ (Th. 1)	✓ (Th. 1)

We examine the possibility and computational complexity of satisfying combinations of these fairness desiderata. We first observe that the round robin mechanism satisfies both Prop1 and RRS (and thus PPS). However, it fails to satisfy even the most basic efficiency property, Pareto optimality (PO), which requires that no other outcome should be able to make a agent strictly better off without making at least one agent strictly worse off.

When insisting on Pareto optimality, we observe that the leximin mechanism — informally, it chooses the outcome that maximizes the minimum utility to any agent — satisfies RRS (therefore PPS) and PO via a simple argument. However, this argument does not extend to establishing Prop1, although we show that RRS implies a 1/2 approximation to Prop1. To that end, we prove that the maximum Nash welfare (MNW) solution — informally, it chooses the outcome that maximizes the *product* of utilities to agents — that is known for its many desirable fairness properties in dividing private goods [49] satisfies Prop1 and PO in our public decision making framework, and simultaneously provides a $1/n$ approximation to both RRS and PPS. We also show that this approximation is tight up to a factor of $O(\log n)$. For division of private goods, these approximations improve significantly: the MNW solution completely satisfies PPS, and provides an $n/(2n-1) > 1/2$ approximation

(but not better than 2/3 approximation in the worst case) to RRS. Table 2.1 provides a summary of these results.

However, both the MNW outcome and the leximin outcome are $\mathcal{NP}$-hard to compute. It is therefore natural to consider whether our fairness properties can be achieved in conjunction with PO in polynomial time. For public decision making, the answer turns out to be negative for PPS and RRS, assuming $\mathcal{P} \neq \mathcal{NP}$. For division of private goods, however, we show that there exists a polynomial time algorithm that satisfies PPS and PO.

2.1.2   Related Work

Two classic fair division mechanisms — the leximin mechanism and the maximum Nash welfare (MNW) solution — play an important role in this paper. Both mechanisms have been extensively studied in the literature on private goods division. In particular, Kurokawa et al. [111] (Section 3.2) show that the leximin mechanism satisfies envy-freeness, proportionality, Pareto optimality, and a strong game-theoretic notion called group strategyproofness, which prevents even groups of agents from manipulating the outcome to their benefit by misrepresenting their preferences, in a broad fair division domain with private goods and a specific form of non-additive utilities. On the other hand, the MNW solution has been well studied in the realm of additive utilities [140, 49]. For divisible goods, the MNW solution coincides with another well-known solution concept called *competitive equilibrium from equal incomes* (CEEI) [156], which also admits an approximate version for indivisible goods [44]. For indivisible goods, the MNW solution satisfies envy-freeness up to one good, Pareto optimality, and approximations to other fairness guarantees. One line of research aims to approximate the optimum Nash welfare [58, 116], although it is unclear if this achieves any of the appealing fairness guarantees of the MNW solution.

Our model is closely related to that of voting in combinatorial domains (see Lang

and Xia [114] for an overview). However, this literature focuses on the case where there is dependency between decisions on different issues. In contrast, our model remains interesting even though the issues are independent, and incorporating dependency is an interesting future direction. Although there is a range of work in the voting literature that focuses on fairness [51, 120, 34, 26], especially in the context of *representation* in multi-winner elections, it focuses on ordinal, rather than cardinal, preferences.[1] Another difference is that fairness concepts in voting apply most naturally when $n >> m$, whereas our notions of fairness are most interesting when $m \geqslant n$.

Our work is also reminiscent of the *participatory budgeting* problem [45], in which there are multiple public projects with different costs, and a set of projects need to be chosen based on preferences of the participants over the projects, subject to a budget constraint. Recently, researchers in computational social choice have addressed this problem from an axiomatic viewpoint [93], including fairness considerations [74], and from the viewpoint of *implicit utilitarian voting* [29]. However, they assume access only to ordinal preferences (that may stem from underlying cardinal utilities), while we assume a direct access to cardinal utilities, as is common in the fair division literature. Also, we do not have a budget constraint that binds the outcomes on different issues.

2.2   Model and Preliminaries

For $k \in \mathbb{N}$, define $[k] \triangleq \{1, \ldots, k\}$. Before we introduce the problem we study in this paper, let us review the standard fair division setting with private goods.

Private goods division.   A private goods division problem consists of a set of

---

[1] That said, there is a recent line of work on *implicit utilitarian voting* that attempts to maximize an objective with respect to the cardinal utilities underlying the ordinal preferences [137, 35], and is therefore closer to our work.

agents $N = [n]$ and a set of m *goods* M. Each agent $i \in N$ is endowed with a utility function $u_i : M \to \mathbb{R}_+$ such that $u_i(g)$ denotes the value agent i derives from good $g \in M$. A standard assumption in the literature is that of additive valuations, i.e., (slightly abusing the notation) $u_i(S) = \sum_{g \in S} u_i(g)$ for $S \subseteq M$. An *allocation* $\mathbf{A}$ is a partition of the set of goods among the set of agents, where A_i denotes the bundle of goods received by agent i. Importantly, agents only derive utility from the goods they receive, i.e., the goods private to them. The utility of agent i under allocation $\mathbf{A}$ is $u_i(\mathbf{A}) = u_i(A_i)$.

Public decision making. A *public decision making* problem also has a set of agents $N = [n]$, but instead of private goods, it has a set of *issues* $T = [m]$. Each issue $t \in T$ has an associated set of alternatives $A^t = \{a_1^t, \ldots, a_{k_t}^t\}$, exactly one of which must be chosen. Each agent i is endowed with a utility function $u_i^t : A^t \to \mathbb{R}_+$ for each issue t, and derives utility $u_i^t(a_j^t)$ if alternative a_j^t is chosen for issue t. In contrast to private goods division, a single alternative can provide positive utility to multiple agents.

An outcome $\mathbf{c} = (c_1, \ldots, c_m)$ of a public decision making problem is a choice of an alternative for every issue, i.e., it consists of an outcome $c_t \in A^t$ for each issue $t \in T$. Let $\mathcal{C}$ denote the space of possible outcomes. Slightly abusing the notation, let $u_i^t(\mathbf{c}) = u_i^t(c_t)$ be the utility agent i derives from the outcome of issue t. We also assume additive valuations: let $u_i(\mathbf{c}) = \sum_{t \in T} u_i^t(\mathbf{c})$ be the utility agent i derives from outcome $\mathbf{c}$.

In this work, we study deterministic outcomes, and in Section 2.5, discuss the implications when randomized outcomes are allowed. Further, we study the *offline* problem in which we are presented with the entire problem up front, and need to choose the outcomes on all issues simultaneously.

Private goods versus public decisions. To see why public decision making gen-

eralizes private goods division, take an instance of private goods division, and create an instance of public decision making as follows. Create an issue t_g for each good g. Let there be n alternatives in A^{t_g}, where alternative $a_i^{t_g}$ gives agent i utility $u_i(g)$ while giving zero utility to all other agents. It is easy to see that choosing alternative $a_i^{t_g}$ is equivalent to allocating good g to agent i. Hence, the constructed public decision making problem effectively mimics the underlying private goods division problem.

2.2.1 Efficiency and Fairness

In this paper, we not only adapt classical notions of efficiency and fairness defined for private goods division to our public decision making problem, but also introduce three fairness axioms that are novel for both public decision making and private goods division. First, we need additional notation that we will use throughout the paper.

Let $p \triangleq \lfloor m/n \rfloor$. For issue $t \in T$ and agent $i \in N$, let $a_{max}^t(i) \in \arg\max_{a \in A^t}\{u_i^t(a)\}$ and $u_{max}^t(i) = u_i^t(a_{max}^t(i))$. That is, $a_{max}^t(i)$ is an alternative that gives agent i the most utility for issue t, and $u_{max}^t(i)$ is the utility agent i derives from $a_{max}^t(i)$. Let the sequence $\langle u_{max}^{(1)}(i), \ldots, u_{max}^{(m)}(i) \rangle$ represent the maximum utilities agent i can derive from different issues, sorted in a non-ascending order. Hence, $\{u_{max}^{(k)}(i)\}_{k \in [m]} = \{u_{max}^t(i)\}_{t \in T}$ and $u_{max}^{(k)}(i) \geqslant u_{max}^{(k+1)}(i)$ for $k \in [m-1]$.

Efficiency. In this paper, we focus on a popular notion of economic efficiency. We say that an outcome $\mathbf{c}$ is *Pareto optimal* (PO) if there does not exist another outcome $\mathbf{c}'$ that can provide at least as much utility as $\mathbf{c}$ to every agent, i.e., $u_i(\mathbf{c}') \geqslant u_i(\mathbf{c})$ for all $i \in N$, and strictly more utility than $\mathbf{c}$ to some agent, i.e., $u_{i*}(\mathbf{c}') > u_{i*}(\mathbf{c})$ for some $i^* \in N$.

Fairness. For private goods division, perhaps the most prominent notion of fairness

is envy-freeness [76]. An allocation $\mathbf{A}$ is called *envy-free* (EF) if every agent values her bundle at least as much as she values any other agent's bundle, i.e., $u_i(A_i) \geqslant u_i(A_j)$ for all $i, j \in N$. Because envy-freeness cannot in general be guaranteed, prior work also focuses on its relaxations. For instance, an allocation $\mathbf{A}$ is called *envy-free up to one good* (EF1) if no agent envies another agent after removing at most one good from the latter agent's bundle, i.e., for all $i, j \in N$, either $u_i(A_i) \geqslant u_i(A_j)$ or $\exists g_j \in A_j$ such that $u_i(A_i) \geqslant u_i(A_j \backslash \{g_j\})$.

Unfortunately, as argued in Section 2.1, the notion of envy is not well defined for public decisions. Hence, for public decision making, we focus on another fairness axiom, Proportionality, and its relaxations. For private goods division, proportionality is implied by envy-freeness.[2]

Proportionality (Prop). At a high level, proportionality requires that each agent must receive at least her "proportional share", which is a $1/n$ fraction of the utility she would derive if she could act as the dictator. For a public decision making problem, the *proportional share* of agent i (Prop$_i$) is $1/n$ times the sum of the maximum utilities the agent can derive across all issues, i.e.,

$$\text{Prop}_i = \frac{1}{n} \sum_{t \in T} u_{max}^t(i).$$

For $\alpha \in (0, 1]$, we say that an outcome $\mathbf{c}$ satisfies α-*proportionality* (α-Prop) if $u_i(\mathbf{c}) \geqslant \alpha \cdot \text{Prop}_i$ for all agents $i \in N$. We refer to 1-Prop simply as Prop.

Proportionality up to one issue (Prop1). We introduce a novel relaxation of proportionality (more generally, of α-proportionality) in the same spirit as envy-freeness up to one good, which is a relaxation of envy-freeness. For $\alpha \in (0, 1]$, we say that an outcome $\mathbf{c}$ satisfies α-*proportionality up to one issue* (α-Prop1) if for every agent $i \in N$, there exists an issue $t \in T$ such that, ceteris paribus, changing

---

[2] This assumes non-wastefulness, i.e., that all goods are allocated. We make this assumption throughout the paper.

the outcome of t from c_t to $a^t_{max}(i)$ ensures that agent i achieves an α fraction of her proportional share, i.e., if

$$\forall i \in N \;\; \exists t \in T \;\; \text{s.t.} \;\; u_i(\mathbf{c}) - u_i^t(\mathbf{c}) + u_{max}^t(i) \geqslant \alpha \cdot \text{Prop}_i.$$

We refer to 1-Prop1 simply as Prop1.

Round robin share (RRS). Next, we introduce another novel fairness axiom that is motivated from the classic round robin method that, for private goods, lets agents take turns and in each turn, pick a single most favorite item left unclaimed. For public decision making, we instead let agents make a decision on a single issue in each turn. The utility guaranteed to the agents by this approach is captured by the following fairness axiom.

Recall that the sequence $\langle u_{max}^{(1)}(i), \ldots, u_{max}^{(m)}(i) \rangle$ represents the maximum utility agent i can derive from different issues, sorted in a non-ascending order. Then, we define the *round robin share* of agent i (RRS_i) as

$$\text{RRS}_i = \sum_{k=1}^{p} u_{max}^{(k \cdot n)}(i).$$

This is agent i's utility from the round robin method, if she is last in the ordering and all issues she does not control give her utility 0. For $\alpha \in (0, 1]$, we say that an outcome $\mathbf{c}$ satisfies α-*round robin share* (α-RRS) if $u_i(\mathbf{c}) \geqslant \alpha \cdot \text{RRS}_i$ for all agents $i \in N$. Again, we refer to 1-RRS simply as RRS.

Pessimistic proportional share (PPS). We introduce another novel fairness axiom that is a further relaxation of round robin share. Note that the round robin method, by letting agents make a decision on a single issue per turn, allows each agent to make decisions on at least $p = \lfloor m/n \rfloor$ issues. The following axiom captures the utility agents would be guaranteed if each agent still made decisions on a "proportional share" of p issues, but if these issues were chosen pessimistically.

25

We define the *pessimistic proportional share* of agent i (PPS_i) to be the sum of the maximum utilities the agent can derive from a set of p issues, chosen adversarially to minimize this sum:

$$\text{PPS}_i = \sum_{k=m-p+1}^{m} u_{max}^{(k)}(i).$$

For $\alpha \in (0, 1]$, we say that an outcome $\mathbf{c}$ satisfies α-*pessimistic proportional share* (α-PPS) if $u_i(\mathbf{c}) \geqslant \alpha \cdot \text{PPS}_i$ for all agents $i \in N$. Again, we refer to 1-PPS simply as PPS.

Connections among fairness properties. Trivially, proportionality (Prop) implies proportionality up to one issue (Prop1). In addition, it can also be checked that the following sequence of logical implications holds: Prop $\implies$ MMS $\implies$ RRS $\implies$ PPS.

Here, MMS is the *maximin share guarantee* [44, 138]. Adapting the definition naturally from private goods division to public decision making, the maximin share of a agent is the utility the agent can guarantee herself by dividing the set of issues into n bundles, if she gets to make the decisions best for her on the issues in an adversarially chosen bundle. The maximin share (MMS) guarantee requires that each agent must receive utility that is at least her maximin share. We do not focus on the maximin share guarantee in this paper.

2.2.2 Mechanisms

A mechanism for a public decision making problem (resp. a private goods division problem) maps each input instance of the problem to an outcome (resp. an allocation). We say that a mechanism satisfies a fairness or efficiency property if it always returns an outcome satisfying the property. There are three prominent mechanisms that play a key role in this paper.

Round robin method. As mentioned earlier, the round robin method first fixes

an ordering of the agents. Then the agents take turns choosing their most pre-ferred alternative on a single issue of their choice whose outcome has not yet been determined.

The leximin mechanism. The leximin mechanism chooses an outcome which maximizes the utility of the worst off agent, i.e., $\min_{i \in N} u_i(\mathbf{c})$. Subject to this con-straint, it maximizes the utility of the second least well off agent, and so on. Note that the leximin mechanism is trivially Pareto optimal because if it were possible to improve some agent's utility without reducing that of any other, it would improve the objective that the leximin mechanism optimizes.

Maximum Nash welfare (MNW). The *Nash welfare* of an outcome $\mathbf{c}$ is the product of utilities to all agents under $\mathbf{c}$: $NW(\mathbf{c}) = \prod_{i \in N} u_i(\mathbf{c})$. When there exists an outcome $\mathbf{c}$ with $NW(\mathbf{c}) > 0$, then the MNW solution chooses an arbitrary outcome $\mathbf{c}$ that maximizes the Nash welfare. When all outcomes have zero Nash welfare, it finds a largest cardinality set S of agents that can be given non-zero utility, and selects an outcome maximizing the product of their utilities, i.e., $\prod_{i \in S} u_i(\mathbf{c})$.

2.2.3 Examples

We illustrate the fairness properties through two examples.

Example 1. *Consider a public decision making problem with two agents ($N = [2]$) and two issues ($T = [2]$). Each issue has two alternatives ($|A^1| = |A^2| = 2$). The utilities of the two agents for the two alternatives in both issues are as follows.*

$$
\begin{array}{c|cc}
 & a_1^t & a_2^t \\
\hline
u_1^t & 1 & 0 \\
u_2^t & 0 & 1
\end{array}
\quad \textit{for } t \in [2].
$$

The various fair shares of the two agents are $Prop_1 = RRS_1 = PPS_1 = Prop_2 = RRS_2 = PPS_2 = 1$. Now, outcome $\mathbf{c} = (a_1^1, a_1^2)$ gives utilities $u_1(\mathbf{c}) = 2$ and $u_2(\mathbf{c}) =$

0, *and therefore violates Prop, RRS, and PPS. It satisfies Prop1 because switching the decision on either issue in favor of agent 2 makes her achieve her proportional share. On the other hand, outcome $\mathbf{c} = (a_1^1, a_2^2)$ gives utility 1 to both agents, and thus satisfies Prop (as well as Prop1, RRS, and PPS, which are relaxations of Prop).*

Example 2. *Consider a public decision making problem with two agents ($N = [2]$) and eight issues ($T = [8]$). Once again, each issue has two alternatives, for which the utilities of the two agents are as follows.*

$$
\begin{array}{c|cc}
 & a_1^t & a_2^t \\
\hline
u_1^t & 1 & 0 \\
u_2^t & 0 & 1
\end{array}
\quad \text{for } t \in \{1,2,3,4\}, \quad \text{and} \quad
\begin{array}{c|cc}
 & a_1^t & a_2^t \\
\hline
u_1^t & 1 & 0 \\
u_2^t & 0 & 0
\end{array}
\quad \text{for } t \in \{5,6,7,8\}.
$$

In this case, we have $Prop_1 = RRS_1 = PPS_1 = 4$, whereas $Prop_2 = RRS_2 = 2$ and $PPS_2 = 0$. Consider outcome $\mathbf{c} = (a_1^1, a_1^2, a_1^3, a_1^4, a_1^5, a_1^6, a_1^7, a_1^8)$. Then, we have $u_1(\mathbf{c}) = 8$ while $u_2(\mathbf{c}) = 0$, which satisfies PPS but violates RRS. Further, $\mathbf{c}$ also violates Prop1 because switching the outcome of any single issue can only give agent 2 utility at most 1, which is less than $Prop_2 = 2$. On the other hand, outcome $\mathbf{c} = (a_2^1, a_2^2, a_2^3, a_2^4, a_1^5, a_1^6, a_1^7, a_1^8)$ achieves $u_1(\mathbf{c}) = u_2(\mathbf{c}) = 4$, and satisfies Prop (and thus its relaxations Prop1, RRS, and PPS).

2.3 (Approximate) Satisfiability of Axioms

If we are willing to sacrifice Pareto optimality, then we can easily achieve both RRS (and therefore PPS) and Prop1 simultaneously with the round robin mechanism. This is not a surprising result. RRS is defined based on the guarantee provided by the round robin mechanism, and PPS is a relaxation of RRS. The round robin mechanism is also known to satisfy EF1 for private goods division, which is similar in spirit to Prop1.

Theorem 1. *The round robin mechanism satisfies RRS (and therefore PPS) and Prop1, and runs in polynomial time.*

Proof. The round robin mechanism clearly runs in polynomial time (note that it is easy for a agent to choose the next issue on which to determine the outcome). To see why it satisfies RRS, note that the mechanism allows every agent to make a decision on one issue once every n turns. Thus, for each $k \in [p]$, every agent gets to make decisions on at least k of her "top" $k \cdot n$ issues, when issues are sorted in the descending order of the utility her favorite alternative in the issue gives her. It is easy to see that this implies every agent i gets utility at least RRS_i. Because RRS implies PPS, the mechanism also satisfies PPS. It remains to show that it satisfies Prop1 as well.

Fix a agent i and let $\mathbf{c}$ be the outcome produced by the round robin mechanism for some choosing order of the agents. Because the round robin mechanism satisfies RRS, agent i gets utility at least

$$u_i(\mathbf{c}) \geqslant \sum_{k=1}^{p} u_{max}^{(k \cdot n)}(i).$$

For $k \in [m]$, let the k^{th} favorite issue of agent i be the issue t for which $u_{max}^t(i)$ is the k^{th} highest. Let $\ell \in \mathbb{N} \cup \{0\}$ be the largest index such that for every $k \in [\ell]$, outcome $\mathbf{c}$ chooses agent i's most preferred alternative on her k^{th} favorite issue. Let t^* be her $(\ell + 1)^{\text{th}}$ favorite issue. To show that $\mathbf{c}$ satisfies Prop1, we construct outcome $\mathbf{c}'$ from $\mathbf{c}$ by only changing the outcome of issue t^* to $a_{max}^{t^*}(i)$, and show that $u_i(\mathbf{c}') \geqslant \text{Prop}_i$. Note that if $\ell \geqslant p$, then

$$u_i(\mathbf{c}') \geqslant \sum_{k=1}^{\ell+1} u_{max}^{(k)}(i) \geqslant \sum_{k=1}^{p+1} u_{max}^{(k)}(i) \geqslant \frac{1}{n} \sum_{k=1}^{m} u_{max}^{(k)}(i) = \text{Prop}_i.$$

Let $\ell < p$. Then, using the fact that the round robin mechanism lets agent i choose her most preferred alternative for at least k of her favorite $k \cdot n$ issues for every $k \leqslant p$

29

(and her $(\ell+1)^{\text{th}}$ favorite issue was not one of these), we have

$$u_i(\mathbf{c}') \geq \sum_{k=1}^{\ell+1} u_{max}^{(k)}(i) + \sum_{k=\ell+1}^{p} u_{max}^{(k \cdot n)}(i)$$

$$\geq \frac{1}{n} \sum_{k=1}^{(\ell+1) \cdot n} u_{max}^{(k)}(i) + \frac{1}{n} \sum_{k=(\ell+1)n+1}^{m} u_{max}^{(k)}(i) = \text{Prop}_i.$$

Therefore, the round robin mechanism satisfies Prop1. ∎

While this result seems to reflect favorably upon the round robin mechanism, recall that it violates Pareto optimality even for private goods division. For public decision making, a simple reason for this is that the round robin mechanism, for each issue, chooses an alternative that is some agent's favorite, while it could be unanimously better to choose compromise solutions that make many agents happy. Imagine there are two agents and two issues, each with two alternatives. The "extreme" alternative in each issue $i \in \{1, 2\}$ gives utility 1 to agent i but 0 to the other, while the "compromise" alternative in each issue $i \in \{1, 2\}$ gives utility $2/3$ to both agents. It is clear that both agents prefer choosing the compromise alternative in both issues to choosing the extreme alternative in both issues. Because such "Pareto improvements" which make some agents happier without making any agent worse off are unanimously preferred by the agents, the round robin outcome becomes highly undesirable. We therefore seek mechanisms that provide fairness guarantees while satisfying Pareto optimality.

A natural question is whether there exists a mechanism that satisfies RRS, Prop1, and PO. An obvious approach is to start from an outcome that already satisfies RRS and Prop1 (e.g., the round robin outcome), and make Pareto improvements until no such improvements are possible. While Pareto improvements preserve RRS as the utilities to the agents do not decrease, Prop1 can be lost as it depends on the exact alternatives chosen and not only on the utilities to the agents. We leave it as

an important open question to determine if RRS, Prop1, and PO can be satisfied simultaneously.

We therefore consider satisfying each fairness guarantee individually with PO. One can easily find an outcome satisfying RRS and PO by following the aforementioned approach of starting with an outcome satisfying RRS, and making Pareto improvements while possible. There is also a more direct approach to satisfying RRS and PO. Recall that the leximin mechanism chooses the outcome which maximizes the minimum utility to any agent, subject to that maximizes the second minimum utility, and so on. It is easy to see that this mechanism is always Pareto optimal. Now, let us normalize the utilities of all agents such that $\text{RRS}_i = 1$ for every agent $i \in N$.[3] Because the round robin mechanism gives every agent i utility at least $\text{RRS}_i = 1$, it must be the case that the leximin mechanism operating on these normalized utilities must also give every agent utility at least 1, and thus produce an outcome that is both RRS and PO.

Theorem 2. *The leximin mechanism satisfies RRS, PO, and (1/2)-Prop1.*

That leximin satisfies (1/2)-Prop1 follows directly from the following lemma, and noting that leximin satisfies RRS.

Lemma 3. *RRS implies (1/2)-Prop1.*

Proof. Note that

$$\text{RRS}_i = \sum_{k=1}^{p} u_{max}^{(k \cdot n)}(i) \geqslant \frac{1}{n} \sum_{t=n+1}^{m} u_{max}^{(t)}(i)$$

and

$$u_{max}^{(1)}(i) \geqslant \frac{1}{n} \sum_{t=1}^{n} u_{max}^{(t)}(i).$$

---

[3] Agents with zero round robin share can be incorporated via a simple extension to the argument.

Summing the two equations, we get

$$\mathrm{RRS}_i + u_{max}^{(1)}(i) \geqslant \frac{1}{n} \sum_{t=1}^{m} u_{max}^{(t)}(i) = \mathrm{Prop}_i.$$

Therefore, $\max\{\mathrm{RRS}_i, u_{max}^{(1)}(i)\} \geqslant \frac{1}{2}\mathrm{Prop}_i$.

Suppose that $u_i(\mathbf{c}) \geqslant \mathrm{RRS}_i$ for some outcome $\mathbf{c}$. Then either i already receives her most valued item, in which case she receives utility at least $\max\{\mathrm{RRS}_i, u_{max}^{(1)}(i)\} \geqslant \frac{1}{2}\mathrm{Prop}_i$, or she does not receive her most valued item. If she does not, then after giving it to her, she receives utility at least $\frac{1}{2}\mathrm{Prop}_i$. Therefore, $\mathbf{c}$ satisfies $(1/2)$-Prop1. ∎

Next, we study whether we can achieve Prop1 and PO simultaneously. Neither of the previous approaches seems to work: we already argued that following Pareto improvements could lose Prop1, and the normalization trick is difficult to apply because Prop1 is not defined in terms of any fixed share of utility.

One starting point to achieving Prop1 and PO is the maximum Nash welfare (MNW) solution, which, for private goods division, is known to satisfy the similar guarantee of EF1 and PO [49]. It turns out that the MNW solution is precisely what we need.

Theorem 4. *The MNW solution satisfies proportionality up to one issue (Prop1) and Pareto optimality (PO).*

Before we prove this, we need a folklore result, which essentially states that if the sum of n terms is to be reduced by a fixed quantity δ that is less than each term, then their product reduces the most when δ is taken out of the lowest term. The following lemma proves this result when all initial terms are 1, which is sufficient for our purpose. The proof of the lemma appears in the appendix.

32

Lemma 5. *Let $\{x_1, \ldots, x_n\}$ be a set of n non-negative real numbers such that $\sum_{i=1}^{n} \max\{0, 1 - x_i\} \leqslant \delta$, where $0 < \delta < 1$. Then, $\prod_{i=1}^{n} x_i \geqslant 1 - \delta$.*

Proof of Theorem 4. Fix an instance of the public decision making problem. Let $S \subseteq N$ be the set of agents that the MNW outcome $\mathbf{c}$ gives positive utility to. Then, by the definition of the MNW outcome, S must be a largest set of agents that can simultaneously be given positive utility, and $\mathbf{c}$ must maximize the product of utilities of agents in S.

First, we show that $\mathbf{c}$ is PO. Note that a Pareto improvement over $\mathbf{c}$ must either give a positive utility to a agent in $N \backslash S$ or give more utility to a agent in S, without reducing the utility to any agent in S. This is a contradiction because it violates either optimality of the size of S or optimality of the product of utilities of agents in S. Hence, MNW satisfies PO.

We now show that MNW also satisfies Prop1. Suppose for contradiction that Prop1 is violated for agent i under $\mathbf{c}$. First, note that we must have $\text{Prop}_i > 0$. Further, it must be the case that $u_{max}^t(i) > 0$ for at least $n + 1$ issues. Were this not the case, Prop1 would be trivially satisfied for agent i since we can give her utility

$$u_{max}^{(1)}(i) \geqslant \frac{1}{n} \sum_{t=1}^{n} u_{max}^{(t)}(i) = \frac{1}{n} \sum_{t=1}^{m} u_{max}^{(t)}(i) = \text{Prop}_i$$

by changing the outcome on a single issue.

We now show that $u_i(\mathbf{c}) > 0$ (i.e., $i \in S$). For contradiction, suppose otherwise. For each of the (at least) $n + 1$ issues with $u_{max}^t(i) > 0$, there must exist another agent $j \neq i$ that gets positive utility *only* from that issue under $\mathbf{c}$ (otherwise we could use that issue to give positive utility to i while not reducing any other agents' utility to zero, contradicting the maximality of S). But this is impossible, since there are at least $n + 1$ issues and only $n - 1$ agents (other than i).

Because MNW outcomes and the Prop1 property are invariant to individual scal-

ing of utilities, let us scale the utilities such that $\text{Prop}_i = 1$ and $u_j(\mathbf{c}) = 1$ for all $j \in S\backslash\{i\}$. Select issue $t^* \in T$ as

$$t^* \in \arg\min_{t \in T} \frac{\sum_{j \in N\backslash\{i\}} u_j^t(\mathbf{c})}{u_{max}^t(i) - u_i^t(\mathbf{c})}.$$

Note that t^* is well defined because $u_{max}^t(i) > u_i^t(\mathbf{c})$ for at least one $t \in T$, otherwise Prop1 would not be violated for agent i.

We now show that outcome $\mathbf{c}'$ such that $c'_{t*} = a_{max}^{t^*}(i)$ and $c'_t = c_t$ for all $t \in T\backslash\{t^*\}$ achieves strictly greater product of utilities of agents in S than outcome $\mathbf{c}$ does, which is a contradiction as $\mathbf{c}$ is an MNW outcome. First, note that

$$\frac{\sum_{j \in N\backslash\{i\}} u_j^{t^*}(\mathbf{c})}{u_{max}^{t^*}(i) - u_i^{t^*}(\mathbf{c})} \leqslant \frac{\sum_{t \in T} \sum_{j \in N\backslash\{i\}} u_j^t(\mathbf{c})}{\sum_{t \in T} (u_{max}^t(i) - u_i^t(\mathbf{c}))} = \frac{\sum_{j \in N\backslash\{i\}} u_j(\mathbf{c})}{n\text{Prop}_i - u_i(\mathbf{c})} \leqslant \frac{(n-1)}{(n-1)\text{Prop}_i} = 1,$$

$$(2.1)$$

where the penultimate transition follows because we normalized utilities to achieve $u_j(\mathbf{c}) = 1$ for every $j \in S\backslash\{i\}$, every $j \in N\backslash S$ satisfies $u_j(\mathbf{c}) = 0$, and agent i does not receive her proportional share. The final transition holds due to our normalization $\text{Prop}_i = 1$.

Let $\delta = \sum_{j \in S\backslash\{i\}} u_j^{t^*}(\mathbf{c})$. Then, Equation (2.1) implies $u_i(\mathbf{c}') - u_i(\mathbf{c}) = u_{max}^{t^*}(i) - u_i^{t^*}(\mathbf{c}) \geqslant \delta$. Thus,

$$u_i(\mathbf{c}) + \delta \leqslant u_i(\mathbf{c}') < 1, \tag{2.2}$$

where the last inequality follows because the original outcome $\mathbf{c}$ violated Prop1 for agent i. In particular, this implies $\delta < 1$. Our goal is to show that $\prod_{j \in S} u_j(\mathbf{c}') > u_i(\mathbf{c}) = \prod_{j \in S} u_j(\mathbf{c})$, where the last equality holds due to our normalization $u_j(\mathbf{c}) = 1$ for $j \in S\backslash\{i\}$ and because $i \in S$. This would be a contradiction because $\mathbf{c}$ maximizes the product of utilities of agents in S. Now,

$$\sum_{j \in S\backslash\{i\}} \max\{0, 1 - u_j(\mathbf{c}')\} = \sum_{j \in S\backslash\{i\}} \max\{0, u_j^{t^*}(\mathbf{c}) - u_j^{t^*}(\mathbf{c}')\} \leqslant \sum_{j \in S\backslash\{i\}} u_j^{t^*}(\mathbf{c}) = \delta,$$

34

where the first transition follows from setting $1 = u_j(\mathbf{c})$ (by our normalization) and noting that $\mathbf{c}$ and $\mathbf{c}'$ are identical for all issues except t^*, and the second because all utilities are non-negative.

Hence, Lemma 5 implies that $\prod_{j \in S \setminus \{i\}} u_j(\mathbf{c}') \geqslant 1 - \delta$. Thus,

$$\prod_{j \in S} u_j(\mathbf{c}') \geqslant (1 - \delta) \cdot (u_i(\mathbf{c}) + \delta) = u_i(\mathbf{c}) + \delta \cdot (1 - u_i(\mathbf{c})) - \delta^2 > u_i(\mathbf{c}) + \delta^2 - \delta^2 = u_i(\mathbf{c}),$$

where the inequality holds because $u_i(\mathbf{c}) + \delta < 1$ from Equation (2.2). ∎

For private goods division, this result can be derived in a simpler fashion. Caragiannis et al. [49] already show that MNW satisfies PO. In addition, they also show that MNW satisfies EF1, which implies Prop1 due to our next result. To be consistent with the goods division literature, we use proportionality up to one good (rather than one issue) in the private goods division context.

Lemma 6. *For private goods division, envy-freeness up to one good (EF1) implies proportionality up to one good (Prop1).*

Proof. Take an instance of private goods division with a set of agents N and a set of goods M. Let $\mathbf{A}$ be an allocation satisfying EF1. Fix a agent $i \in N$.

Due to the definition of EF1, there must exist[4] a set of goods $X = \{g_j\}_{j \in N \setminus \{i\}}$ such that $u_i(A_i) \geqslant u_i(A_j) - u_i(g_j)$ for every $j \in N \setminus \{i\}$. Summing over all $j \in N \setminus \{i\}$, we get

$$(n-1) \cdot u_i(A_i) \geqslant \left(\sum_{j \in N \setminus \{i\}} u_i(A_j) \right) - u_i(X) \implies n \cdot u_i(A_i) \geqslant u_i(M) - u_i(X)$$

$$\implies u_i(A_i) + \frac{u_i(X)}{n} \geqslant \frac{u_i(M)}{n}. \quad (2.3)$$

---

[4] If $A_j = \varnothing$, we can add a dummy good g_j that every agent has utility 0 for, and make $A_j = \{g_j\}$.

Note that X has less than n goods. Suppose agent i receives good $g^* \in \arg\max_{g \in X} u_i(g)$. Note that $g^* \notin A_i$. Then, we have $u_i(A_i \cup \{g^*\}) \geqslant u_i(M)/n = \text{Prop}_i$, which implies that Prop1 is satisfied with respect to agent i. Because agent i was chosen arbitrarily, we have that EF1 implies Prop1. ■

Equation 2.3 in the proof of Lemma 6 directly implies the following lemma because the set X in the equation contains at most $n-1$ goods.

Lemma 7. *Let* **A** *be an allocation of private goods that satisfies EF1. Then, for every agent i,*

$$u_i(A_i) \geqslant Prop_i - \frac{1}{n} \sum_{t=1}^{n-1} u_{max}^{(t)}(i),$$

where $u_{max}^{(t)}(i)$ is the utility agent i derives from her t^{th} most valued good.

Next, we turn our attention to RRS and PPS. While MNW does not satisfy either of them, it approximates both.

Theorem 8. *The MNW solution satisfies $1/n$-RRS (and therefore $1/n$-PPS). The approximation is tight for both RRS and PPS up to a factor of $O(\log n)$.*

Proof. We first show the lower bound. Fix an instance of public decision making, and let **c** denote an MNW outcome. Let $S \subseteq N$ denote the set of agents that achieve positive utility under **c**.

Without loss of generality, let us normalize the utilities such that $u_j(\mathbf{c}) = 1$ for every $j \in S$. Suppose for contradiction that for some agent i, $u_i(\mathbf{c}) < (1/n) \cdot \text{RRS}_i$. First, this implies that $\text{RRS}_i > 0$, which in turn implies that agent i must be able to derive a positive utility from at least n different issues. By an argument identical to that used to argue that $u_i(\mathbf{c}) > 0$ in the proof of Theorem 4, it can be shown that we must also have $u_i(\mathbf{c}) > 0$ in this case (i.e., $i \in S$).

Now, recall that the sequence $\langle u_{max}^{(1)}(i), \ldots, u_{max}^{(m)}(i) \rangle$ contains the maximum utility agent i can derive from different issues, sorted in a non-ascending order. For every $q \in [p]$, let

$$t_q = \underset{(q-1)n+1 \leqslant t \leqslant qn}{\arg\min} \sum_{j \in S \setminus \{i\}} u_j^t(\mathbf{c}).$$

That is, we divide the public decision making into sets of n issues, grouped by agent i's maximum utility for them, and for each set of issues, we let t_q be the one that the remaining agents derive the lowest total utility from. Note that $t_q \leqslant qn$ for each $q \in [p]$, and therefore $u_{max}^{(t_q)}(i) \geqslant u_{max}^{(qn)}(i)$.

We will show that outcome $\mathbf{c}'$, where $c'_{t_q} = a_{max}^{t_q}(i)$ for all $q \in [p]$ and $c'_t = c_t$ for all other issues t, achieves a higher product of utilities to agents in S than $\mathbf{c}$ does, which is a contradiction because $\mathbf{c}$ is an MNW outcome. First, note that

$$u_i(\mathbf{c}') \geqslant \sum_{q=1}^{p} u_{max}^{(t_q)}(i) \geqslant \sum_{k=1}^{p} u_{max}^{(k \cdot n)}(i) = \mathrm{RRS}_i > n.$$

Further, we have

$$\sum_{j \in S \setminus \{i\}} \max\{0, 1 - u_j(\mathbf{c}')\} = \sum_{j \in S \setminus \{i\}} \max\{0, u_j(\mathbf{c}) - u_j(\mathbf{c}')\}$$

$$= \sum_{j \in S \setminus \{i\}} \sum_{q=1}^{p} \max\{0, u_j^{t_q}(\mathbf{c}) - u_j^{t_q}(\mathbf{c}')\} \leqslant \sum_{j \in S \setminus \{i\}} \sum_{q=1}^{p} u_j^{t_q}(\mathbf{c}),$$

where the first equality follows from our normalization, the second because $\mathbf{c}$ and $\mathbf{c}'$ only differ on issues $\{t_q\}_{q \in [p]}$, and the last because all utilities are non-negative.

Reversing the order of the summation and further manipulating the expression, we have

$$\sum_{q=1}^{p} \sum_{j \in S \setminus \{i\}} u_j^{t_q}(\mathbf{c}) \leqslant \sum_{q=1}^{p} \frac{1}{n} \sum_{t=(q-1)n+1}^{qn} \sum_{j \in S \setminus \{i\}} u_j^t(\mathbf{c}) = \frac{1}{n} \sum_{t=1}^{pn} \sum_{j \in S \setminus \{i\}} u_j^t(\mathbf{c}) \leqslant \frac{n-1}{n},$$

where the first transition follows from the definition of t_q. By Lemma 5, we have

$$\prod_{j\in S} u_j(\mathbf{c}') = u_i(\mathbf{c}') \prod_{j\in S\setminus\{i\}} u_j(\mathbf{c}') > n \cdot \left(1 - \frac{n-1}{n}\right) = 1 = \prod_{j=1}^{n} u_j(\mathbf{c}),$$

where the inequality holds because $u_i(\mathbf{c}') \geqslant \mathrm{RRS}_i > n \cdot u_i(\mathbf{c}) = n$, as agent i receives her round robin share under $\mathbf{c}'$ but did not even receive a $1/n$ fraction of it under $\mathbf{c}$. Hence, outcome $\mathbf{c}'$ achieves a higher product of utilities to agents in S than $\mathbf{c}$ does, which is a contradiction.

For the upper bound, Consider a public decision making problem with n issues, where each issue t has two alternatives a_1^t and a_2^t with the following utilities to the agents. Let $x = (\log n - \log\log n)/n$.

$$
\begin{array}{c|cc}
 & a_1^1 & a_2^1 \\
\hline
u_1^1 & 1 & d \\
u_2^1 & 0 & \mathrm{x} \\
\vdots & \vdots & \vdots \\
u_n^1 & 0 & \mathrm{x}
\end{array}
\quad \text{and} \quad
\begin{array}{c|cc}
 & a_1^t & a_2^t \\
\hline
u_1^t & 1 & d \\
u_2^t & 0 & 0 \\
\vdots & \vdots & \vdots \\
u_t^t & 0 & 1 \\
\vdots & \vdots & \vdots \\
u_n^t & 0 & 0
\end{array}
\quad \text{for } t \in \{2,\ldots,n\}.
$$

We choose the value of d later. Note that $\mathrm{PPS}_1 = 1$ as long as $d < 1$. Our goal is to make the MNW outcome choose alternative a_2^t for every issue t. Let us denote this outcome by $\mathbf{c}$. Then, the Nash welfare under $\mathbf{c}$ is

$$(n \cdot d) \cdot (1 + x)^{n-1}. \tag{2.4}$$

Let us find conditions on d under which this is greater than the Nash welfare that other possible outcomes $\mathbf{c}'$ could achieve.

Clearly, if $c_1' = a_1^1$ and $c_t' = a_1^t$ for any $t \neq 1$, then $u_t(\mathbf{c}') = 0$, and therefore $NW(\mathbf{c}') = 0$. Let us consider $\mathbf{c}'$ under which $c_1' = a_1^1$ and $c_t' = a_2^t$ for all $t \neq 1$. The Nash welfare produced by this outcome is

$$1 + (n-1)d. \tag{2.5}$$

38

Next, consider $\mathbf{c}'$ with $c_1' = a_2^1$, $c_{t*}' = a_1^{t*}$ for some $t^* \neq 1$, and $c_t' = a_2^t$ for all remaining t. The Nash welfare under this outcome is

$$(1 + (n-1)d) \cdot x \cdot (1 + x)^{n-2}. \tag{2.6}$$

We do not need to consider outcomes $\mathbf{c}'$ with $c_1' = a_2^1$ and $c_t' = a_1^t$ for multiple $t \neq 1$. This is because if switching the outcome from a_2^{t*} to a_1^{t*} for even a single $t^* \neq 1$ decreases the Nash welfare, switching the outcomes on other $t \neq 1$ would only further decrease the Nash welfare, as it would reduce the utility to another agent t by the same factor $1/(1+x)$, while increasing the utility to agent 1 by an even smaller factor.

Let us identify the conditions on d required for the quantity in Equation (2.4) to be greater than the quantities in Equations (2.5) and (2.6). We need

$$(n \cdot d) \cdot (1+x)^{n-1} > (1 + (n-1)d) \cdot 1$$

$$\Leftrightarrow n \cdot d > \frac{1}{(1+x)^{n-1} - 1 + 1/n}, \tag{2.7}$$

and

$$(n \cdot d)(1+x)^{n-1} > (1 + (n-1)d) \cdot x \cdot (1+x)^{n-2}$$

$$\Leftrightarrow n \cdot d > \frac{n \cdot x}{n + x}. \tag{2.8}$$

It is easy to check that for $x = (\log n - \log \log n)/n$, the quantities on the RHS of both Equations (2.7) and (2.8) are $O(\log n/n)$. Hence, we can set d to be sufficiently low for $n \cdot d$ to be $\Theta(\log n/n)$. However, note that $n \cdot d$ is precisely the approximation to PPS achieved for agent 1 under $\mathbf{c}$, as required. $\blacksquare$

For private goods, we can show that the MNW solution provides much better approximations to both RRS and PPS, as a result of its strong fairness guarantee of EF1.

Lemma 9. *For private goods division, envy-freeness up to one good (EF1) implies PPS and $n/(2n-1)$-RRS, but does not imply $n/(2n-2)$-RRS.*

Proof. Let **A** be an allocation of private goods that satisfies EF1. First, we show that **A** must also satisfy PPS. Suppose for contradiction that it violates PPS. Then, there exists a agent i such that $u_i(A_i) < \text{PPS}_i$, which in turn implies that $|A_i| < p$. Because the average number of goods per agent is $\frac{m}{n} \geqslant p$, there must exist a agent j such that $|A_j| > p$. Hence, for any good $g \in A_j$, agent j has at least p goods even after removing g from A_j, which implies $u_i(A_j \backslash \{g\}) \geqslant \text{PPS}_i > u_i(A_i)$. However, this contradicts the fact that **A** is EF1.

We now show that **A** also satisfies 1/2-RRS. By Lemma 7, we have

$$u_i(A_i) \geqslant \frac{1}{n} \sum_{t=n}^{m} u_{max}^{(t)}(i) \geqslant \frac{1}{n} u_{max}^{(n)}(i) + \sum_{k=2}^{p} u_{max}^{(k \cdot n)}(i). \tag{2.9}$$

Further, since **A** satisfies EF1, it must be the case that

$$u_i(A_i) \geqslant u_{max}^{(n)}(i). \tag{2.10}$$

To see this, suppose for contradiction that $u_i(A_i) < u_{max}^{(n)}(i)$, which implies that agent i is not allocated any of her n most valued goods. Therefore, by the pigeonhole principle, there must exist a agent $j \in N \backslash \{i\}$ that is allocated at least two of these goods. Hence, for any $g \in A_j$, we have $u_i(A_j \backslash \{g\}) \geqslant u_{max}^{(n)}(i) > u_i(A_i)$, which violates EF1. Finally, adding n times Equation (2.9) with $n-1$ times Equation (2.10), we obtain

$$(2n-1) \cdot u_i(A_i) \geqslant n \cdot u_{max}^{(n)}(i) + n \cdot \sum_{k=2}^{p} u_{max}^{(k \cdot n)}(i) = n \cdot \text{RRS}_i,$$

which implies the desired $n/(2n-1)$-RRS guarantee.

For the upper bound, consider an instance with n agents and n^2 goods, and define

agent 1's utility function to be

$$u_1(g_j) = \begin{cases} 1 & 1 \leqslant j \leqslant n, \\ \frac{1}{n-1} & n+1 \leqslant j \leqslant n^2. \end{cases}$$

Note that $\text{RRS}_1 = 1 + (n-1)\frac{1}{n-1} = 2$. Consider the allocation $\mathbf{A}$ with $A_1 = \{g_{n+1}, \ldots, g_{2n}\}$, $A_2 = \{g_1, g_2\}$, and $A_i = \{g_i, g_{(i-1)n+1}, \ldots, g_{i \cdot n}\}$ for all agents $i > 2$. Let the utilities of agents 2 through n be positive for the goods they receive and zero for the remaining goods. Hence, they clearly do not envy any agents. For agent 1, we have $u_1(A_1) = \frac{n}{n-1}$, $u_1(A_2 \backslash \{g_2\}) = 1$, and $u_1(A_i \backslash \{g_i\}) = \frac{n}{n-1}$ for all $i > 2$. That is, agent 1 does not envy any other agent up to one good. Hence, $\mathbf{A}$ satisfies EF1, and agent 1 obtains a $\frac{n}{2n-2}$ fraction of her RRS share, as required. ∎

As a corollary of Lemma 9, EF1 implies 1/2-RRS, and this approximation is asymptotically tight. Further, because the MNW solution satisfies EF1, Lemma 9 immediately provides guarantees (lower bounds) for the MNW solution. However, the upper bound in the proof of Lemma 9 does not work for the MNW solution. Next, we establish a much weaker lower bound, leaving open the possibility that the MNW solution may achieve a constant approximation better than 1/2 to RRS.

Theorem 10. *For private goods division, the MNW solution satisfies PPS and $n/(2n-1)$-RRS. For every $\varepsilon > 0$, the MNW solution does not satisfy $(2/3 + \varepsilon)$-RRS.*

Proof. The lower bounds follow directly from Lemma 9 and the fact that the MNW solution satisfies EF1. For the upper bound, consider an instance with two agents and four goods. Agent 1 has utilities $(1 - \delta, 1 - \delta, 1/2, 1/2)$ and agent 2 has utilities $(1, 1, 0, 0)$ for goods (g_1, g_2, g_3, g_4), respectively. Note that $\text{RRS}_1 = 3/2 - \delta$. The MNW allocation $\mathbf{A}$ is given by $A_1 = \{g_3, g_4\}$ and $A_2 = \{g_1, g_2\}$. Thus, $\frac{u_1(A_1)}{\text{RRS}_1} = \frac{2}{3 - 2\delta}$. The upper bound follows by setting δ sufficiently small. ∎

41

2.4 Computational Complexity

In Section 2.3, we showed that without requiring Pareto optimality, we can achieve both RRS (thus PPS) and Prop1 in polynomial time using the round robin method (Theorem 1). In contrast, when we require PO, the leximin mechanism (with an appropriate normalization of utilities) provides RRS (thus PPS) and PO, while the MNW solution provides Prop1 and PO. However, both these solutions are $\mathcal{NP}$-hard to compute [140, 28]. This raises a natural question whether we can efficiently find outcomes satisfying our fairness guarantees along with PO. For PPS, the answer is negative.

Theorem 11. *It is $\mathcal{NP}$-hard to find an outcome satisfying PPS and PO.*

Note that it is the *search problem* of finding an outcome (any outcome) satisfying PPS and PO for which we prove computational hardness; the *decision problem* of testing the existence of such an outcome is trivial as we know it always exists. Before we prove this result, we need to introduce a new (to our knowledge) decision problem and show that it is $\mathcal{NP}$-complete.

Exact Triple-Cover by 3-sets (X33C): An instance $(Y, \mathcal{T})$ of X33C is given by a set Y of r vertices and a set $\mathcal{T} = \{T_1, T_2, \ldots, T_m\}$, where each T_i is a set of three vertices. The decision problem is to determine whether it is possible to choose r sets, with repetition allowed, such that every vertex v is contained in exactly three of the chosen sets (an exact triple-cover).

Let us contrast this with the definition of the popular NP-complete problem, Exact Cover by 3-sets (X3C): An instance $(X, \mathcal{S})$ of X3C is given by a set X of $3q$ vertices and a set $\mathcal{S} = \{S_1, \ldots, S_n\}$, where each S_i is a set of three vertices. The decision problem is to determine if there exists a subset of $\mathcal{S}$ of size q that covers every vertex $x \in X$ exactly once (an exact cover).

Lemma 12. *X33C is NP-complete.*

Proof. Clearly, X33C lies in NP because a triple-cover can be checked in polynomial time. To show hardness, we reduce from X3C. Given an instance $(X, \mathcal{S})$ of X3C, divide X into q sets of 3 vertices arbitrarily, indexed by k. For every one of these q sets of three vertices $s^k = \{s_1^k, s_2^k, s_3^k\}$, create 8 new vertices, $\{s_{i,j}^k : i \in [2], j \in [4]\}$, and 10 new sets $\{T_{i,j}^k : i \in [2], j \in [5]\}$. The sets $T_{i,j}^k$ are defined as follows: $T_{i,1}^k = \{s_{i,1}^k, s_{i,2}^k, s_{i,3}^k\}$, $T_{i,2}^k = \{s_{i,2}^k, s_{i,3}^k, s_{i,4}^k\}$, $T_{i,3}^k = \{s_{i,1}^k, s_{i,2}^k, s_{i,4}^k\}$, $T_{i,4}^k = \{s_{i,1}^k, s_{i,3}^k, s_1^k\}$, and $T_{i,5}^k = \{s_{i,4}^k, s_2^k, s_3^k\}$.

The X33C instance is given by $(Y = X \cup \{s_{i,j}^k : i \in [2], j \in [4], k \in [q]\}, \mathcal{T} = \mathcal{S} \cup \{T_{i,j}^k : i \in [2], j \in [5], k \in [q]\})$. Note that $|Y| = 11q$. We show that $(Y, \mathcal{T})$ has an exact triple-cover if and only if there exists an exact cover for $(X, \mathcal{S})$.

First, suppose that there exists an exact cover for $(X, \mathcal{S})$. Then there exists an exact triple-cover for $(Y, \mathcal{T})$ by selecting sets $T_{i,j}^k$ for every $k \in [q]$, $i \in [2]$, and $j \in [5]$. It is easy to check that these $10q$ sets cover each $s_{i,j}^k$ exactly three times, as well as covering s_k^k exactly twice, for all $k \in [q]$ and $k \in [3]$. Hence, all we need to do is add the solution to the original X3C instance.

Now, suppose that there exists an exact triple-cover by 3-sets for the X33C instance. This implies that, for any k and i, exactly three sets from $\{T_{i,1}^k, T_{i,2}^k, T_{i,3}^k\}$ must be chosen (recall that we can choose the same set more than once), because these are the only sets that contain $s_{i,2}^k$, which must be covered exactly three times. We now consider how we can choose these three sets. Suppose that $T_{i,2}^k$ is chosen more than once. Then only (at most) one of $T_{i,1}^k$ and $T_{i,3}^k$ is chosen, so we still need to cover $s_{i,1}^k$ (at least) twice. The only way to do this is by choosing $T_{i,4}^k$ twice. But then $s_{i,3}^k$, which is contained in both $T_{i,2}^k$ and $T_{i,4}^k$, is covered more than three times, a violation of the conditions of an exact triple-cover. By similar reasoning, we can show that $T_{i,3}^k$ cannot be chosen more than once. Now suppose that $T_{i,1}^k$ is chosen

twice. Then it remains to choose exactly one of $T_{i,2}^k$ and $T_{i,3}^k$; suppose WLOG that we choose $T_{i,2}^k$. Then we still need to cover $s_{i,1}^k$ an additional time. The only way to do this is to choose $T_{i,4}^k$, which also covers $s_{i,3}^k$, meaning that $s_{i,3}^k$ is covered four times, violating the condition of the exact triple-cover. Finally, suppose that $T_{i,1}^k$ is chosen three times. Then we still need to cover $s_{i,4}^k$ three times without covering any of $s_{i,1}^k$, $s_{i,2}^k$, or $s_{i,3}^k$ again. We therefore need to choose $T_{i,5}^k$ three times. Otherwise, we can choose each of $T_{i,1}^k$, $T_{i,2}^k$, $T_{i,3}^k$, $T_{i,4}^k$, and $T_{i,5}^k$ once each, which covers each of $s_{i,1}^k$, $s_{i,2}^k$, and $s_{i,3}^k$ once each. All other options have been ruled out. In particular, it is necessary to choose $T_{i,5}^k$ at least once.

So there are two options. Regardless of which option we choose, we still have to cover each of $s_{i',1}^k$, $s_{i',2}^k$, and $s_{i',3}^k$ three times each, for $i' \neq i$. Since the options for i' are symmetric to those for i, it is again necessary to choose $T_{i',5}^k$ at least once. If we choose $T_{i,1}^k$ three times and $T_{i,5}^k$ three times, as well as $T_{i',5}^k$ at least once (as we must), then s_2^k and s_3^k are covered at least four times, a violation of the exact triple-cover. Therefore the only possibility is to choose each of $T_{i,1}^k$, $T_{i,2}^k$, $T_{i,3}^k$, $T_{i,4}^k$, and $T_{i,5}^k$ once. Similarly, we must choose each of $T_{i',1}^k$, $T_{i',2}^k$, $T_{i',3}^k$, $T_{i',4}^k$, and $T_{i',5}^k$ once each also. And, since k was arbitrary, this holds for all $k \in [q]$.

So, for all $k \in [q]$ and all $i \in [2]$, each of $T_{i,1}^k$, $T_{i,2}^k$, $T_{i,3}^k$, $T_{i,4}^k$, and $T_{i,5}^k$ is chosen exactly once, a total of $10q$ sets chosen. We therefore have q more sets to choose, which necessarily cover each of $v \in \mathcal{S}$ exactly once (note that each $v \in \mathcal{S}$ corresponds to an s_j^k for some $k \in [q]$ and $j \in [3]$, and these are covered exactly once by either $T_{1,4}^k$ or $T_{1,5}^k$, and exactly once again by either $T_{2,4}^k$ or $T_{2,5}^k$). The only way to choose q sets that cover each $v \in \mathcal{S}$ exactly once is by choosing an exact cover for the instance $(X, \mathcal{S})$. $\blacksquare$

Using this lemma, we can now show that finding an outcome satisfying PPS and PO is $\mathcal{NP}$-hard through a reduction from X33C.

Proof of Theorem 11. Let $(Y, \mathcal{T})$ be an instance of X33C, with $|Y| = r$. Let $\varepsilon \in (0, 1/(3r))$. We define a public decision making problem as follows. There are r agents, one corresponding to each vertex $v \in Y$, and r issues. For each issue, there are $|Y| + |\mathcal{T}|$ alternatives. For each issue t and each agent i, there is an alternative $a_{t,i}$ which is valued at $1 - \varepsilon$ by agent i, and 0 by all other agents. The remaining $|\mathcal{T}|$ alternatives correspond to the 3-sets from the X33C instance. For a set $T_j \in \mathcal{T}$, the corresponding alternative is valued at $\frac{1}{3}$ by agents $i \in T_j$, and valued at 0 by all other agents. Note that $\mathrm{PPS}_i = 1 - \varepsilon$ for each agent i, because there are exactly as many issues as agents, and each agent values its most preferred alternative for each issue at $1 - \varepsilon$.

We now show that there exists an exact triple-cover by 3-sets if and only if all outcomes to the public decision making problem that satisfy PPS and PO have $u_i(\mathbf{c}) = 1$ for all i. First, suppose that there exists an exact triple-cover by 3-sets. We need to show that all outcomes satisfying PPS and PO have $u_i(\mathbf{c}) = 1$ for all i. So suppose otherwise – that there exists an outcome satisfying PPS and PO with $u_i(\mathbf{c}) \neq 1$ for some agent i. In particular, some agent must have $u_i(\mathbf{c}) > 1$, otherwise $\mathbf{c}$ is not PO (because it is possible to choose an outcome corresponding exactly to an exact triple-cover, which gives each agent utility 1). But agents only derive utility in discrete amounts of $1 - \varepsilon$ or $\frac{1}{3}$, which means that any agent with $u_i(\mathbf{c}) > 1$ has $u_i(\mathbf{c}) \geqslant \frac{4}{3} - \varepsilon$.

$$\sum_{i=1}^{r} u_i(\mathbf{c}) \geqslant \frac{4}{3} - \varepsilon + \sum_{i=1}^{r-1}(1 - \varepsilon) = \frac{1}{3} + r - r\varepsilon > r,$$

where the last inequality holds because $\varepsilon < 1/(3r)$. However, this is a contradiction because each alternative in each of the r issues contributes at most 1 to the social welfare. Therefore, every outcome satisfying PPS and PO has $u_i(\mathbf{c}) = 1$ for all i.

Next, suppose that there does not exist an exact triple-cover by 3-sets. So if

45

we choose an alternative corresponding to a 3-set for every issue, it is not possible for every agent to derive utility 1. Therefore, some agent must derive utility $\frac{2}{3}$ (or lower), which violates PPS. Thus, every outcome that satisfies PPS must include at least one issue where the chosen alternative is one that corresponds to a agent, not to a 3-set. Such an alternative only contributes $1 - \varepsilon$ to social welfare. Therefore, the social welfare is strictly less than r, which means that some agent gets utility strictly less than 1. Therefore, there is no outcome satisfying PPS (either with or without PO) such that $u_i(\mathbf{c}) = 1$ for all i. Since the set of outcomes satisfying PPS is always non-empty, it is therefore not the case that all outcomes satisfying PPS and PO have $u_i(\mathbf{c}) = 1$ for all i. ∎

Because every outcome satisfying RRS also satisfies PPS, we have the following corollary.

Corollary 13. *It is $\mathcal{NP}$-hard to find an outcome satisfying RRS and PO.*

For private goods division, we show, in stark contrast to Theorem 11, that we can find an allocation satisfying PPS and PO in polynomial time. This is achieved using Algorithm 1. Interestingly, it produces not an arbitrary allocation satisfying PPS and PO, but an allocation that assigns at least $p = \lfloor m/n \rfloor$ goods to every agent — implying PPS, and maximizes weighted (utilitarian) social welfare according to some weight vector — implying PO.

At a high level, the algorithm works as follows. It begins with an arbitrary weight vector $\mathbf{w}$, and an allocation $\mathbf{A}$ maximizing the corresponding weighted (utilitarian) social welfare. Then, it executes a loop (Lines 3-22) while there exists a agent receiving less than p goods, and each iteration of the loop alters the allocation in a way that one of the agents who received more than p goods loses a good, one of the agents who received less than p goods gains a good, and every other agent retains the same number of goods as before.

ALGORITHM 1: Polynomial time algorithm to achieve PPS and PO for private goods

Input: The set of agents N, the set of private goods M, and agents' utility functions $\{u_i\}_{i \in N}$

Output: A deterministic allocation $\mathbf{A}$ satisfying PPS and PO

1   $\mathbf{w} \leftarrow (1/n, \ldots, 1/n) \in \mathbb{R}^n$;

2   $\mathbf{A} \leftarrow \arg\max_{\mathbf{A}'} \sum_{i \in N} w_i \cdot u_i(\mathbf{A}')$;

3   **while** $\exists i \in N, |A_i| < p$ **do**   /* Until every agent receives at least $p = \lfloor m/n \rfloor$ goods */

4     $GT \leftarrow \{i \in N : |A_i| > p\}$;   /* Partition agents by the number of goods they receive */

5     $EQ \leftarrow \{i \in N : |A_i| = p\}$;

6     $LS \leftarrow \{i \in N : |A_i| < p\}$;

7     $DEC = GT$;   /* Agents whose weights we will decrease */

8   **while** $DEC \cap LS = \varnothing$ **do**

         /* Minimally reduce weights of agents in DEC so a agent in DEC loses a good   */

9        $(i^*, j^*, g^*) \leftarrow \arg\min_{i \in DEC, j \in N \setminus DEC, g \in A_i} (w_i \cdot v_{i,g})/(w_j \cdot v_{j,g})$;

10       $r \leftarrow (w_{i*} \cdot v_{i*,g*})/(w_{j*} \cdot v_{j*,g*})$;

11       $\forall i \in DEC, \ w_i \leftarrow w_i/r$;

12       $DEC \leftarrow DEC \cup \{j^*\}$;

13       $D(j^*) \leftarrow (i^*, g^*)$;   /* Bookkeeping: j^* can receive g^* from i^* */

14   **end**

15     $j^* \leftarrow DEC \cap LS$;   /* Agent from LS who receives a good */

16   **while** $j^* \notin GT$ **do**

17       $(i^*, g^*) \leftarrow D(j^*)$;

18       $A_{i*} \leftarrow A_{i*} \setminus \{g^*\}$;

19       $A_{j*} \leftarrow A_{j*} \cup \{g^*\}$;

20       $j^* \leftarrow i^*$;

21   **end**

22   **end**

23   **return A**;

Each iteration of the loop maintains a set DEC of agents whose weight it reduces. Initially, DEC consists of agents who have more than p goods (Line 7). When the weights are reduced enough so that a agent in DEC is about to lose a good to a agent, necessarily outside DEC (Lines 9-11), the latter agent is added to DEC (Line 12) before proceeding further. When a agent who has less than p goods is added to DEC, this process stops and the algorithm leverages the set of ties it created along the way to make the aforementioned alteration to the allocation (Lines 16-21).

We now formally state that this produces an allocation satisfying PPS and PO, and that it runs in polynomial time; the proof is deferred to the appendix.

Theorem 14. *For private goods division, PPS and PO can be satisfied in polynomial time.*

The complexity of finding an allocation (of private goods) satisfying the stronger guarantee RRS along with PO in polynomial time remains open, as does the complexity of finding an allocation satisfying Prop1 and PO.

We note that the convenient approach of weighted welfare maximization we use in Theorem 14 cannot be used for finding an outcome satisfying RRS and PO, as the following example shows. This leads us to conjecture that it may be $\mathcal{NP}$-hard to find such an outcome.

Example 3. *Consider a private goods division problem with two agents and four goods. Agent 1 has utilities $(4, 4, 1, 1)$ and agent 2 has utilities $(3, 3, 2, 2)$ for goods (g_1, g_2, g_3, g_4), respectively. Note that $RRS_1 = RRS_2 = 5$. Consider assigning weights w_1 and w_2 to agents 1 and 2, respectively. If $4w_1 > 3w_2$, i.e., $w_1 > 3w_2/4$ then agent 1 receives both g_1 and g_2, which means that agent 2 receives utility less than her RRS share. On the other hand, if $3w_2 > 4w_1$, i.e., $w_1 < 3w_2/4$ then agent 2 receives both g_1 and g_2, which means that agent 1 receives utility less than her RRS share.*

The only remaining possibility is that $w_1 = 3w_2/4$, but in that case, agent 2 receives both g_3 and g_4. Regardless of how we divide goods g_1 and g_2, one of the two agents still receives utility less than her RRS share.

In contrast, a simple modification of Algorithm 1 seems to quickly find an allocation satisfying Prop1 and PO in hundreds of thousands of randomized simulations. At each iteration of this version, the set DEC initially consists of agents who attain their proportional share (it is easy to show using the Pigeonhole principle that this

set is non-empty for any weighted welfare maximizing allocation), and ends when a agent is added to DEC that is not currently achieving Prop1. Thus, at every loop, a agent that was receiving her proportional share may lose a good (but will still achieve at least Prop1), the agent added to DEC that was not achieving Prop1 gains a good, and some agents that were achieving Prop1 but not their proportional share may lose a good, but only if they gain one too. These three classes of agents are therefore analogous to agents with more than p goods, less than p goods, and exactly p goods in Algorithm 1. Unfortunately, we are unable to prove termination of this algorithm because it is possible that a agent who achieves Prop1 but not her proportional share loses a high-valued good while gaining a low-valued good, thus potentially sacrificing Prop1. Thus we do not get a property parallel to the key property of Algorithm 1, that no agent's utility ever drops below her PPS share, after she attains it. However, our algorithm always seems to terminate quickly and finds an allocation satisfying Prop1 and PO in our randomized simulations, which leads us to conjecture that it may be possible to find an allocation satisfying Prop1 and PO in polynomial time, either from our algorithm directly or via some other utilization of weighted welfare maximization.

2.5 Discussion

We introduced several novel fairness notions for public decision making and considered their relationships to existing mechanisms and fairness notions. Throughout the paper, we highlighted various open questions including the existence (and complexity) of a mechanism satisfying RRS, Prop1, and PO, the complexity of finding an outcome satisfying Prop1 and PO (for public decisions and private goods), the complexity of finding an outcome satisfying RRS and PO (for private goods), and whether MNW provides a constant approximation to RRS better than 1/2.

So far we only considered deterministic outcomes. If randomized outcomes are

allowed (an alternative interpretation in the private goods case is that the goods are infinitely divisible), then the MNW solution satisfies Prop as a direct consequence of it satisfying Prop1 for deterministic outcomes (Theorem 4).[5] To see this, consider replicating each issue K times and dividing utilities by K. The relative effect of granting a single agent control of a single issue becomes negligible. Thus, as K approaches infinity, the utility of each agent i in an MNW outcome approaches a value that is at least their proportional share $Prop_i$. The fraction of copies of issue t in which outcome a_j^t is selected can be interpreted as the weight placed on a_j^t in the randomized outcome. Because RRS, PPS, and Prop1 are relaxations of Prop, the randomized MNW outcome also satisfies all of them.

For private goods division, this can be seen as a corollary of the fact that the randomized MNW outcome satisfies envy-freeness, which is strictly stronger than proportionality. This hints at a very interesting question: Is there a stronger fairness notion than proportionality in the public decision making framework that generalizes envy-freeness in private goods division? Although such a notion would not be satisfiable by deterministic mechanisms, it may be satisfied by randomized mechanisms, or it could have novel relaxations that may be of independent interest. Recent work by Fain et al. [75] considered exactly this problem, using the stabiliity notion of the *core* as a fairness primitive. Although an outcome in the core is not guaranteed to exist, it can be approximated.

At a high level, our work provides a framework bringing together two long-studied branches of social choice theory — fair division theory and voting theory. Both have at their heart the aim to aggregate individual preferences into a collective outcome that is fair and agreeable, but approach the problem in different ways. Fair division theory typically deals with multiple private goods, assumes access to cardinal

---

[5] Of course, the *realization* may fail to satisfy Prop (and other desiderata), but the *lottery* is fair if we consider expected utilities.

utilities, and focuses on notions of fairness such as envy-freeness and proportionality. Voting theory, in contrast, typically deals with a single public decision (with the exception of combinatorial voting mentioned earlier), assumes access only to less expressive ordinal preferences, has the "one voter, one vote" fairness built inherently into the voting rules, and focuses on different axiomatic desiderata such as Condorcet consistency and monotonicity.

Of course, one can use a voting approach to fair division, since we can have agents express preferences over complete outcomes, and this approach has been used successfully to import mechanisms from voting to fair division and vice versa [24, 23]. However, not only does this approach result in an exponential blowup in the number of alternatives, it also does not provide a convenient way to express fair-division-like axioms. Continuing to explore connections between the two fields remains a compelling direction for future work.

3

Public Goods, Online Allocation

3.1   Introduction

In the previous chapter we studied fairness in settings where a single outcome, possibly consisting of several different issues, must be chosen at a fixed point in time. Fairness can be achieved by sacrificing some agents' utility on certain issues, but rewarding them with desirable outcomes on other issues.

Unfortunately, in real settings, it is often not possible to observe the full set of issues upfront. For instance, suppose that Alice and Bob go to lunch together every week. Neither knows exactly what they will want to eat in a week's time, so the best they can do is make decisions week-by-week. Some weeks, their preferences will differ, and every option will leave one of them unhappy. Fortunately, we can often address this unfairness over time—Alice gets her most preferred restaurant today, and Bob gets his next week.

Achieving fairness over time is the topic of this chapter. We adopt the public decision making model of Chapter 2, except that the issues arrive one at a time, and an alternative for issue t must be chosen before issue $t + 1$ arrives. Motivated

by the performance of the MNW solution in the offline model of Chapter 2, as well as its extensive use in the fair division setting [58, 64, 140], we adopt Nash welfare maximization as our goal in this work.[1]

When we make multiple decisions over time, we could simply maximize the Nash welfare in each round separately. But it is easy to see that this can lead to dominated outcomes. For example, suppose there are two agents, and we can choose an alternative that gives one a reward of 3, and the other a reward of 0; or vice versa; or an alternative that gives each of them 1. Within a round, the last alternative maximizes Nash welfare; but if this scenario is repeated every round, then it would be better to alternate between the first two alternatives, so that each agent obtains 1.5 per round on average. Of course, *initially*, say in the first round, we may not realize we will have these options every round, and so we may choose the last alternative; but if we do have these options every round, we should *eventually* catch on to this pattern and start alternating. Ideally, we would maximize the long-term Nash welfare, that is, the product of the long-run utilities (which are the sums of each agent's rewards), rather than, for example, the sum of the products within the rounds.

The rest of the chapter is organized as follows. In Section 3.2 we introduce notation and preliminaries. In Section 3.3 we present two simple greedy algorithms for choosing alternatives, and provide intuitive interpretations of them, including an axiomatic justification for one of them. After presenting the algorithms, we evaluate them on data from a computer systems application in Section 3.4.

Related work: In addition to work cited in Chapter 2, there is a body of literature studying social choice and fair division in dynamic settings. Parkes and Procaccia [129] examine a similar problem by modeling agents' evolving preferences with Markov Decision Processes, with a reward function defined over states and

---

[1] The leximin mechanism also performs well in Chapter 2, but it relies on a very particular scaling that requires knowing the entire instance upfront.

actions (alternatives). However, their goal is to maximize the sum of (discounted) rewards and they do not explicitly consider fairness as an objective. Kash et al. [108] examine a model of dynamic fair division where agents arrive at different times and must be allocated resources; however, they do not allow for the preferences of agents to change over time as we do. Aleksandrov et al. [17] consider an online fair division problem in a setting where items appear one at a time, and agents declare yes/no preferences over that item. In our setting, each round has many alternatives and agents express more general utilities. Our work is related to the literature on dynamic mechanism design (Parkes et al. [130] provide an overview), except that we do not consider monetary transfers. Guo et al. [99] consider a setting similar to ours, also without money, except that they are not explicitly interested in fairness, only welfare, and focus on incentive compatibility.

3.2   Preliminaries

We consider an online version of the model from Chapter 2. Recall that we have a set of agents $N = [n]$ and issues $T = [m]$. Each issue $t \in T$ has an associated set of alternatives $A^t = \{a_1^t, \ldots, a_{k_t}^t\}$, exactly one of which must be chosen. At round t, we observe a matrix $U^t = (u_i^t(a_j^t))_{i \in [n], j \in [k_t]}$ of utilities; $u_i^t(a_j^t)$ is the utility that agent i receives from alternative a_j^t being chosen. In this chapter, we will require $u_i^t(a_j^t) \in \mathbb{N}$, which is necessary for some of our results in Section 3.3. This is still sufficient for agents to express their preferences to arbitrary levels of precision. Let $\mathbf{u^t}(\mathbf{a_j^t})$ denote the j-th column of matrix U^t, the vector of valuations for alternative a_j^t.

For every round t, a (dynamic) mechanism chooses a set of alternatives C_t, from which a single alternative c_t is chosen arbitrarily. Importantly, the problem is *online*, so we may only use information up to time t in order to choose C_t.

We define a vector of *accrued rewards at round t*, $\mathbf{r_t}$, where the accrued reward of agent i at round t is the sum of i's utilities for the chosen alternatives up to and

including round t: $r_t(i) = \sum_{t'=1}^{t} u_i^{t'}(c_{t'})$. We will often be interested in an agent's accrued reward *before* the start of round t, $r_{t-1}(i)$. For convenience, we will refer to the set of agents with $r_{t-1}(i) = 0$ by I_0 when the round, t, is clear. The average utility of the agents over the first t rounds is $\mathbf{r_t^{avg}} = \frac{1}{t}\mathbf{r_t}$.

A dynamic mechanism is *anonymous* if applying permutation σ to the agents, for all t, does not change the set of chosen alternatives C_t, for any t. A dynamic mechanism is *neutral* if applying permutation σ to the alternatives, for all t, results in choosing alternatives $\sigma(C_t)$ for all t. For the rest of this chapter we only consider anonymous, neutral DSCFs.

The *Nash welfare* of valuation vector $\mathbf{r}$, $NW(\mathbf{r})$, is defined to be the product of the agents' utilities, $NW(\mathbf{r}) = \prod_{i=1}^{n} r(i)$. We also define $NW^+(\mathbf{r}) = \prod_{i:r(i) \neq 0} r(i)$ to be the product of all positive entries of $\mathbf{r}$. Our aim is to maximize the Nash welfare of the average utility across all T rounds, $NW(\mathbf{r_T^{avg}})$. Note that while our setting allows for discounting, we do not need to explicitly address it since the input matrices can be pre-multiplied by the necessary factor before being passed as input to the mechanism.

The benchmark algorithm is the optimal algorithm for the offline problem, where an instance is given by the set $\{U^t\}_{t \in T}$, and can be solved by a mixed integer convex program. We denote the optimal Nash welfare by OPT.

Our algorithms and analysis use a formal infinitesimal quantity ε. Numbers involving ε take the form $\sum_{i=-\infty}^{i=\infty} x_i \varepsilon^i$.[2] For two such numbers $x = \sum_{i=-\infty}^{i=\infty} x_i \varepsilon^i$ and $y = \sum_{i=-\infty}^{i=\infty} y_i \varepsilon^i$, let i' be the smallest index for which $x_i \neq y_i$, if it exists. Then $x > y$ if and only if $x_{i'} > y_{i'}$. That is, we compare numbers lexicographically by the lowest powers of ε. Two numbers are equal if all coefficients are equal.

---

[2] While our framework allows for unbounded powers of ε, here we utilize only powers of ε between ε^{-1} and ε^n.

3.3 Greedy Algorithms

3.3.1 Algorithm Definitions

In this section we present two greedy algorithms. We note that, although these algorithms are designed to give an approximate solution to that which maximizes Nash welfare, much of this section is devoted to showing that they satisfy desirable properties as algorithms in their own right. Such an approach is not new in computational social choice – several papers treat approximation algorithms as distinct voting rules [47, 48, 73]. The first algorithm, GREEDY, simply chooses c_t to maximize $NW(\mathbf{r_t^{\mathrm{avg}}})$, the Nash welfare at the end of the round. The second algorithm is a linearized version of greedy known as PROPORTIONALFAIR (PF) in the networking community [158, 103], which maximizes the sum of percentage increases in accrued reward at each round. Equivalently, it works by assigning each agent a weight w_i (denote the vector of weights by $\mathbf{w}$) equal to the inverse of her accrued reward at the start of each round and chooses $C_t = \arg\max_{a_j^t \in A^t} \mathbf{w} \cdot \mathbf{u^t}(\mathbf{a_j^t})$, the alternatives that maximize the weighted sum of valuations. Note that w_i is proportional to the product of the other agents' accrued rewards.

Example 4. *Let $n = m = 2$ and suppose that $r_{t-1}(1) = 1$, $r_{t-1}(2) = 3$, and $U^t = \left(\begin{smallmatrix} 2 & 3 \\ 3 & 1 \end{smallmatrix}\right)$. That is, agent 1 has valuation 2 for alternative a_1^t and valuation 3 for alternative a_2^t. Agent 2 has valuation 3 for alternative a_1^t and valuation 1 for alternative a_2^t. Choosing a_1^t results in Nash welfare of $(1 + 2) \cdot (3 + 3) = 18$, while choosing a_2^t results in Nash welfare of $(1 + 3) \cdot (3 + 1) = 16$. Thus GREEDY chooses a_1^t.*

Under PF, each agent is given weight inversely proportional to their own accrued utility. That is, agent 1 has weight 1 and agent 2 has weight $\frac{1}{3}$. Now, taking the weighted sum of valuations yields $(1 \cdot 2) + (\frac{1}{3} \cdot 3) = 3$ for alternative a_1^t, and $(1 \cdot 3) + (\frac{1}{3} \cdot 1) = \frac{10}{3}$ for alternative a_2^t. Thus PF chooses a_2^t.
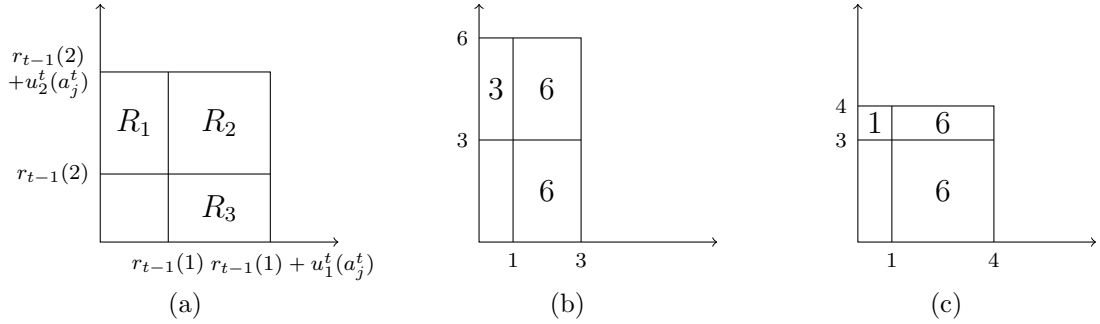
FIGURE 3.1: Illustration of the difference between GREEDY and PF for an instance with two agents. The horizontal axis represents agent 1's reward, and the vertical axis represents agent 2's reward. Figure 3.1(a) shows a general instance. GREEDY chooses the alternative that maximizes area $R_1 + R_2 + R_3$, while PF chooses the alternative that maximizes $R_1 + R_3 = v_2^t(a_j)r_{t-1}(1) + u_1^t(a_j^t)r_{t-1}(2) = r_{t-1}(1)r_{t-1}(2)\left[\frac{u_1^t(a_j^t)}{r_{t-1}(1)} + \frac{u_2^t(a_j^t)}{r_{t-1}(2)}\right]$. Figures 3.1(b) and 3.1(c) illustrate the choice of alternative a_1^t and a_2^t in Example 4, respectively.

A graphical illustration of the difference between the two algorithms is given in Figure 3.1.

Unfortunately, both algorithms encounter problems while there exist agents with zero accrued reward. For GREEDY, it can (and, unless some alternative is valued positively by all agents, will) be the case that $NW(\mathbf{r_t^{avg}}) = 0$ for all choices of c_t, even when one alternative is weakly preferred to all other alternatives by all agents. For PF, it is impossible to set a weight $w_i = \frac{1}{r_{t-1}(i)}$ for an agent with $r_{t-1}(i) = 0$.

As a general framework for addressing this issue, we endow each agent $i \in I_0$ with some arbitrary, infinitesimal reward at the start of each round. This is a natural way to allow the algorithms to give high priority to agents with zero accrued reward while avoiding mathematical inconsistencies, and it allows us to efficiently choose an alternative c_t if we are happy with selecting any member of the choice set C_t.

However, once we endow rewards (even infinitesimal ones), we immediately lose scale-freeness, one of the appealing properties of using Nash welfare. Further, if we want to choose a member of the choice set C_t uniformly at random, there is no obvi-

ous distribution over endowed rewards that allows us to do this – choosing endowed rewards uniformly at random from some interval will not, in general, result in drawing uniformly from C_t. So, while the technique of randomly endowing infinitesimal reward is a general and intuitive way for the algorithms to handle all situations, we also want an algorithm to compute the entire choice set C_t.

In the following, for both GREEDY and PF, we first present the algorithm to select a single alternative via nondeterministically endowing infinitesimal reward, followed by an algorithm to compute the entire choice set C_t.

ALGORITHM 2: GREEDY (select one alternative)

Input: $\mathbf{r_{t-1}}$

1 **for** $i = 1, \ldots, n$ **do**
2 $\quad$ Randomly choose $0 < x_i \leqslant 1$;
3 **end**
4 **return** $c_t \in \arg\max_{a_j^t \in A^t} \prod_{i=1}^{n} \max\{r_{t-1}(i) + u_i^t(a_j^t), x_i \varepsilon\}$;

The alternatives chosen by Algorithm 2 are exactly the alternatives that result in a maximal number of agents with positive accrued reward and, subject to holding fixed the set of agents with positive accrued reward, maximizes the product of these agents' rewards. For a single round, this reproduces the definition of the Maximum Nash Welfare mechanism from Chapter 2.

ALGORITHM 3: GREEDY (select all alternatives)

Input: r_{t-1}

1 $C_t \leftarrow \arg\max_{a_j^t \in A^t} |\{i : r_{t-1}(i) + u_i^t(a_j^t) > 0\}|$;
2 **for** $a_j^t \in C_t$ **do**
3 $\quad$ **if** $\exists j'$ such that $\{i : r_{t-1}(i) + u_i^t(a_j^t) > 0\} = \{i : r_{t-1}(i) + u_i^t(a_{j'}^t) > 0\}$ and
$\quad\quad NW^+(\mathbf{r_{t-1}} + \mathbf{u^t(a_j^t)}) < NW^+(\mathbf{r_{t-1}} + \mathbf{u^t(a_{j'}^t)})$ **then**
4 $\quad\quad$ $C_t \leftarrow C_t \backslash \{a_j^t\}$;
5 $\quad$ **end**
6 **end**
7 **return** C_t;

The version of PF for selecting a single alternative is presented as Algorithm 4.

ALGORITHM 4: PROPORTIONALFAIR (select one alternative)

Input: $\mathbf{r_{t-1}}$

1 **for** $i \in I_0$ **do**

2 $\quad$ Randomly choose $0 < x_i \leqslant 1$;

3 $\quad$ Randomly choose $y_i \in \mathbb{R}$;

4 **end**

5 $w_i \leftarrow \begin{cases} x_i \frac{1}{\varepsilon} + y_i, & \text{if } r_{t-1}(i) = 0 \\ \frac{1}{r_{t-1}(i)}, & \text{if } r_{t-1}(i) > 0 \end{cases}$;

6 **return** $c_t \in \arg\max_{a_j^t \in A^t} \mathbf{w} \cdot \mathbf{u^t(a_j^t)}$;

To determine the complete choice set C_t, we solve a linear program for each alternative that explicitly determines whether there is some infinitesimal endowment that results in the alternative being chosen by PF.

ALGORITHM 5: PROPORTIONALFAIR (select all alternatives)

Input: r_{t-1}

1 $C_t \leftarrow \varnothing$;

2 **for** $j = 1, \ldots, m$ **do**

3 $\quad$ **if** the following linear program is unbounded

$$\text{Maximize } L$$
$$\text{subject to } \boldsymbol{w}' \cdot \mathbf{u^t(a_j^t)} \geqslant \boldsymbol{w}' \cdot \mathbf{u^t(a_{j'}^t)} \qquad \forall j'$$
$$w_i' = \frac{1}{r_{t-1}(i)} \qquad \forall i \text{ such that } r_{t-1}(i) > 0$$
$$w_i' \geqslant L \qquad \forall i \text{ such that } r_{t-1}(i) = 0$$

$\quad$ **then**

4 $\quad\quad$ | $C_t \leftarrow C_t \cup \{a_j^t\}$;

5 $\quad$ **end**

6 **end**

7 **return** C_t;

A notable difference in the algorithms is that unlike GREEDY, PF may leave some agents with zero accrued utility even when it was possible to give positive utility to all agents.

Example 5. *Let $n = 2$, $m = 3$, and $t = 1$. Suppose that $U^1 = \left(\begin{smallmatrix} 3 & 0 & 1 \\ 0 & 3 & 1 \end{smallmatrix} \right)$. Because $t = 1$, $r_{t-1}(1) = r_{t-1}(2) = 0$.*

GREEDY *chooses* a_3^1 *since it is the only alternative that provides non-zero reward to both agents. However,* PF *assigns the agents weights* w_1, w_2 *and chooses* $\arg\max_{j \in \{1,2,3\}} \mathbf{w} \cdot \mathbf{u^t}(\mathbf{a_j^1})$. *Since it must be the case that either* $3w_1 > w_1 + w_2$ *or that* $3w_2 > w_1 + w_2$, *it is not possible for* a_3^1 *to be chosen by* PF.

For each algorithm, we prove equivalence of the two versions in the sense that the set generated by the 'select all' version consists exactly of the alternatives that the 'select one' version generates for some nondetermistic choices.

Theorem 15. *The set of alternatives* C_t *chosen by Algorithm 3 at round* t *is exactly the set of alternatives that can be chosen at round* t *by Algorithm 2.*

The proof uses the fact that the product on Line 4 of Algorithm 2 is maximized when the number of ε terms appearing in the product is minimized.

Proof. We begin by showing that every alternative that can be selected by Algorithm 2 is also selected by Algorithm 3. Let c_t be an alternative chosen by Algorithm 2 for some choices of $\{x_i\}$ and let $p = |\{i : r_{t-1}(i) + u_i^t(c_t) > 0\}|$. Therefore, the lowest power of ε with non-zero coefficient in the product on Line 4 of Algorithm 2 is ε^{n-p}. If some other alternative a_j^t has $|\{i : r_{t-1}(i) + u_i^t(a_j^t) > 0\}| > p$ then the corresponding product has non-zero coefficient on a lower power of ε, contradicting optimality of c_t. That is, $c_t \in \arg\max_{a_j^t \in A^t} |\{i : r_{t-1}(i) + u_i^t(a_j^t) > 0\}|$.

Next, let $a_{j'}^t$ be an alternative with $\{i : r_{t-1}(i) + u_i^t(c_t) > 0\} = \{i : r_{t-1}(i) + u_i^t(a_{j'}^t) > 0\}$. The product on Line 4 of Algorithm 2 is

$$NW^+(\mathbf{r_{t-1}} + \mathbf{u^t}(\mathbf{c_t}))\varepsilon^{n-p} \prod_{i:r_{t-1}(i)+u_i^t(c_t)=0} x_i$$

60

for c_t and

$$NW^+(\mathbf{r_{t-1}} + \mathbf{u^t}(\mathbf{a_{j'}^t}))\varepsilon^{n-p} \prod_{i: r_{t-1}(i)+u_i^t(a_{j'}^t)=0} x_i$$

$$=NW^+(\mathbf{r_{t-1}} + \mathbf{u^t}(\mathbf{a_{j'}^t}))\varepsilon^{n-p} \prod_{i: r_{t-1}(i)+u_i^t(a_{j'}^t)=0} x_i$$

for alternative $a_{j'}^t$. Since c_t is chosen by Algorithm 2, it must be the case that $NW^+(\mathbf{r_{t-1}}+\mathbf{u^t}(\mathbf{c_t})) \geqslant NW^+(\mathbf{r_{t-1}}+\mathbf{u^t}(\mathbf{a_{j'}^t}))$. Therefore, c_t is chosen by Algorithm 3.

To complete the proof, we show that every alternative selected by Algorithm 3 can also be selected by Algorithm 2. To that end, let $c_t \in C_t$. We exhibit a specific choice of $\{x_i\}$ which results in c_t being selected by Algorithm 2. Let K be some integer greater than the largest entry in U^t and let

$$x_i = \begin{cases} \frac{1}{2(K+1)^n}, & \text{if } r_{t-1}(i) + u_i^t(c_t) > 0 \\ 1 & \text{if } r_{t-1}(i) + u_i^t(c_t) = 0. \end{cases}$$

Then the product on Line 4 of Algorithm 2 that results from c_t being selected is

$$NW^+(\mathbf{r_{t-1}} + \mathbf{u^t}(\mathbf{c_t}))\varepsilon^{n-p},$$

where $p = |\{i : r_{t-1}(i) + u_i^t(c_t) > 0\}|$. Now consider some alternative $a_j^t \neq c_t$. If $\{i : r_{t-1}(i) + u_i^t(c_t)) > 0\} = \{i : r_{t-1}(i) + u_i^t(a_j^t) > 0\}$ and $NW^+(\mathbf{r_{t-1}} + \mathbf{u^t}(\mathbf{c_t})) \geqslant NW^+(\mathbf{r_{t-1}} + \mathbf{u^t}(\mathbf{a_j^t}))$ then the leading term in the product on Line 4 of Algorithm 2 that results from a_j^t being selected is

$$NW^+(\mathbf{r_{t-1}} + \mathbf{u^t}(\mathbf{a_j^t}))\varepsilon^{n-p} \leqslant NW^+(\mathbf{r_{t-1}} + \mathbf{u^t}(\mathbf{c_t}))\varepsilon^{n-p}.$$

Similarly, an alternative a_j^t with $|\{i : r_{t-1}(i)+u_i^t(a_j^t) > 0\}| < |\{i : r_{t-1}(i)+u_i^t(c_t) > 0\}|$ has coefficient 0 for the ε^{n-p} term (and larger terms) in the corresponding product on Line 4. In both cases, this product is greater for c_t than for a_j^t.

The final case is when $|\{i : r_{t-1}(i) + u_i^t(c_t) > 0\}| = |\{i : r_{t-1}(i) + u_i^t(a_j^t) > 0\}|$ but the two sets are not equal. In this case, the dominant term in the product on Line 4 of Algorithm 2 that results from a_j^t being selected is at most

$$NW^+(\mathbf{r_{t-1}} + \mathbf{u^t}(\mathbf{a_j^t}))\frac{1}{2(K+1)^n}\varepsilon^{n-p}$$

by the choice of $\{x_i\}$ and noting that at least one agent with $r_{t-1}(i) + u_i^t(c_t) > 0$ has $r_{t-1}(i) + u_i^t(a_j^t) = 0$. But, since the maximum reward any agent derives from any alternative is K,

$$NW^+(\mathbf{r_{t-1}} + \mathbf{u^t}(\mathbf{a_j^t})) \leqslant (K+1)^p NW^+(\mathbf{r_{t-1}} + \mathbf{u^t}(\mathbf{c_t}))$$

$$\leqslant (K+1)^n NW^+(\mathbf{r_{t-1}} + \mathbf{u^t}(\mathbf{c_t})).$$

Therefore,

$$NW^+(\mathbf{r_{t-1}} + \mathbf{u^t}(\mathbf{a_j^t}))\frac{1}{2(K+1)^n} \leqslant (K+1)^n NW^+(\mathbf{r_{t-1}} + \mathbf{u^t}(\mathbf{c_t}))\frac{1}{2(K+1)^n}$$

$$< NW^+(\mathbf{r_{t-1}} + \mathbf{u^t}(\mathbf{c_t})),$$

so the product from Line 4 of Algorithm 2 is larger for c_t than for a_j^t. Hence the particular choice of $\{x_i\}$ results in c_t being chosen by Algorithm 2. ∎

Theorem 16. *The set of alternatives C_t chosen by Algorithm 5 at round t is exactly the set of alternatives that can be chosen at round t by Algorithm 4.*

Proof. We begin by showing that every alternative that can be selected by Algorithm 4 is also selected by Algorithm 5. Let c_t be an alternative chosen by Algorithm 4 for some choices of $\{x_i\}_{i \in I_0}$ and $\{y_i\}_{i \in I_0}$. For all $i \notin I_0$, set $w_i' = \frac{1}{r_{t-1}(i)}$, and for all $i \in I_0$, set $w_i' = \frac{x_i}{\delta} + y_i$ for any $\delta > 0$. As $\delta \to 0$, the variables w_i' grow arbitrarily large. Therefore, to show feasibility of the variables $\{w_i'\}$ we need to show that the first set of constraints in the LP in Algorithm 5 hold for sufficiently small δ.

Fix an alternative a_j^t. From Line 6 of Algorithm 4, we know that $\mathbf{w} \cdot \mathbf{u^t}(\mathbf{c_t}) \geqslant \mathbf{w} \cdot \mathbf{u^t}(\mathbf{a_j^t})$. The dominant coefficient in this expression is that of ε^{-1}. Comparing these coefficients gives us

$$\sum_{i \in I_0} x_i u_i^t(c_t) \geqslant \sum_{i \in I_0} x_i u_i^t(a_j^t). \tag{3.1}$$

If Inequality 3.1 is strict, then we know that

$$\sum_{i \in I_0} \frac{x_i}{\delta} u_i^t(c_t) > \sum_{i \in I_0} \frac{x_i}{\delta} u_i^t(a_j^t)$$

for any $\delta > 0$, and we can make the gap arbitrarily large by setting δ sufficiently small. In particular, we can force the gap to be large enough that the following inequality holds for any fixed values of $\{y_i\}_{i \in I_0}$ and $\{r_{t-1}(i)\}_{i \notin I_0}$:

$$\sum_{i \in I_0} \left(\frac{x_i}{\delta} + y_i \right) u_i^t(c_t) + \sum_{i \notin I_0} \frac{1}{r_{t-1}(i)} u_i^t(c_t) > \sum_{i \in I_0} \left(\frac{x_i}{\delta} + y_i \right) u_i^t(a_j^t) + \sum_{i \notin I_0} \frac{1}{r_{t-1}(i)} u_i^t(a_j^t),$$

which is precisely the first constraint in the linear program from Algorithm 5.

If Inequality 3.1 holds with equality, then we turn attention to the coefficient of ε^0 in the dot product from Line 6 of Algorithm 4. This tells us that

$$\sum_{i \in I_0} y_i u_i^t(c_t) + \sum_{i \notin I_0} \frac{1}{r_{t-1}(i)} u_i^t(c_t) \geqslant \sum_{i \in I_0} y_i u_i^t(a_j^t) + \sum_{i \notin I_0} \frac{1}{r_{t-1}(i)} u_i^t(a_j^t). \tag{3.2}$$

Dividing Inequality 3.1 by δ and adding Inequality 3.2 gives

$$\sum_{i=1}^{n} w_i' u_i^t(c_t) \geqslant \sum_{i=1}^{n} w_i' u_i^t(a_j^t),$$

satisfying the first constraint of the LP, so the weights $\{w_i'\}$ are feasible. These weights allow us to set L to arbitrarily large values as $\delta \to 0$, so the LP is unbounded and Algorithm 5 selects c_t.

We now show the other direction, that every alternative selected by Algorithm 5 can also be selected by Algorithm 4. Let $c_t \in C_t$. That is, the optimal value for the LP in Algorithm 5 is unbounded. Then it is the case that there exist vectors $\mathbf{p}$ and $\mathbf{q} \neq \mathbf{0}$ for the values of the variables in the LP such that $\mathbf{p} + k\mathbf{q}$ is feasible for all $k > 0$ and $\mathbf{q}$ has positive objective value (this is a known fact about linear programs with unbounded value; see, e.g., [125], Theorem 4.7). We use these to exhibit values of $\{x_i\}_{i \in I_0}$ and $\{y_i\}_{i \in I_0}$ so that c_t is chosen by Algorithm 4.

Set $y_i = p_i$ and $x_i = q_i$ for all $i \in I_0$. Let $a_j^t \in A^t$. By the first set of constraints from the LP,

$$\sum_{i \in I_0} (p_i + kq_i) u_i^t(c_t) + \sum_{i \notin I_0} \frac{1}{r_{t-1}(i)} u_i^t(c_t) \geqslant \sum_{i \in I_0} (p_i + kq_i) u_i^t(a_j^t) + \sum_{i \notin I_0} \frac{1}{r_{t-1}(i)} u_i^t(a_j^t) \quad (3.3)$$

for all $k > 0$. In particular, this implies that it can not be the case that $\sum_{i \in I_0} q_i u_i^t(c_t) < \sum_{i \in I_0} q_i u_i^t(a_j^t)$, or else Inequality 3.3 would be violated for large enough values of k. There are two possiblities.

First, suppose that $\sum_{i \in I_0} q_i u_i^t(c_t) > \sum_{i \in I_0} q_i u_i^t(a_j^t)$. Then, by our choice of $x_i = q_i$ for all $i \in I_0$, we have that

$$\sum_{i \in I_0} x_i u_i^t(c_t) > \sum_{i \in I_0} x_i u_i^t(a_j^t).$$

But, as discussed earlier, $\sum_{i \in I_0} x_i u_i^t(a_j^t)$ is exactly the dominant term in Line 6 of Algorithm 4. Therefore, this dot product is maximized by c_t, so c_t is chosen by Algorithm 4.

Finally, suppose that $\sum_{i \in I_0} q_i u_i^t(c_t) = \sum_{i \in I_0} q_i u_i^t(a_j^t)$. So the dominant term in Line 6 of Algorithm 4 is equal for c_t and a_j^t. By Inequality 3.3, it must be the case that

$$\sum_{i \in I_0} p_i u_i^t(c_t) + \sum_{i \notin I_0} \frac{u_i^t(c_t)}{r_{t-1}(i)} \geqslant \sum_{i \in I_0} p_i u_i^t(a_j^t) + \sum_{i \notin I_0} \frac{u_i^t(a_j^t)}{r_{t-1}(i)}.$$

By the choice of $y_i = p_i$ for all $i \in I_0$, the above inequality holds when we substitute y_i for every instance of p_i. After making that substitution, we are left with exactly the expression for the coefficient of ε^0 in Line 6 of Algorithm 4. Since the coefficient is at least as large for c_t as for a_j^t, and the ε^{-1} coefficients are equal (and there are no further non-zero terms), c_t may be chosen by Algorithm 4. ∎

3.3.2 Axiomatization of PROPORTIONALFAIR

Now that we have given a precise definition of the PF mechanism and justified it, in this section we provide an axiomatization of the PF mechanism.

A dynamic mechanism is scale-free if it is not affected by a uniform (multiplicative) scaling of some agent's valuations. This property is desirable because it means we do not require any sort of agreement or synchronization as to the units of measurement used by the agents in their reporting.

Definition 1. *Let $k > 0$. Say that a dynamic mechanism satisfies* scale-free-ness *(SF) if C_t is unchanged (for the same choice of tiebreaking in earlier rounds) if we replace $u_i^t(a_j^t)$ by $k \cdot u_i^t(a_j^t)$ for all $a_j^t \in A^t$ for every $t \in T$.*

Lemma 17. PF *satisfies SF.*

Proof. Let $c \in C_t$ and suppose that agent i scales all her valuations by $k > 0$. We show by induction that PF still chooses c at round t. Consider a round t such that the chosen alternative is unchanged in all previous rounds.

Suppose that $r_{t-1}(i) = 0$. So for any L there exists vector of weights $\mathbf{w}'$ such that $w_i' \geqslant L$ and alternative c maximizes the weighted sum of valuations. After i scales her valuations by a factor of k, we can simply scale w_i' by a factor of $\frac{1}{k}$ (this will still allow unbounded values of w_i'). Therefore, the value $\mathbf{w}' \cdot \mathbf{u^t}(\mathbf{a_j^t})$ is unchanged for every alternative a_j^t. Thus, alternative c still maximizes this expression.

Now suppose that $r_{t-1}(i) > 0$. Then i's weight w_i' in the scaled instance is a factor of k smaller than in the un-scaled instance, but $u_i^t(a_j^t)$ is a factor of k larger than in the un-scaled instance for all alternatives a_j^t. Thus, for any setting of weights $\{w_{i'}'\}_{i' \neq i}$ in the un-scaled instance, the value $\mathbf{w'} \cdot \mathbf{u^t(a_j^t)}$ is unchanged in the scaled instance. Thus, the existence of a feasible set of weights such that c is chosen in the unscaled instance implies that c is chosen in the scaled instance also, for the same choice of weights.

Finally we need to rule out the possibility that some new alternative, $a_j^t \notin C_t$, is chosen at round t in the scaled instance. But if this were the case, then we can just scale the scaled instance by $\frac{1}{k}$ and return to the original instance where, by the above proof, $a_j^t \in C_t$. ∎

A dynamic mechanism is separable into single-minded agents if the chosen alternative at a round is unchanged by replacing an agent by several new agents with the same accrued reward, each of which has unit positive valuation for only one alternative. The axiom reflects that we can interpret utilities cardinally rather than just ordinally.

Definition 2. *Say that a dynamic mechanism is* separable into single-minded agents (SSMA) *if, when all agents have the same accrued reward $r_{t-1}(i) = r > 0$, C_t is unchanged if we replace each agent with several new agents (denoted generically by x) according to the following scheme: For every $u_i^t(a_j^t) \in U^t$, create $u_i^t(a_j^t)$ agents each with $r_{t-1}(x) = r$, $u_x^t(a_j^t) = 1$, and $u_x^t(a_{j'}^t) = 0$ for all $j' \neq j$.*

Lemma 18. PF *satisfies SSMA.*

Proof. Consider round t with valuation matrix U^t. PF chooses all alternatives that maximize the expression

$$\sum_{i=1}^{n} \frac{1}{r} u_i^t(a_j^t). \tag{3.4}$$

66

Now consider the instance expanded as defined by Definition 2. For every alternative a_j^t, there are exactly $\sum_{i=1}^n u_i^t(a_j^t)$ agents that have valuation 1 for a_j^t being chosen, while all other agents have valuation 0. Since each new agent has accrued utility r, PF chooses all alternatives which maximize Equation 3.4. ∎

The plurality axiom says that if all agent valuation vectors are unit vectors, and we have no reason to distinguish between agents, then the alternatives favored by the most agents should be chosen.

Definition 3. *Say that a dynamic mechanism satisfies* plurality (P) *if, when all agents have unit valuation for only a single alternative, and all agents have the same (non-zero) accrued reward, then C_t consists of the alternatives with non-zero valuation from the most agents.*

Plurality says nothing about the case when some agent has $r_{t-1}(i) = 0$. The idea of the axiom (in combination with SF) is that we should choose the alternative which provides the greatest utility, relative to what agents already have. However, if agents have zero accrued reward then it is not possible to make accurate comparisons as to the relative benefit each agent receives.

Observation 1. PF *satisfies plurality.*

The final axiom says that, if we restrict attention to only agents with zero accrued reward, alternatives which are dominated by a mixture of other alternatives should not be played. In the case that two alternatives are equivalent with respect to agents with $r_{t-1}(i) = 0$, we should only choose an alternative if it would still be chosen in the absence of the agents with $r_{t-1}(i) = 0$. The definition is inspired by mixed strategy dominance in game theory and, intuitively, formalizes that we should prioritize agents with zero utility above all others.

We first define the notion of 0-dominance.

Definition 4. *Let $z_1, \ldots, z_m$ be nonnegative coefficients with $\sum_{j'} z_{j'} = 1$. We say that an alternative a_j^t is* strictly 0-dominated *by the mixture of alternatives $\sum_{j'} z_{j'} a_{j'}^t$ at round t if $\sum_{j'} z_{j'} u_i^t(a_{j'}^t) \geqslant u_i^t(a_j^t)$ for all agents i with $r_{t-1}(i) = 0$, with at least one of these inequalities being strict. If all inequalities hold with equality, then we say that a_j^t is* weakly 0-dominated *by the mixture $\sum_{j'} z_{j'} a_{j'}^t$.*

We say that a_j^t is (strictly, weakly) 0-dominated if there exists some mixture of alternatives that (strictly, weakly) 0-dominates it.

Definition 5. *A dynamic mechanism f satisfies* No 0-Dominated Alternatives (NZDA) *if it never chooses a strictly 0-dominated alternative, and chooses a weakly 0-dominated alternative a_j^t only if a_j^t would be chosen by f under a scenario where U^t was modified to include (1) only the agents with $r_{t-1}(i) > 0$, and (2) only the (mixtures of) alternatives that weakly 0-dominate a_j^t (including a_j^t itself).*

Lemma 19. PF *satisfies NZDA.*

Proof. Let a_j^t be a strictly 0-dominated alternative. Note that the dominant coefficient in Line 6 of Algorithm 4 is that of ε^{-1}, which is determined by the values of $\{x_i\}_{i \in I_0}$. Therefore, an alternative is chosen by PF only if it maximizes $\sum_{i \in I_0} x_i u_i^t(a_j^t)$. So, to show that a_j^t is not selected by PF, it suffices to show that there does not exist any allowed choice of $\{x_i\}$ for which

$$\sum_{i \in I_0} x_i u_i^t(a_j^t) \geqslant \sum_{i \in I_0} x_i u_i^t(a_{j'}^t)$$

for all other alternatives $a_{j'}^t$.

Fix $\{x_i\}_{i \in I_0}$, and consider drawing an alternative $a_{j'}^t$ from the distribution defined by the weights $z_1, \ldots, z_m$. By the dominance condition and the fact that all x_i are positive,

$$\sum_{i \in I_0} x_i u_i^t(a_j^t) < \sum_{i \in I_0} x_i u_i^t(a_{j'}^t)$$

in expectation. Thus there must exist a particular j' for which the above inequality holds, so a_j^t is not chosen by PF.

Fix a choice of $\{x_i, y_i\}_{i \in I_0}$ and let a_j^t be a weakly 0-dominated alternative – suppose that it is weakly 0-dominated by alternative $a_{j'}^t$ (which may be a mixture of several alternatives). Since $u_i^t(a_{j'}^t) = u_i^t(a_j^t)$ for all agents $i \in I_0$,

$$\sum_{i \in I_0} (\frac{x_i}{\varepsilon} + y_i) u_i^t(a_j^t) = \sum_{i \in I_0} (\frac{x_i}{\varepsilon} + y_i) u_i^t(a_{j'}^t).$$

Suppose that a_j^t is chosen by PF. Then, by the definition of PF,

$$\sum_{i \in I_0} (\frac{x_i}{\varepsilon} + y_i) u_i^t(a_j^t) + \sum_{i \notin I_0} \frac{1}{r_{t-1}(i)} u_i^t(a_j^t) \geqslant \sum_{i \in I_0} (\frac{x_i}{\varepsilon} + y_i) u_i^t(a_{j'}^t) + \sum_{i \notin I_0} \frac{1}{r_{t-1}(i)} u_i^t(a_{j'}^t),$$

which requires that

$$\sum_{i \notin I_0} \frac{1}{r_{t-1}(i)} u_i^t(a_j^t) \geqslant \sum_{i \notin I_0} \frac{1}{r_{t-1}(i)} u_i^t(a_{j'}^t). \tag{3.5}$$

Equation 3.5 exactly says that PF would still choose a_j^t if only alternatives that weakly 0-dominated a_j^t were included in U^t, and in the absence of all agents with $r_{t-1}(i) = 0$, which completes the proof. $\blacksquare$

We now show that any mechanism that achieves SF, SSMA, P, and NZDA simultaneously must agree with PF. We note that of the four axioms, GREEDY satisfies only SF and P. Despite GREEDY being (arguably) simpler than PF, we do not know a good axiomatization for it.

Theorem 20. *Let f be a dynamic mechanism that satisfies SF, SSMA, P, and NZDA. Suppose that f chooses alternative c_t at round t. Then* PF *must also choose c_t at round t (for the same history up to that point).*

Proof. We have already shown that PF satisfies SF, SSMA, P, and NZDA.

It remains to show that f's choice of alternative can also be chosen by PF.

First suppose that all agents have $r_{t-1}(i) > 0$. Without loss of generality, let $r_{t-1}(i) = r$ for all agents i. We may assume this because, by SF, f and PF would choose the same alternatives at round t even if the valuation vectors of some agent(s) were multiplied by a constant across all rounds. Multiplying each agent i's valuations by $\prod_{i' \neq i} r_{t-1}(i')$, we obtain an instance in which all agents have the same accrued utility, $\prod_i r_{t-1}(i)$.

By SSMA, we can replace the agent i with $\sum_{j=1}^m u_i^t(a_j^t)$ agents, such that $u_i^t(a_j^t)$ of them have unit valuation for alternative a_j^t (and 0 valuation for all other alternatives), all with accrued reward r. Then, by plurality, f chooses $c_t \in \arg\max_{a_j^t \in A^t} \sum_{i=1}^n u_i^t(a_j^t)$. Note that PF assigns equal weight w_i to each agent since $r_{t-1}(i) = r_{t-1}(i')$ for all i, i'. Thus PF chooses precisely the alternatives which maximize $\sum_{i=1}^n u_i^t(a_j^t)$, which includes any alternative chosen by f.

The more intricate case is when there exists at least one agent with $r_{t-1}(i) = 0$. Since f satisfies NZDA, we know that f never chooses a strictly 0-dominated alternative and only chooses a weakly 0-dominated alternative if f would still choose that alternative when U^t is modified according to Definition 5. To complete the proof, we show that PF selects all alternatives that can possibly be chosen by f. Specifically, we show that PF can choose all alternatives that are not (strictly or weakly) 0-dominated, as well as any weakly 0-dominated alternative a_{j*}^t that is chosen by PF for the modified U^t. That is, when all alternatives are removed other than a_{j*}^t and (mixtures of) alternatives that weakly 0-dominate it, and all agents with $r_{t-1}(i) = 0$ are removed. This is sufficient since we have shown that PF chooses every alternative chosen by f when all agents have $r_{t-1}(i) > 0$ (which is the case when all agents with $r_{t-1}(i) = 0$ are removed).

70

An alternative a^t_{j*} is either (a) strictly 0-dominated, or (b) weakly 0-dominated and not chosen by PF when U^t is modified according to Definition 5, if and only if the optimal value of the following LP is negative for arbitrarily large values of H.

$$\text{Minimize } H \sum_{i \in I_0} \sum_{a^t_j \in A^t} (u^t_i(a^t_{j*}) - u^t_i(a^t_j))z_j + \sum_{i \notin I_0} \sum_{a^t_j \in A^t} \frac{1}{r_{t-1}(i)} (u^t_i(a^t_{j*}) - u^t_i(a^t_j))z_j$$

(3.6)

$$\text{subject to } \sum_{a^t_j \in A^t} u^t_i(a^t_j)z_j \geqslant u^t_i(a^t_{j*}) \qquad \forall i \in I_0$$

$$\sum_{a^t_j \in A^t} z_j = 1$$

$$z_j \geqslant 0 \qquad \forall j$$

If a^t_{j*} is strictly dominated then the first term of the objective can be made negative (and therefore the whole objective can be made negative when H is large enough). If a^t_{j*} is only weakly dominated, then the first term can be set to 0, and the second term to be negative when there exists a mixture of alternatives that is chosen by PF (ahead of a^t_{j*}) according to the modified U^t. Conversely, if the optimal value of the objective is negative then either there exist values for $\{z_j\}$ such that the first term is negative (which, combined with the first set of constraints, says that a^t_{j*} is strictly 0-dominated), or there exist values for $\{z_j\}$ such that the first term is zero and the second term is negative. If the second term is negative then the weighted sum of valuations for the mixed alternative defined by $\{z_j\}$ is higher than the weighted sum of valuations for a^t_{j*}, for the weights defined by PF when restricted to agents $i \notin I_0$. This proves correctness of the LP.

We want to show that PF can choose any alternative for which the the optimal value of LP (3.6) is nonnegative. Let a^t_{j*} be such an alternative. We show that a^t_{j*} can be chosen by PF by considering the dual, which has variables w_i for all $i \in I_0$ (one for each constraint) and s (corresponding to the constraint on the sum of the

71

z_j):

$$\text{Maximize} \quad \sum_{i \in I_0} u_i^t(a_{j*}^t) w_i - s$$

$$\text{subject to} \quad \sum_{i \in I_0} u_i^t(a_j^t) w_i - s \leqslant H \sum_{i \in I_0} (u_i^t(a_{j*}^t) - u_i^t(a_j^t))$$

$$+ \sum_{i \notin I_0} \frac{u_i^t(a_{j*}^t) - u_i^t(a_j^t)}{r_{t-1}(i)} \quad \forall j \in \{1, \ldots, m\}$$

$$w_i \geqslant 0 \quad \forall i \in I_0$$

Let $\overline{s} = \sum_{i \in I_0} u_i^t(a_{j*}) w_i - s$ denote the objective. The first set of constraints can now be rewritten as

$$\overline{s} + \sum_{i \in I_0} (w_i + H) u_i(a_j^t) + \sum_{i \notin I_0} \frac{1}{r_{t-1}(i)} u_i^t(a_j^t) \leqslant \sum_{i \in I_0} (w_i + H) u_i^t(a_{j*}^t) + \sum_{i \notin I_0} \frac{1}{r_{t-1}(i)} u_i^t(a_{j*}^t).$$

Since a_{j*}^t is not 0-dominated, the optimal value of LP (3.6) is at least zero for any arbitrarily large value of H. By strong duality, the optimal value of the dual is therefore also at least zero for arbitrarily large values of H. Thus, if we set $w_i' = w_i + H$ for all $i \in I_0$ and $w_i' = \frac{1}{r_{t-1}(i)}$ for all $i \in I_0$, we have an unbounded and feasible set of weights for the linear program to choose a_{j*}^t in the definition of Algorithm 5. Therefore, a_{j*}^t can be chosen by PF. ∎

3.4 Simulations

We ran the algorithms on data gathered from a power boost allocation problem. In this problem, n computer applications are each allocated a base level of power, and compete for $k < n$ additional (indivisible) units of extra power (*power boosts*) at each of m rounds (each application gets at most one boost per round). We obtain our instance from Apache Spark [168] benchmarks.

Table 3.1 lists the twelve Spark applications in our instance, each of which is defined by a fixed number of tasks. We profile tasks' completion time. We take the

Table 3.1: Spark applications and categories.

Category	Applications
Statistics	`Correlation`
Classification	`DecisionTree, GradientBoostedTrees, SVM, LinearRegression, NaiveBayesian`
Pattern Mining	`FP Growth`
Clustering	`KMeans`
Collaborative Filtering	`ALS`
Graph Processing	`Pagerank, ConnectedComponents, TriangleCounting`

number of tasks completed in a round by an application as that application's utility. Thus, for each application x, we estimate the base and boosted power utility ($u_{x,t}^{\text{base}}$ and $u_{x,t}^{\text{boost}}$) in each round.

In our instance, there are two power boosts to be allocated. So at each round there are $\binom{12}{2}$ alternatives, one for each pair of applications. For an alternative a_j^t corresponding to power boosts for applications x and y, we have that $u_x^t(a_j^t) = u_{x,t}^{\text{boost}}$, $u_y^t(a_j^t) = u_{y,t}^{\text{boost}}$, and $u_z^t(a_j^t) = u_{z,t}^{\text{base}}$ for all other applications $z \neq x, y$. We have 497 rounds in the instance we tested.

We evaluate the performance of GREEDY and PF against the optimal offline solution, and also against an algorithm designed for online stochastic convex programming[3] [15] - a class of problems that includes the one we study. To our knowledge this algorithm is the state of the art for such problems in terms of theoretical guarantees. We refer to this algorithm as STOCHASTIC. The algorithm works by maximizing a weighted sum of valuations at each round, where the weights are updated at each round using online convex optimization techniques. The theoretical guarantees for STOCHASTIC are in expectation over instances where the order of the input matrices is *randomly permuted*. In the instance we test, however, the utilities are highly correlated over time. Applications that would benefit from a power boost in some round t are more likely to also benefit from a power boost in round $t + 1$, because application phases may span multiple rounds. Due to this and other technical rea-

---

[3] Of course, there are other online scheduling algorithms but they do not pursue Nash welfare as an objective.
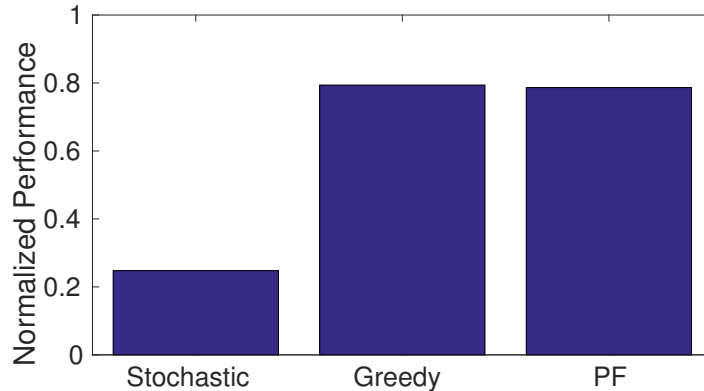
FIGURE 3.2: Nash welfare achieved by the algorithms, normalized against OPT (which has performance 1).

sons, the theoretical guarantees do not apply here. The performance of the three algorithms is shown in Figure 3.2.

We see that STOCHASTIC performs relatively poorly, while GREEDY and PF each achieve about 80% of the performance of OPT. This motivates us to examine the difference in performance between GREEDY and PF for smaller values of m, as the difference between these two algorithms is most pronounced while a single decision has a relatively large effect.

To generate smaller instances, we sample starting rounds from the full set of 497 rounds. For each value of m in Figure 3.3, we randomly generate a starting round $t \in [1, 497 - m]$ and consider the m rounds starting at t, for 100 random choices of t. Our measure of performance is $NW(\mathbf{u}_{\mathbf{m}}^{\mathrm{avg}})$, allowing for fair comparisons between different values of m.

We note that PF consistently performs slightly worse than GREEDY, which is consistent with the performance on the full instance. The difference is most pronounced on small values of m, since this is where the two algorithms differ the most. Performance increases with m, as we would expect, since more rounds allow the algorithms to choose more flexibly once all applications have positive accrued re-
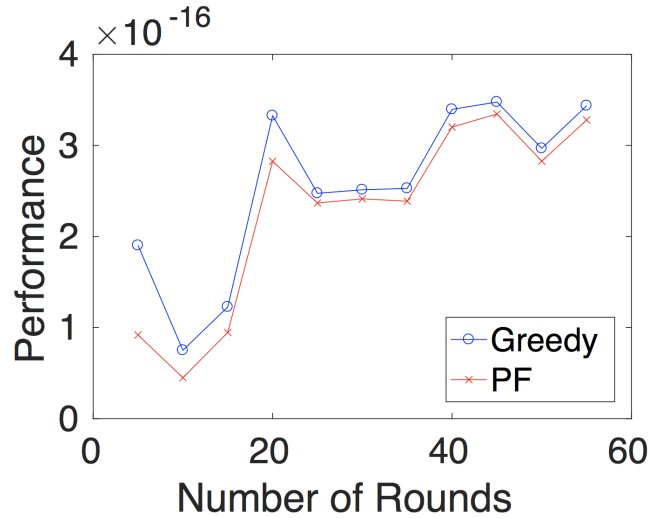
FIGURE 3.3: Nash welfare achieved by GREEDY and PF as a function of the number of rounds.

ward. However, the increase is not monotonic. One explanation for this is because we throw away any choice of starting round t for which it is impossible to achieve $NW(\mathbf{u_m^{avg}}) > 0$ (it might be the case that for all m rounds, some application cannot receive positive utility). Since smaller values of m result in more choices of t being disqualified, there is a sense in which we are selecting for 'easier' instances for smaller values of m.

3.5 Discussion

Election designers and social choice researchers often do not consider the fact that many elections do not occur in isolation, but rather are repeated over time. In this work, we have provided a framework to allow for the design and analysis of dynamic election protocols, and repeated decision making rules generally. We have presented two candidate online algorithms for solving these dynamic problems. Our simulations show that both algorithms perform well, but do not determine that either is clearly a better choice than the other. While GREEDY achieves slightly higher performance,

PROPORTIONALFAIR has the advantage of being justified by the axiomatization given in this paper.

Note that in neither this chapter nor Chapter 2 have we focused on strategic concerns. This is an important direction that we consider in Chapter 4, and discuss in Chapter 8. However, there are also important contexts where strategic concerns do not come into play. For example, instead of considering a setting where there are multiple *agents* that have different utility functions, we can consider a setting where there are multiple *objectives* that each alternative contributes towards. For example, consider faculty hiring. Suppose the three objectives that we want our faculty hires to contribute to are research, teaching, and service; moreover, suppose that at the time of hiring we can predict well how much each candidate would contribute to each of these objectives, if hired. Then, it stands to reason that, one year, we may hire a top researcher that we do not expect to contribute much to the other objectives. But we would be loath to make such a decision *every* year; having hired a few top researchers who are not good at teaching or service, pressure will mount to address these needs. This fits well into our framework, if we simply treat each of the three objectives as an agent that is "happy" with an alternative to the extent to which it addresses the corresponding objective. In particular, note that the fact that objectives are measured in incomparable units – for example, we might measure research crudely by number of top-tier publications, and teaching crudely by course evalation scores – poses no problem to our methodology, since this methodology can anyway address agents measuring their utilities in different units.

4

Private Goods, Online Allocation

4.1 Introduction

So far, we have not considered agents that act strategically for private benefit. In our general public decisions model, accounting for these concerns is a difficult task due to the well known *free-rider problem*: agents with high utility for some popular alternative can report low utility for that alternative and 'free-ride' off its popularity.

In this chapter, we address strategic behavior by considering a less general, private good setting. Our motivation comes primarily from sharing computational resources, but the model applies to any setting where a community of agents contribute resources to a pool that is shared over time. Examples include supercomputers for scientific computing [115], datacenters for Internet services [36, 157], and clusters for academic research [12, 70].

In this chapter, we consider a model in which each user owns and operates some number of identical resources (their *endowment*). At every time period (*round*), they have some resource requirement, which may or may not be met by their endowment alone. If agents act individually, some will waste unused resources at a given round,

while others will fail to meet their full demand.

We seek to design centralized allocation mechanisms to rectify this problem. Of course, if agents were to willingly put their resources under the control of a benevolent and omniscient dictator, that dictator could allocate resources to the agents with highest utility for them. Unfortunately, agents are strategic and their true utilities are private information that must be extracted by the mechanism. Strategic agents act selfishly to pursue their own objectives. Agents will determine whether misreporting demands can improve their performance even at the expense of others in the system. For example, an agent is likely to over-report her demand in the current time period to obtain more resources, unless doing so leads to a reduction in the resources allocated to her in later periods.

We seek allocation mechanisms that satisfy strategy-proofness (SP), which ensures that no agent benefits by misreporting her demand for resources. Incentive compatiblity is a key feature contributing to efficiency as it allows the mechanism to optimize performance according to agents' true utilities. Without SP, agents' reports may not represent their true utility and allocating based on reported demands may not produce any meaningful performance guarantee. Moreover, strategy-proof mechanisms reduce the cognitive load on agents by eliminating the need to optimally construct resource demands or preemptively respond to misreports by other agents in the system.

Strategy-proofness is complemented by sharing incentives (SI), which ensures that agents perform at least as well as they would have by not participating in the allocation mechanism (i.e., using their own resources as a smaller, private system). With sharing incentives, agents willingly share their resources and manage them according to the commonly agreed upon policy. Without sharing incentives, we need to be able to either force agents to cooperate, or be content with the possibility that rational agents will choose not to participate in the mechanism.

In this chapter, we consider agents who derive high utility per unit of resource up until some amount of resource allocation (i.e., their demand) and derive low utility beyond that allocation. The high-low formulation is appropriate for varied resources such as processor cores, cache and memory capacity, or virtual machines in a datacenter. For example, an agent could derive high utility when additional processors permit her to dequeue more tasks from a highly critical job. Once the job's queue is empty, she derives low utility from using additional processors to replicate tasks, which guards against stragglers or failures. In another example, an agent that is allocated more power can turn on more processors, each of which provides high utility from task parallelism. Once the agent exhausts her job's parallelism, it can use additional power to boost processor voltage and frequency for lower, non-zero utility.

We propose allocation mechanisms for dynamic proportional sharing to address limitations in existing approaches. We begin by proving that policies used in state-of-the-art schedulers [11, 10, 13] fail to satisfy SP or SI. We then propose two alternative mechanisms. First, as our main contribution, we propose the flexible lending mechanism to satisfy SP, guarantee at least 50% of SI performance, and provide an asymptotic efficiency guarantee. The mechanism uses tokens to enable these theoretical guarantees. In practice, our simulations show that performance is comparable to that of state-of-the-art mechanisms and achieves 98% of SI performance, much better than the lower bound. Second, for situations where SI is a hard constraint, we propose the T-Period mechanism to satisfy SP and SI while still outperforming static allocations.

4.2 Preliminaries

Although the setting considered in this chapter is captured by the general public decisions model of Chapters 2 and 3, it is more conveniently represented by a more
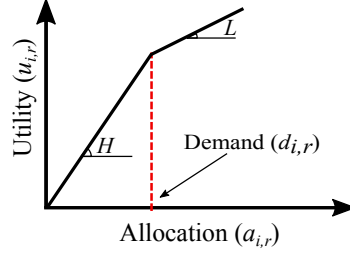
FIGURE 4.1: A user derives high utility from resources up to her demand and derives low utility from resources beyond her demand.

compact model. Therefore, we introduce a new set of notation for this chapter.

Consider a dynamic system with n agents and R discrete rounds. Agent i contributes $e_i > 0$ units of a resource at each round, which we refer to as her *endowment*. In other words, e_i is agent i's contribution to the federated system, which does not vary over time. Let $[n] = \{1, \ldots, n\}$ and $E = \sum_{i \in [n]} e_i$ denote the total number of units to be allocated at each round. At round r, agent i has a true demand of $d_{i,r} \geqslant 0$ units and reports a demand of $d'_{i,r} \geqslant 0$. Let $\mathbf{d}'_i = (d'_{i,1}, \ldots, d'_{i,R})$ denote the vector of agent i's reports, and $\mathbf{d}'_{-i}$ denote the reports of all agents other than i.

A dynamic allocation mechanism M assigns each agent an allocation $a_{i,r}^M(\mathbf{d}'_i, \mathbf{d}'_{-i})$ using only information from the first r entries in the demand vectors. We will often write simply $a_{i,r}$ when the exact mechanism and the demands are clear from context. Let $\mathbf{a_i^M}(\mathbf{d}'_i, \mathbf{d}'_{-i})$, often simply $\mathbf{a_i}$, denote the vector of agent i's allocations. Agents have high (H) utility per resource up to their demand, and low (L) utility per resource that exceeds their demand. Formally, the utility of agent i at round r for $a_{i,r}$ units is denoted by $u_{i,r}(a_{i,r})$ and modeled as the following.

$$
u_{i,r}(a_{i,r}) = \begin{cases} a_{i,r} H & \text{if } a_{i,r} \leqslant d_{i,r}, \\ d_{i,r} H + (a_{i,r} - d_{i,r}) L & \text{if } a_{i,r} > d_{i,r}. \end{cases}
$$

Figure 4.1 shows $u_{i,r}$ for user i with demand $d_{i,r}$ at round r. For simplicity, we assume H and L are the same for all agents, but all our results extend to the case

80

where agents have different values of H and L (with the exception of §4.5.5).

While resources and demands are discrete, we allow the allocations $a_{i,r}$ to be real-valued. Real-valued allocations can be thought of as probabilistic—the realized allocation is a random allocation where agent i is allocated $a_{i,r}$ resources in expectation, which is always possible as a result of the Birkhoff-von Neumann theorem [33]. Agent i's overall utility after R rounds for allocation $\mathbf{a_i}$ is calculated additively as follows.

$$U_{i,R}(\mathbf{a_i}) = \sum_{r=1}^{R} u_{i,r}(a_{i,r}).$$

We do not consider discounting for simplicity of presentation, but our mechanisms readily extend to the case where agents discount their utilities over time.

In this paper, we focus on three main properties: strategy-proofness, sharing incentives, and efficiency. First, strategy-proofness says that agents never benefit from lying about their demands. In other words, agent i's utility decreases if she reports $\mathbf{d_i'} \neq \mathbf{d_i}$.

Definition 6. *Mechanism M satisfies* strategy-proofness (SP) *if*

$$U_{i,R}(\mathbf{a_i^M}(\mathbf{d_i}, \mathbf{d_{-i}'})) \geqslant U_{i,R}(\mathbf{a_i^M}(\mathbf{d_i'}, \mathbf{d_{-i}'})) \qquad \forall i, \forall R, \forall \mathbf{d_i}, \forall \mathbf{d_i'}, \text{ and } \forall \mathbf{d_{-i}'}.$$

Next, sharing incentives says that by participating in the mechanism, agents receive at least the utility they would have received by not participating.

Definition 7. *Mechanism M satisfies* sharing incentives (SI) *if*

$$U_{i,R}(\mathbf{a_i^M}(\mathbf{d_i}, \mathbf{d_{-i}'})) \geqslant U_{i,R}(\mathbf{e_i}) \qquad \forall i, \forall R, \forall \mathbf{d_i}, \text{ and } \forall \mathbf{d_{-i}'}.$$

We also define a relaxed notion of α-sharing incentives, which says that every agent gets at least an α fraction of the utility that she would have received without taking part in the mechanism. Note that 1-SI is equivalent to SI.

Definition 8. *Mechanism M satisfies α-SI if*

$$U_{i,R}(\mathbf{a_i^M}(\mathbf{d_i}, \mathbf{d'_{-i}})) \geqslant \alpha \ U_{i,R}(\mathbf{e_i}) \qquad \forall i, \forall R, \forall \mathbf{d_i}, \ and \ \forall \mathbf{d'_{-i}}.$$

Finally, efficiency says that all resources should be allocated, and an agent with L valuation should never receive a resource while there are agents with H valuation for that resource.

Definition 9. *Mechanism M satisfies* efficiency *if*

$$\sum_{i \in [n]} a_{i,r}^M = E,$$

and if $a_{i,r}^M > d'_{i,r}$ for some agent i and round r, then $a_{j,r}^M \geqslant d'_{j,r}$ for all agents.

Note that efficiency is relative to the agents' *reports*, not their actual valuations, which are hidden from the mechanism. Therefore, in situations where agents lie about their valuations, it is possible that even an efficient mechanism allocates a unit inefficiently with respect to the actual valuations. With this in mind, there is little value in a mechanism that is efficient but not SP. Similarly, if a mechanism does not satisfy SI, then agents may not want to participate in it. So an efficient mechanism that does not satisfy SI may not actually exhibit efficiency gains in practice because agents choose not to participate. In some contexts, SI may not be of concern because agents are forced to participate or are willing to risk participation if gains are likely large and losses are likely small.

For readability, some proofs are omitted and appear in the appendix.

4.3   Existing Mechanisms

In this section, we focus on the (weighted) max-min fairness policy, which is one of the most widely used policies in computing systems. It is deployed in many state-of-the-art datacenter schedulers such as the Hadoop Fair Scheduler [11], Hadoop Capacity

Scheduler [10] and Spark Dynamic Allocator [13]. And it has been extensively studied in the literature [89, 91, 149].

A dynamic allocation mechanism could deploy the max-min policy for two different objectives: maximizing the minimum accumulated allocations up to a round, or maximizing the minimum allocation at each round, independently of previous rounds. We call the first mechanism *Dynamic Max-Min (DMM)* and the second mechanism *Static Max-Min (SMM)*. First, at each round r, DMM selects the allocation that maximizes $\min_i \sum_{r'=1}^{r} a_{i,r'}/e_i$, the minimum weighted cumulative allocation; subject to this, it maximizes the second lowest weighted cumulative allocation, and so on. This maximization is subject to the constraint that no resource is allocated to an agent with low valuation as long as there are agents with high valuation.

Second, at each round r, SMM selects the allocation that maximizes $\min_i a_{i,r}/e_i$, the minimum weighted allocation at that round; subject to this, it maximizes the second lowest weighted allocation, and so on. This maximization is also subject to the constraint that no resource is allocated to an agent with low valuation as long as there are agents with high valuation. Under SMM, agents are guaranteed to receive their demands as long as they are less than or equal to their endowment. Agents with demands higher than their endowments receive extra resources from agents with demands lower than their endowments. Unlike DMM, SMM allocates resources locally at round r, regardless of agents' allocations prior to round r.

In the rest of this section, we study properties of these two mechanisms. In particular, we focus on three properties: strategy-proofness, sharing incentives, and efficiency. We examine whether the existing mechanisms satisfy these properties for the special case when $L = 0$ and for the general case when $L > 0$.

4.3.1 Properties of Mechanisms for $\mathbf{L} = \mathbf{0}$

When $L = 0$, one might think that agents do not have any incentive to misreport their demands. However, we show that DMM fails to satisfy SI and SP.

Theorem 21. *Dynamic max-min mechanism violates sharing incentives, even when $L = 0$.*

Proof. Suppose that $R = 10$ and there are three agents, each with $e_i = 3$. For all rounds $r \neq 10$, the demands are $d_{1,r} = 1$, $d_{2,r} = 2$, and $d_{3,r} = 6$. For rounds $r = 1, \ldots, 9$, each agent is allocated exactly her demand. After round 9, utilities for agents 1, 2 and 3 are $9H$, $18H$ and $54H$, respectively. At round 10, demands are $d_{1,10} = 9$, $d_{2,10} = 9$, and $d_{3,10} = 6$. DMM allocates all 9 units to agent 1, which maximizes the minimum weighted cumulative allocation. Consider agent 2. Under DMM, agent 2's allocation is $a_{2,r} = 2$ for all $r \neq 10$ and $a_{2,10} = 0$. If she had not participated in the mechanism, then she would have obtained the same utility in each round $r \neq 10$, but a strictly higher utility in round $r = 10$. ∎

Theorem 22. *Dynamic max-min mechanism violates strategy-proofness, even when $L = 0$ [17].*

Proof. Consider three agents with equal endowments $m_1 = m_2 = m_3 = 1$ sharing three units of a resource for three rounds. The demand of agent 1 is 3 for all three rounds. Agent 2's demand is 3 for rounds 1 and 3 and 0 for round 2. And agent 3 has a demand of 3 for round 2 and 0 for rounds 1 and 3. Agent 1 achieves utility of $3.375H$ by truthful reporting. If agent 1 misreports 0 for round 1, her utility would increase to $3.75H$. ∎

Since DMM does not satisfy SP, it cannot guarantee any meaningful notion of efficiency, as explained in §4.2. Next, we show that SMM satisfies SI, SP, and efficiency.

Theorem 23. *Static max-min mechanism satisfies strategy-proofness, sharing incentives, and efficiency when $L = 0$.*

Proof. We start by proving that SMM satisfies SP. Under SMM, allocations at round r are independent of allocations at previous rounds. Suppose that agent i reports $d'_{i,r} \neq d_{i,r}$ at round r. Let $a'_{i,r}$ and $a_{i,r}$ denote i's allocations at round r for reporting $d'_{i,r}$ and $d_{i,r}$, respectively. If $a_{i,r} \geqslant d_{i,r}$, then i already receives her highest possible utility, $d_{i,r}H$ (because $L = 0$), and she cannot benefit from misreporting.

If $a_{i,r} < d_{i,r}$, then for all $j \neq i$, we have: (1) $a_{j,r} \leqslant d_{j,r}$ and (2) $a_{i,r}/e_i \geqslant a_{j,r}/e_j$. The former holds by SMM's definition. The latter holds because SMM maximizes the minimum weighted allocations in a lexicographical order. If there is j with $a_{j,r}/e_j > a_{i,r}/e_i$, then SMM should decrease $a_{j,r}$ and increase $a_{i,r}$. Now, suppose for contradiction that $a'_{i,r} > a_{i,r}$. Since $\sum_k a'_{k,r} = \sum_k a_{k,r}$, there should be an agent ℓ with $a'_{\ell,r} < a_{\ell,r} \leqslant d_{\ell,r}$. Therefore, we have:

$$a'_{\ell,r}/e_\ell < a_{\ell,r}/e_\ell \leqslant a_{i,r}/e_i < a'_{i,r}/e_i.$$

This is a contradiction because SMM could improve its objective value by decreasing $a'_{i,r}$ and increasing $a'_{\ell,r}$.

To see that SMM satisfies SI, note that an agent can guarantee herself at least e_i resources (her utility from not participating) at each round by reporting $d'_{i,r} = e_i$ for all r. By SP, truthful reporting achieves at least this utility. Therefore, truthful reporting achieves at least as much utility as not participating in SMM, which proves SI. Finally, SMM satisfies efficiency by definition, since it either completely fulfills all demands or allocates all resources to agents that value them highly. ∎

4.3.2 Properties of Mechanisms for **L > 0**

We now consider the general setting where an agent's low valuation is still positive. Unfortunately, SMM no longer retains its properties from the $L = 0$ case. Agents

are no longer indifferent to forsaking low-valued resources and may lie in order to receive them.

Theorem 24. *When $L > 0$, static max-min mechanism violates strategy-proofness and sharing incentives.*

Proof. Consider an instance with 2 agents, each with endowment $e_i = 1$, and a single round. Agent 1 has demand 2 and agent 2 has demand 0. SMM allocates both resources to agent 1 and nothing to agent 2. However, had agent 2 not participated in the mechanism, she would have received one resource and utility $L > 0$. Similarly, had she misreported her demand to be 1, she would have received one resource and utility $L > 0$. ∎

Indeed, in this general setting, no mechanism can simultaneously satisfy efficiency and either of the two other desired properties.

Theorem 25. *When $L > 0$, there is no dynamic mechanism that satisfies α-sharing incentives and efficiency, for any $\alpha > 0$.*

Proof. Consider an instance with two agents, each with endowment $e_i = 1$, and a single round. Agent 1 has demand 2 and agent 2 has demand 0. Efficiency dictates that we allocate both resources to agent 1, which would violate α-SI for agent 2 for any $\alpha > 0$. ∎

Theorem 26. *When $L > 0$, there is no dynamic mechanism that satisfies strategy-proofness and efficiency.*

Proof. Consider an instance with two agents, each with endowment $e_i = 1$, and a single round. Both agents have demand 0. For efficiency, the mechanism must allocate all the resources so that at least one agent receives $a_{i,1} > 0$. Supposing without loss of generality that $a_{1,1} > 0$, then $a_{2,1} < 2$. If agent 2 misreports $d'_{2,1} = 2$,

by efficiency, the mechanism must allocate both resources to agent 2, which is an improvement over her utility from reporting truthfully. ∎

Note that SP and SI are compatible. A mechanism that statically allocates agents their endowments satisfies SP and SI; agents have no incentive to misreport because allocations do not depend on reports and agents receive their fair share of resources. This mechanism clearly fails to satisfy efficiency and does not extract any benefit from sharing. In §4.5, we propose a mechanism that satisfies strategy-proofness, guarantees each user at least 50% of their utilities from sharing incentives, and provides an asymptotic efficiency guarantee.

4.4 Proportional Sharing With Constraints Procedure

The mechanisms we present in the remainder of this paper have, at their core, a procedure we call *Proportional Sharing With Constraints (PSWC)*. The procedure allocates some amount of resources among agents proportionally to their (exogenous) weights subject to (agent-dependent) minimum and limit constraints: (1) each agent receives at least her minimum allocation, and (2) each agent should receives no more than her limit allocation.

Formally, PSWC takes as input an amount to allocate A, weights $\mathbf{w} = (w_1, \ldots, w_n)$, minimum allocations $\mathbf{m} = (m_1, \ldots, m_n)$, and limit allocations $\mathbf{l} = (l_1, \ldots, l_n)$. PSWC outputs a vector of allocations $\mathbf{a} = (a_1, \ldots, a_n)$ defined as the solution to the following program.
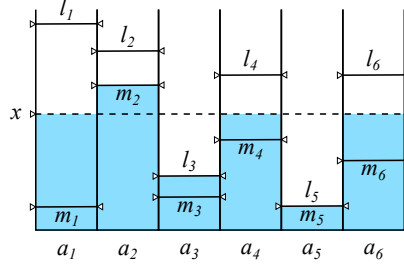
FIGURE 4.2: There are six agents with equal weights. The allocations are represented by the height of the corresponding blue vertical bar. The allocations can be thought of as the 'most equal' allocations, subject to no agent receiving less than her minimum constraint or more than her maximum.

Minimize x,

$$\text{s.t. } a_i/w_i \leqslant x \qquad\qquad \text{if } m_i < a_i \leqslant l_i,$$

$$a_i \leqslant l_i \qquad\qquad \forall i,$$

$$a_i \geqslant m_i \qquad\qquad \forall i,$$

$$\sum_{i\in[n]} a_i = A.$$

PSWC is illustrated in Figure 4.2. The program can be solved in $O(n\log(n))$ time by the Divvy algorithm [98]. The Divvy algorithm proceeds by sorting the limit and minimum allocation bounds in $O(n\log(n))$ time, and then conducting a linear time search for the optimal value of x by increasing the allocations in discrete steps until all resources have been allocated.

The following lemma characterizes the allocations produced by the PSWC procedure and will be useful in our later proofs.

Lemma 27. *Under PSWC, for every agent i, $a_i = \max(m_i, \min(l_i, xw_i))$.*

Proof. First, we show that if $m_i < xw_i$, then $a_i = \min(l_i, xw_i)$. If $a_i > \min(l_i, xw_i)$, then at least one constraint is violated. If $a_i < \min(l_i, xw_i)$, then there exists at least one agent ℓ such that $a_\ell = xw_\ell$ because otherwise, x is not optimal. In this case, a_i

88

can be increased while a_ℓ for all ℓ with $a_\ell = xw_\ell$ decreases. This allows for a smaller value of x, which contradicts the optimality of x. Next, we show that if $m_i \geqslant xw_i$, then $a_i = m_i$. Since a_i cannot be less than m_i, if a_i is not equal to m_i, then $a_i > m_i$, which means $a_i > xw_i$. However, since $a_i > m_i$, the first constraint dictates that $a_i \leqslant xw_i$, a contradiction. Combining these two cases gives the desired result. $\blacksquare$

Our proposed mechanisms all have similar structure. First, agents always receive exactly the same number of resources that they contribute to the system (over the entire R rounds). This is a fairness primitive in its own right, but is primarily a design feature that helps us provide desirable properties. Second, all our proposed dynamic mechanisms call the PSWC procedure to allocate resources at each round. Our mechanisms are determined primarily by how we set the minimum and maximum constraints.

4.5   Flexible Lending Mechanism

We now turn to designing mechanisms that satisfy our game-theoretic desiderata while increasing efficiency significantly over static allocation. The static allocation mechanism satisfies both SP and SI, but it does not exhibit any gains from sharing. DMM and SMM sacrifice SP and SI in exchange for efficiency. However, in the absence of SP, any guarantee on efficiency based on agents' demands is not meaningful as agents have incentives to misreport their demands when $L > 0$. In this section, we present the flexible lending (FL) mechanism. The flexible lending mechanism achieves strategy-proofness and an asymptotic efficiency guarantee. FL satisfies a theoretical 0.5 approximation to SI and our simulation results show that it significantly outperforms this bound in practice (see §4.7).

4.5.1 Definition

For a fixed number of rounds R, FL allocates exactly Re_i resources to each agent i, which is exactly her contribution to the shared pool over all R rounds. The mechanism enforces this constraint by simply removing agent i from the list of eligible agents once she receives Re_i resources in total. We keep track of the resources each agent has received with a running token count t_i, effectively 'charging' each agent a token for every resource she receives. We denote by $t_{i,r}$ the number of tokens that agent i holds at the start of round r. Thus, the number of tokens that an agent holds puts a hard limit on the number of resources she can receive at any given round.

Algorithm 6 presents the flexible lending mechanism. We define $\bar{d}_i$ to be the allocatable demand of agent i at each round, which is simply the minimum of her reported demand $d'_{i,r}$ and the number of tokens she has remaining t_i. We distinguish between two cases depending on whether the total allocatable demand is higher or lower than the total supply of resources.

First, if the total allocatable demand is at least as high as the total supply, then FL runs PSWC with the minimum allocation for each agent set to 0, and the limit allocation set to $\bar{d}_i$. This way, resources are allocated proportionally among all agents that want them. Second, if the total allocatable demand is less than the total supply, then agents receive their full allocatable demand. Therefore, FL runs PSWC with minimum allocation for each agent i set to $\bar{d}_i$, and limit allocations set to her number of tokens t_i (which is always at least as large as her allocatable demand). This way, FL allocates resources proportionally among all agents, subject to the condition that no agent receives fewer resources than her demand.

We illustrate FL with an example.

Example 6. *Consider a system with three agents and four rounds. Each agent has endowment $e_i = 1$. Suppose that agents' (truthful) reports are given by the following*

ALGORITHM 6: Flexible Lending Mechanism

```
1  t = Re;                                          /* Initialize token count */
2  for r ∈ {1, ..., R} do
3  |    d̄ ← min(d'_{·,r}, t);                       /* d̄_i is i's allocatable demand */
4  |    D ← Σ_{i∈[n]} d̄_i;
5  |    if D ⩾ E then
6  |    |    a_{·,r} ← PSWC(A = E, l = d̄, m = 0, w = e);
7  |    else
8  |    |    a_{·,r} ← PSWC(A = E, l = t, m = d̄, w = e);
9  |    end
10 |    t ← t − a_{·,r};
11 end
```

table:

	$d_{i,1}$	$d_{i,2}$	$d_{i,3}$	$d_{i,4}$
$i = 1$	3	1	1	0
$i = 2$	0	2	1	2
$i = 3$	0	0	0	4

FL allocations are given by the following table:

	$a_{i,1}^{FL}$	$a_{i,2}^{FL}$	$a_{i,3}^{FL}$	$a_{i,4}^{FL}$
$i = 1$	3	1	0	0
$i = 2$	0	2	1.5	0.5
$i = 3$	0	0	1.5	2.5

While all agents have tokens remaining, FL efficiently allocates resources. However, in round 3, agent 1 has no tokens remaining and therefore the supply of resources exceeds the allocatable demand. In this case, resources are evenly divided between agents 2 and 3. In the final round, agent 2 can receive only 0.5 resources before running out of tokens, so the rest of the resources are allocated to agent 3.

4.5.2   Basic Properties

Next, we study the properties of FL. We first show that FL satisfies strategy-proofness. We then show that FL guarantees at least 50% of SI performance. And finally we show that FL provides an asymptotic efficiency guarantee. Throughout

this section, we extensively use the following lemma which characterizes FL allocations.

Lemma 28. *Let x denote the objective value of FL's call to PSWC at round r. If $D \geqslant E$, then $a_{i,r} = \min(xe_i, d_{i,r}, t_{i,r})$. If $D < E$, then $a_{i,r} = \min(t_{i,r}, \max(d_{i,r}, xe_i))$.*

Proof. Suppose first that $D \geqslant E$. Substituting the relevant terms into Lemma 27, we have

$$a_{i,r} = \max(0, \min(\min(d_{i,r}, t_{i,r}), xe_i)) = \min(xe_i, d_{i,r}, t_{i,r}).$$

If instead $D < E$, then again substituting into Lemma 27 gives

$$a_{i,r} = \max(\min(d_{i,r}, t_{i,r}), \min(t_{i,r}, xe_i)) = \min(t_{i,r}, \max(d_{i,r}, xe_i)).$$

The final equality, $\max(\min(A, B), \min(A, C)) = \min(A, \max(B, C))$ can easily be checked to hold case by case for any relative ordering of A, B, and C. ∎

We next prove a basic monotonicity result, which states that if we shift some tokens to a single agent from all other agents, then the agent with more tokens achieves a (weakly) higher allocation. The proof follows easily from Lemma 28 and is deferred to the Appendix.

Lemma 29. *Consider some agent i, and suppose that $t'_{i,r} \geqslant t_{i,r}$, $t'_{j,r} \leqslant t_{j,r}$ for all $j \neq i$, and $d'_{k,r} = d_{k,r}$ for all $k \in [n]$. Then $a'_{i,r} \geqslant a_{i,r}$.*

As our main technical result, we show in the following subsection that FL is strategy-proof. At a high level, we show that if an agent receives fewer high-valued resources as a result of misreporting, then her allocations in all future rounds are weakly higher. This means that she cannot receive fewer low-valued resources at any future round, relative to her allocations had she not misreported. Therefore, because the total number of resources allocated to each agent is fixed (by the initial token count), her misreport can only result in trading high-valued resources at an early round for other, potentially low-valued, resources at later rounds.

Suppose agent i reports demands that are not equal to her true demands. Let r' be the latest round for which i misreports. That is, $r' = \max\{r : d'_{i,r} \neq d_{i,r}\}$. Suppose that $d'_{i,r'} < d_{i,r'}$. We show that, all else being equal, i could (weakly) improve her utility by instead reporting $d'_{i,r'} = d_{i,r'}$. The proof that reporting $d'_{i,r'} > d_{i,r'}$ is also (weakly) worse than reporting $d'_{i,r'} = d_{i,r'}$ is almost identical and can be found in Appendix B.2. It follows from this that FL is strategy-proof, since any non-truthful reports can be converted to truthful reports one round at a time, (weakly) improving i's utility.

We consider parallel universes: one in which agent i misreports $d'_{i,r'}$ at round r' (the 'misreported instance') and one in which she truthfully reports $d_{i,r}$ (the 'truthful instance,' even though i's reports prior to r' may yet be non-truthful). All other reports are identical in both universes. We denote allocations and tokens in the misreported instance using a' and t', respectively, and in the truthful instance by a and t. We denote by D_r and D'_r the total demand D at round r in the truthful and misreported instances, respectively.

We first note that for all rounds prior to r', the allocations in the truthful and misreported instances are the same.

Lemma 30. *For all rounds $r < r'$ and for all agents j, $a'_{j,r} = a_{j,r}$.*

Proof. The mechanism does not take future reports into account, so because agents' demands in both instances are identical up to round r', so are the allocations. ∎

We next show a monotonicity lemma, which says that agent i's allocation at round r' is (weakly) smaller in the misreported instance than the truthful instance, and all other agents' allocations are (weakly) larger.

Lemma 31. *For all agents $j \neq i$, we have that $a'_{j,r'} \geq a_{j,r'}$. Further, $a'_{i,r'} \leq a_{i,r'}$.*

Proof. We prove the statement for all $j \neq i$. The statement for i follows immediately because the total number of allocated resources is fixed. Observe first that

$$D'_{r'} = \sum_{k \in [n]} \min(d'_{k,r'}, t_{k,r'}) \leqslant \sum_{k \in [n]} \min(d_{k,r'}, t_{k,r'}) = D_{r'},$$

since i's demand decreases in the misreported instances but all other demands and token counts stay the same. Let x' denote the objective value in FL's call to PSWC in the misreported instance, and x in the truthful instance.

Suppose that $E \leqslant D'_{r'} \leqslant D_{r'}$. Suppose first that $x' > x$. Then, by Lemma 28, for all $j \neq i$, we have

$$a'_{j,r'} = \min(x'e_j, d_{j,r'}, t_{j,r'}) \geqslant \min(xe_j, d_{j,r'}, t_{j,r'}) = a_{j,r'}.$$

Next, suppose that $x' \leqslant x$. Then, again by Lemma 28 and the fact that $d'_{i,r'} < d_{i,r'}$,

$$a'_{i,r'} = \min(x'e_i, d'_{i,r'}, t_{i,r'}) \leqslant \min(xe_i, d_{i,r'}, t_{i,r'}) = a_{i,r'}.$$

And, for all $j \neq i$,

$$a'_{j,r'} = \min(x'e_j, d_{j,r'}, t_{j,r'}) \leqslant \min(xe_j, d_{j,r'}, t_{j,r'}) = a_{j,r'}.$$

Because $a'_{k,r'} \leqslant a_{k,r'}$ for all agents k, and $\sum_{k \in [n]} a_{k,r'} = \sum_{k \in [n]} a'_{k,r'}$, it must be the case that $a'_{k,r'} = a_{k,r'}$ for all k, which satisfies the statement of the lemma.

Next, suppose that $D'_{r'} < E \leqslant D_{r'}$. By the definition of FL, $a'_{k,r'} \geqslant \min(d'_{k,r'}, t_{k,r'})$ for all k, and $a_{k,r'} \leqslant \min(d_{k,r'}, t_{k,r'})$ for all k. Since $\min(d'_{j,r'}, t_{j,r'}) = \min(d_{j,r'}, t_{j,r'})$ for all $j \neq i$, we have that $a'_{j,r'} \geqslant a_{j,r'}$, implying also that $a'_{i,r'} \leqslant a_{i,r'}$.

Finally, suppose that $D'_{r'} \leqslant D_{r'} < E$. Suppose first that $x' \leqslant x$. Then, by Lemma 28 and the assumption that $d'_{i,r'} < d_{i,r'}$, we have

$$a'_{i,r'} = \min(t_{i,r'}, \max(x'e_i, d'_{i,r'})) \leqslant \min(t_{i,r'}, \max(xe_i, d_{i,r'})) = a_{i,r'}$$

and

$$a'_{j,r'} = \min(t_{j,r'}, \max(x'e_j, d_{j,r'})) \leqslant \min(t_{j,r'}, \max(xe_j, d_{j,r'})) = a_{j,r'}$$

94

for all $j \neq i$. Because $a'_{k,r'} \leqslant a_{k,r'}$ for all agents k, and $\sum_{k \in [n]} a_{k,r'} = \sum_{k \in [n]} a'_{k,r'}$, it must be the case that $a'_{k,r'} = a_{k,r'}$ for all k, which satisfies the lemma's statement. Next, suppose that $x' > x$. Then, again by Lemma 28, for all $j \neq i$, we have

$$a'_{j,r'} = \min(t_{j,r'}, \max(x'e_j, d_{j,r'})) \geqslant \min(t_{j,r'}, \max(xe_j, d_{j,r'})) = a_{j,r'}.$$

∎

If it is the case that $a'_{i,r'} = a_{i,r'}$, then it must also be the case that $a'_{j,r'} = a_{j,r'}$ for all $j \neq i$. That is, allocations at round r' are the same in the misreported instance as the truthful instance. Therefore, for all rounds $r \leqslant r'$, allocations in both universes would be the same. In all rounds $r > r'$, reports in both universes are the same. Together, these imply that allocations for all rounds $r > r'$ would be the same in both universes. In particular, i does not profit from her misreport and could weakly improve her utility by reporting $d'_{i,r'} = d_{i,r'}$. So, for the remainder of this section, we assume that $a'_{i,r'} < a_{i,r'}$.

Our next lemma states that the resources that i sacrifices in round r' are *high-valued* resources for her. The intuition is that if it were the case that i was being forced to receive low-valued resources under truthful reporting, then she will still be forced to receive the same number of resources when she under-reports her demand (since there is no agent with excess demand to absorb extra resources).

Lemma 32. *If $a'_{i,r'} < a_{i,r'}$, then $a_{i,r'} \leqslant d_{i,r'}$.*

Proof. Suppose for contradiction that $a_{i,r'} > d_{i,r'}$. It must therefore be the case that $D'_{r'} \leqslant D_{r'} < E$, where the first inequality holds because $d'_{j,r'} = d_{j,r'}$ for all $j \neq i$ and $d'_{i,r'} < d_{i,r'}$. Let x denote the objective value of FL's call to PSWC in the truthful instance, and x' in the misreported instance. Suppose that $x' \leqslant x$. Then, by Lemma 28 and the assumption that $d'_{i,r'} < d_{i,r'}$,

$$a'_{i,r'} = \min(t_{i,r'}, \max(x'e_i, d'_{i,r'})) \leqslant \min(t_{i,r'}, \max(xe_i, d_{i,r'})) = a_{i,r'},$$

and for all $j \neq i$,

$$a'_{j,r'} = \min(t_{j,r'}, \max(x'e_j, d_{j,r'})) \leqslant \min(t_{j,r'}, \max(xe_j, d_{j,r'})) = a_{j,r'}.$$

Because $a'_{k,r'} \leqslant a_{k,r'}$ for all agents k, and $\sum_{k \in [n]} a_{k,r'} = \sum_{k \in [n]} a'_{k,r'}$, it must be the case that $a'_{k,r'} = a_{k,r'}$ for all k. This contradicts the assumption that $a'_{i,r'} < a_{i,r'}$.

Now suppose that $x' > x$. Note that $xe_i > d_{i,r'} > d'_{i,r'}$, where the first inequality holds because $a_{i,r'} > d_{i,r'}$. Then, again by Lemma 28 and the previous observation, we have

$$a'_{i,r'} = \min(t_{i,r'}, \max(x'e_i, d'_{i,r'})) = \min(t_{i,r'}, x'e_i)$$

$$\geqslant \min(t_{i,r'}, xe_i) = \min(t_{i,r'}, \max(xe_i, d_{i,r'})) = a_{i,r'},$$

which contradicts $a'_{i,r} < a_{i,r}$. Since we arrive at a contradiction in all cases, the lemma statement must be true. ∎

As a corollary, we can write the difference in utility between the truthful and misreported instances that i derives from round r'.

Corollary 33. $u_{i,r'}(a_{i,r'}) - u_{i,r'}(a'_{i,r'}) = H(a_{i,r'} - a'_{i,r'})$.

Proof. Because $a'_{i,r'} < a_{i,r'} \leqslant d_{i,r'}$, we can substitute the utility values from Equation (4.2):

$$u_{i,r'}(a_{i,r'}) - u_{i,r'}(a'_{i,r'}) = a_{i,r'}H - a'_{i,r'}H = H(a_{i,r'} - a'_{i,r'}).$$

∎

For a fixed agent k, denote by r_k the round at which agent k runs out of tokens in the truthful instance. That is, r_k is the first (and only) round with $a_{r_k} = t_{k,r_k} > 0$. Note that $r_i \geqslant r'$, since $a_{i,r'} > 0$. Given this, our next lemma states that, under certain conditions, the effect of i's misreport, $d'_{i,r} < d_{i,r}$, is to increase the objective value of FL's call to PSWC.

Lemma 34. *Let $r < r_i$ (i.e., $a_{i,r} < t_{i,r}$). Suppose $t'_{j,r} \leqslant t_{j,r}$ for all agents $j \neq i$. Suppose that either $\min(D_r, D'_r) \geqslant E$ or $\max(D_r, D'_r) < E$. Then $x' \geqslant x$, where x' denotes the objective value of FL's call to PSWC in the misreported instance and x in the truthful instance.*

Proof. First, suppose that $\min(D_r, D'_r) \geqslant E$. Suppose for contradiction that $x' < x$. By Lemma 28, for all $j \neq i$,

$$a'_{j,r} = \min(x'e_j, d_{j,r}, t'_{j,r}) \leqslant \min(xe_j, d_{j,r}, t_{j,r}) = a_{j,r},$$

where the inequality follows from the assumption that $x' < x$ and that $t'_{j,r} \leqslant t_{j,r}$. Further,

$$a'_{i,r} = \min(x'e_i, d_{i,r}, t'_{i,r}) \leqslant \min(x'e_i, d_{i,r}) \leqslant \min(xe_i, d_{i,r}) = \min(xe_i, d_{i,r}, t_{i,r}) = a_{i,r},$$

where the second inequality follows from the assumption that $x' < x$, and the second to the last equality follows from the assumption that $a_{i,r} < t_{i,r}$. Therefore, $a'_{k,r} \leqslant a_{k,r}$ for all agents k. Since $\sum a'_{k,r} = \sum a_{k,r}$, it must be the case that $a'_{k,r} = a_{k,r}$ for all agents k. Now, by the definition of FL in this case, $a_{k,r}/e_k \leqslant x' < x$ for all agents k with $a_{k,r} > 0$. Therefore x is not the optimal objective value of PSWC in the truthful instance, a contradiction. Thus, $x' \geqslant x$.

Next, suppose that $\max(D_r, D'_r) < E$. Suppose for contradiction that $x' < x$. By Lemma 28,

$$a'_{j,r} = \min(t'_{j,r}, \max(x'e_j, d_{j,r})) \leqslant \min(t_{j,r}, \max(xe_j, d_{j,r})) = a_{j,r},$$

for all $j \neq i$, where the inequality follows from the assumption that $x' < x$ and that $t'_{j,r} \leqslant t_{j,r}$. Further, we have

$$a'_{i,r} = \min(t'_{i,r}, \max(x'e_i, d_{i,r})) \leqslant \max(x'e_i, d_{i,r})$$
$$\leqslant \max(xe_i, d_{i,r}) = \min(t_{i,r}, \max(xe_i, d_{i,r})) = a_{i,r},$$

97

where the second inequality follows from the assumption that $x' < x$ and the second to last equality from the assumption $a_{i,r} < t_{i,r}$. Therefore, $a'_{k,r} \leqslant a_{k,r}$ for all agents k. Since $\sum a'_{k,r} = \sum a_{k,r}$, it must be the case that $a'_{k,r} = a_{k,r}$ for all agents k. Consider all agents with $\min(d_{k,r}, t_{k,r}) < a_{k,r}$ (i.e., those agents for which the first constraint in the PSWC program binds in the truthful instance). For all such agents, we have

$$\min(d_{k,r}, t_{k,r}) < a_{k,r}$$

$$\implies \qquad d_{k,r} < a_{k,r} \leqslant t_{k,r}$$

$$\implies \qquad d_{k,r} < a'_{k,r} \leqslant t'_{k,r}$$

$$\implies \min(d_{k,r}, t'_{k,r}) < a'_{k,r},$$

which implies that the constraints bind in the misreported instance as well. Therefore, $a'_{k,r}/e_k \leqslant x' < x$ for all agents k for which the first constraint binds in the truthful instance. Therefore x is not the optimal objective value of the PSWC program in the truthful instance, a contradiction. Thus, $x' \geqslant x$. ■

Using Lemma 34, we show our main lemma. This lemma allows us to make an inductive argument that, after giving up some resources in round r', i's allocation is (weakly) larger for all future rounds in the misreported instance than the truthful instance.

Lemma 35. *Let $r' < r < r_i$ (i.e., $a_{i,r} < t_{i,r}$). Suppose that $t'_{j,r} \leqslant t_{j,r}$ for all agents $j \neq i$. Then for all $j \neq i$, either: (1) $a'_{j,r} = t'_{j,r}$, or (2) $a'_{j,r} \geqslant a_{j,r}$.*

Proof. Note that $t'_{j,r} \leqslant t_{j,r}$ for all $j \neq i$ implies that $t'_{i,r} \geqslant t_{i,r}$, which we use in the proof. Also, because $r > r'$, we know that $d'_{i,r} = d_{i,r}$, as r' is the last round for which $d'_{i,r} \neq d_{i,r}$. We assume that condition (1) from the lemma statement is false (i.e., $a'_{j,r} < t'_{j,r}$) and show that condition (2) must hold. Suppose first that $D_r < E$. Then, because $a_{i,r} < t_{i,r}$, we know that $d_{i,r} \leqslant t_{i,r} \leqslant t'_{i,r}$. This implies that $\min(d_{i,r}, t_{i,r}) = \min(d_{i,r}, t'_{i,r}) = d_{i,r}$. Let $j \neq i$. Since $t'_{j,r} \leqslant t_{j,r}$, we have

$\min(d_{j,r}, t'_{j,r}) \leqslant \min(d_{j,r}, t_{j,r})$. Therefore, it is the case that $D'_r \leqslant D_r < E$. By Lemma 28 and the assumption that $a'_{j,r} < t'_{j,r}$, it must be the case that $a'_{j,r} = \max(d_{j,r}, x'e_j)$. Further, by Lemma 34, we know that $x' \geqslant x$. Therefore, we have

$$a_{j,r} = \max(d_{j,r}, xe_j) \leqslant \max(d_{j,r}, x'e_j) = a'_{j,r}.$$

That is, condition (2) from the lemma statement holds.

Now suppose that $D_r \geqslant E$. Then, from the definition of the mechanism, we have that $a_{j,r} \leqslant \min(d_{j,r}, t_{j,r}) \leqslant d_{j,r}$. If it is the case that $D'_r < E$, then we have that $a'_{j,r} \geqslant \min(d_{j,r}, t'_{j,r}) = d_{j,r}$, where the equality holds because otherwise we would have $a'_{j,r} \geqslant t'_{j,r}$, violating the assumption that $a'_{j,r} < t'_{j,r}$. Using these inequalities, we have $a'_{j,r} \geqslant d_{j,r} \geqslant a_{j,r}$, so condition (2) from the statement of the lemma holds. Finally, it may be the case that $D_r \geqslant E$ and $D'_r \geqslant E$. By Lemma 28 and the assumption that $a'_{j,r} < t'_{j,r}$, we have

$$a'_{j,r} = \min(d_{j,r}, x'e_k) \geqslant \min(d_{j,r}, xe_k) = a_{j,r},$$

where the inequality follows from Lemma 34. Thus, condition (2) of the lemma statement holds. ∎

Finally, we prove that the flexible lending mechanism is strategy-proof. This proof establishes that misreporting $d'_{i,r}$ is never beneficial for an agent.

Theorem 36. *The flexible lending mechanism satisfies SP.*

Proof. We first observe that for every $r \leqslant r_i$, $t'_{j,r} \leqslant t_{j,r}$ for every $j \neq i$. This is true for every $r \leqslant r'$ because $a'_{j,r} = a_{j,r}$ for $r < r'$, by Lemma 30. For $r = r' + 1$, it follows from Lemma 31, which says that $a'_{j,r'} \geqslant a_{j,r'}$. For all subsequent rounds, up to and including $r = r_i$, it follows inductively from Lemma 35: $t'_{j,r} \leqslant t_{j,r}$ implies that either $a'_{j,r} = t'_{j,r}$, in which case $t'_{j,r+1} = 0 \leqslant t_{j,r+1}$, or $a'_{j,r} \geqslant a_{j,r}$, in which case $t'_{j,r+1} = t'_{j,r} - a'_{j,r} \leqslant t_{j,r} - a_{j,r} = t_{j,r+1}$).

Consider an arbitrary round $r \neq r'$, with $r \leqslant r_i$. By the above argument, we know that $t'_{j,r} \leqslant t_{j,r}$ for all $j \neq i$. Further, because reports in the truthful and misreported instances are identical on all rounds $r \neq r'$, we have that $d_{k,r} = d'_{k,r}$ for all $k \in [n]$. Therefore, by Lemma 29, $a'_{i,r} \geqslant a_{i,r}$. For rounds $r > r_i$, it is also true that $a'_{i,r} \geqslant a_{i,r}$, since $a_{i,r} = 0$ for these rounds by the definition of r_i. Finally,

$$U_{i,R}(\mathbf{a_i}) - U_{i,R}(\mathbf{a'_i}) = \sum_{r=1}^{R}(u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r}))$$

$$= (u_{i,r'}(a_{i,r'}) - u_{i,r'}(a'_{i,r'})) + \sum_{r \neq r'}(u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r}))$$

$$= H(a_{i,r'} - a'_{i,r'}) - \sum_{r \neq r'}(u_{i,r}(a'_{i,r}) - u_{i,r}(a_{i,r}))$$

$$\geqslant H(a_{i,r'} - a'_{i,r'}) - H(a_{i,r'} - a'_{i,r'}) = 0$$

Here, the third transition follows from Lemma 33, and the final transition follows because $\sum_{r \neq r'}(a'_{i,r} - a_{i,r}) = a_{i,r'} - a'_{i,r'}$, and every term in the sum is positive.

The proof for the case where $d'_{i,r'} > d_{i,r'}$ is in the Appendix. Together, they show that i achieves (weakly) higher utility by truthfully reporting her demand $d_{i,r'}$ at round r, rather than misreporting $d'_{i,r'} \neq d_{i,r'}$. By the argument at the start of this subsection, this is sufficient to prove strategy-proofness. ∎

4.5.4 Approximating Sharing Incentives

Unfortunately, FL fails to satisfy SI, and may give an agent as little as half of her SI share.

Theorem 37. *FL does not satisfy α-SI for any $\alpha > 0.5$.*

Proof. Consider an instance with R rounds, and $R+1$ agents, each with endowment $e_i = 1$. Agent 1 has $d_{1,1} = d_{1,R} = 1$ and $d_{1,2} = \ldots = d_{1,R-1} = 0$, agent 2 has $d_{2,r} = R$ for all rounds r, and all other agents have $d_{i,r} = 0$ for all rounds r. In round

1, agent 1 receives allocation $a_{1,1} = 1$ and agent 2 receives $a_{2,1} = R$. For rounds $r = 2, \ldots, R - 1$, each agent $j \neq 2$ receives allocation $a_{j,r} = 1 + 1/R$. Therefore, in round R, agent 1 receives $a_{1,R} = R - 1 - (R-2)(1+1/R) = 2/R$. Her total utility is therefore $((R+2)/R)H + (R - (R+2)/R)L$, compared to total utility $2H + (R-2)L$ that she would have received by not participating in the mechanism. For $L = 0$, the ratio of these utilities approaches 0.5 as $R \to \infty$. ∎

However, FL does provide a 0.5 approximation guarantee to SI, as we show in the remainder of this subsection. We suppose that agent i truthfully reports her demand $d_{i,r}$ for all rounds (since FL is SP, she could do no better by lying), and show that she receives at least half as much utility as she would by not participating.

Recall that for every agent i, we denote by r_i the first round at which $a_{i,r_i} = t_{i,r_i} > 0$. For every agent i, define sets B_i and A_i to be the agents that run out of tokens before and after i, respectively. Formally,

$$B_i = \{j : r_j \leqslant r_i \text{ and } a_{j,r_i}/e_j < a_{i,r_i}/e_i\}$$

$$A_i = \{j : r_j \geqslant r_i \text{ and } r_j = r_i \implies a_{j,r_i}/e_j \geqslant a_{i,r_i}/e_i\}.$$

For a round r, define

$$s_{i,r} = a_{i,r} - e_i \frac{\sum_{j \in A_i} a_{j,r}}{\sum_{j \in A_i} e_j}.$$

That is, $s_{i,r}$ is the number of resources i gets more than the (endowment weighted) average number of resources for agents in A_i. Note further that

$$\sum_{r=1}^{R} s_{i,r} = \sum_{r=1}^{R} a_{i,r} - \frac{e_i}{\sum_{j \in A_i} e_j} \sum_{j \in A_i} \sum_{r=1}^{R} a_{j,r} = e_i - \frac{e_i}{\sum_{j \in A_i} e_j} \sum_{j \in A_i} e_j = 0.$$

Lemma 38. *For every agent i and every round r, $s_{i,r} \leqslant \min(d_{i,r}, a_{i,r})$.*

Proof. If $a_{i,r} \leqslant d_{i,r}$, then the lemma statement says that $s_{i,r} \leqslant a_{i,r}$, which is obviously true by the definition of $s_{i,r}$. If $a_{i,r} > d_{i,r}$, then we know from the definition of FL

101

that $\sum_{j\in[n]} \min(d_{j,r}, t_{j,r}) < E$, and $a_{i,r} = \min(xe_i, t_{i,r})$, where x is the objective value of FL's call to the PSWC program. Further, all agents with $\frac{a_{j,r}}{e_j} < \frac{a_{i,r}}{e_i} \leqslant x$ are those with $a_{j,r} = t_{j,r}$, so by definition, $r_j \leqslant r_i$ and $\frac{a_{j,r}}{e_j} < \frac{a_{i,r}}{e_i}$, which means $j \in B_i$. Therefore, $\frac{a_{j,r}}{e_j} \geqslant \frac{a_{i,r}}{e_i}$ for all $j \in A_i$, which implies $\frac{\sum_{j\in A_i} a_{j,r}}{\sum_{j\in A_i} e_j} \geqslant \frac{a_{i,r}}{e_i}$. To complete the proof, note that

$$s_{i,r} = a_{i,r} - e_i \frac{\sum_{j\in A_i} a_{j,r}}{\sum_{j\in A_i} e_j} \leqslant a_{i,r} - e_i \frac{a_{i,r}}{e_i} = 0 \leqslant d_{i,r} = \min(d_{i,r}, a_{i,r}).$$

■

Theorem 39. *Under FL, agents receive at least half the number of high-valued resources that they would have received under static allocations.*

Proof. Let S denote the number of high-valued resources that agent i receives under static allocations. While i has tokens remaining, under FL, she is guaranteed to get as many resources as she demands up to her endowment e_i. Thus, for these rounds, she would obtain no additional high-valued resources from not participating in the mechanism. However, there is the possibility that by participating in the mechanism, she runs out of tokens prematurely, thus missing out on resources in later rounds that she wants, and would have received by not participating in the mechanism (as in the proof of Theorem 37). The proof proceeds by showing that for every resource that i does not receive due to a lack of tokens, she must have received at least one high-valued resource in an earlier round.

Suppose first that $a_{i,r_i} \geqslant e_i$. We have the following inequality:

$$\sum_{r\leqslant r_i} \min(d_{i,r}, a_{i,r}) \geqslant \sum_{r\leqslant r_i} s_{i,r} = -\sum_{r>r_i} s_{i,r} = \sum_{r>r_i} \left(e_i \frac{\sum_{j\in A_i} a_{j,r}}{\sum_{j\in A_i} e_j} \right)$$

$$= \sum_{r>r_i} \left(\frac{E}{\sum_{j\in A_i} e_j} \right) e_i \geqslant (T - r_i)e_i. \quad (4.1)$$

The first inequality follows from Lemma 38, and the second inequity because $\sum_{j \in A_i} e_j \leqslant E$. The first equality holds because $\sum_{r=1}^{R} s_{i,r} = 0$, and the second equality holds because $a_{i,r} = 0$ for all $r > r_i$. The third equality holds because for rounds $r > r_i$, only agents in A_i remain active, so all resources are allocated to them.

Note that S, the number of high-valued resources that i receives by not sharing, is upper bounded by

$$S \leqslant \sum_{r=1}^{R} \min(d_{i,r}, e_i) \leqslant \sum_{r \leqslant r_i} \min(d_{i,r}, e_i) + \sum_{r > r_i} e_i$$

$$\leqslant \sum_{r \leqslant r_i} \min(d_{i,r}, a_{i,r}) + \sum_{r > r_i} e_i$$

$$= \sum_{r \leqslant r_i} \min(d_{i,r}, a_{i,r}) + (T - r_i)e_i$$

$$\leqslant 2 \sum_{r \leqslant r_i} \min(d_{i,r}, a_{i,r}).$$

The third inequality holds because under FL guarantees each agent $\min(d_{i,r}, e_i)$ resources, provided they have sufficient tokens remaining, which is the case because we assume $a_{i,r_i} \geqslant e_i$. The final inequality follows from Equation (4.1). Since agent i receives exactly $\sum_{r \leqslant r_i} \min(d_{i,r}, a_{i,r}) \geqslant S/2$ resources from participating in FL, the lemma holds in this case.

Second, suppose that $a_{i,r_i} < e_i$. We have the following inequality:

$$\sum_{r \leqslant r_i} \min(d_{i,r}, a_{i,r}) \geqslant \sum_{r < r_i} \min(d_{i,r}, a_{i,r}) \geqslant \sum_{r < r_i} s_{i,r} = -\sum_{r > r_i} s_{i,r} - s_{i,r_i}$$

$$\geqslant e_i(T - r_i) + e_i \frac{\sum_{j \in A_i} a_{j,r_i}}{\sum_{j \in A_j} e_j} - a_{i,r_i} \geqslant e_i(T - r_i) + e_i - a_{i,r_i} = e_i(T - r_i + 1) - a_{i,r_i}$$

$$\tag{4.2}$$

The first inequality holds because $\min(d_{i,r_i}, a_{i,r_i}) \geqslant 0$. The second inequality follows from Lemma 38, and the third inequality holds from Equation (4.1) and the

definition of s_{i,r_i}. The fourth inequality holds because at round r_i, agent i receives allocation $a_{i,r_i} < e_i$, therefore every agent $j \in B_i$ receives allocation $a_{j,r_i} < e_j$, therefore $\sum_{j \in A_i} a_{j,r_i} \geqslant \sum_{j \in A_i} e_j$.

As with the previous case, we can derive an upper bound on S, the number of high-valued resources i would receive by not sharing. First, suppose that $a_{i,r_i} > d_{i,r_i}$. Then we have

$$S \leqslant \sum_{r=1}^{R} \min(d_{i,r}, e_i) \leqslant \sum_{r<r_i} \min(d_{i,r}, e_i) + d_{i,r_i} + \sum_{r>r_i} e_i$$

$$\leqslant \sum_{r<r_i} \min(d_{i,r}, a_{i,r}) + \min(d_{i,r_i}, a_{i,r_i}) + \sum_{r>r_i} e_i$$

$$= \sum_{r \leqslant r_i} \min(d_{i,r}, a_{i,r}) + (T - r_i)e_i$$

$$\leqslant \sum_{r \leqslant r_i} \min(d_{i,r}, a_{i,r}) + (T - r_i + 1)e_i - a_{i,r_i}$$

$$\leqslant 2 \sum_{r \leqslant r_i} \min(d_{i,r}, a_{i,r})$$

The third inequality holds because FL guarantees each agent $\min(d_{i,r}, e_i)$ resources, provided they have sufficient tokens remaining, and by the assumption that $a_{i,r_i} > d_{i,r_i}$, the fourth inequality from the assumption that $a_{i,r_i} < e_i$, and the final inequality from Equation (4.2). Next, suppose that $a_{i,r_i} \leqslant d_{i,r_i}$. Then we have

$$S \leqslant \sum_{r=1}^{R} \min(d_{i,r}, e_i) \leqslant \sum_{r<r_i} \min(d_{i,r}, e_i) + e_i + \sum_{r>r_i} e_i$$

$$\leqslant \sum_{r<r_i} \min(d_{i,r}, a_{i,r}) + a_{i,r_i} + (e_i - a_{i,r_i}) + \sum_{r>r_i} e_i$$

$$= \sum_{r \leqslant r_i} \min(d_{i,r}, a_{i,r}) + (T - r_i + 1)e_i - a_{i,r_i}$$

$$\leqslant 2 \sum_{r \leqslant r_i} \min(d_{i,r}, a_{i,r})$$

104

The third inequality holds because FL guarantees each agent $\min(d_{i,r}, e_i)$ resources, provided they have sufficient tokens remaining, the equality from the assumption that $a_{i,r_i} \leqslant d_{i,r_i}$, and the final inequality from Equation (4.2).

As with the previous case, $e_i(T - r_i + 1) - a_{i,r_i}$ is an upper bound on the number of H valued resources that i may have been able to receive in rounds $r \geqslant r_i$ had she not participated in the mechanism, over and above those she receives by participating. $\sum_{r \leqslant r_i} \min(d_{i,r}, a_{i,r})$ is the number of H valued resources she receives by participating in the mechanism. Therefore $\sum_{r \leqslant r_i} \min(d_{i,r}, a_{i,r}) + e_i(T - r_i + 1) - a_{i,r_i} \leqslant 2 \sum_{r \leqslant r_i} \min(d_{i,r}, a_{i,r})$ is an upper bound on the number of H valued resources i would receive by not participating in the mechanism. Therefore, i receives at least half as many H valued resources from participating as she would have by not participating.
∎

Note that Theorem 39 implies the desired approximation. Suppose that i obtains utility $SH + (Re_i - S)L$ by not participating in the mechanism. Theorem 39 in combination with the fact that she will receive the same number of resources overall whether she participates or not, implies that, by participating, she gets at least $SH/2 + (Re_i - S/2)L \geqslant SH/2 + (Re_i/2 - S/2)L = (SH + (Re_i - S)L)/2$.

4.5.5 Limit Efficiency for Symmetric Agents

In this section, we prove that, under certain assumptions, FL is efficient in the limit as the number of rounds grows large. Suppose that each agent has the same endowment. Without loss of generality, suppose that each agent has $e_i = 1$. Further, suppose that demands are drawn i.i.d. across rounds and that the distribution within rounds treats agents symmetrically, either demands are drawn i.i.d. across agents, or there is correlation that treats all agents symmetrically.

Theorem 40. *When demands are drawn i.i.d. across rounds and agents are sym-*

metric, FL achieves an $(R - R^{2/3})/R$ fraction of the optimal efficiency with probability at least $1 - n^3/R^{1/3}$. In particular, FL approaches full efficiency with high probability in the limit as the number of rounds grows large.

Proof. Suppose we are in a world where tokens are unlimited. Let Q be a random variable denoting how many tokens a single agent i would spend (i.e., how many resources i would be allocated) in a single round. Note that Q can never take a value larger than n, since only n resources are allocated per round. Note that by the symmetry of the agents, Q is independent of the identity of any single agent, and independent of the particular round since FL allocates independently of the round. By symmetry, $\mathbb{E}(Q) = 1$. Let $\mathrm{StdDev}(Q) = \sigma \leqslant n$, where the inequality holds because Q is bounded by n. Let $r = R - R^{2/3}$ and let Q_r be a random variable denoting the number of tokens i would spend before the start of round $r+1$. Because demands are drawn independently across rounds, and no agent runs out of tokens, $\mathbb{E}(Q_r) = r$ and $\mathrm{StdDev}(Q_r) = \sqrt{r}\sigma$.

Consider the probability that agent i spends at least R tokens in the first r rounds:

$$P(Q_r \geqslant R) = P(Q_r - \mathbb{E}(Q_r) \geqslant R - r)$$

$$= P(Q_r - \mathbb{E}(Q_r) \geqslant R^{2/3})$$

$$= P\left(Q_r - \mathbb{E}(Q_r) \geqslant \frac{R^{1/6}}{\sigma}\sqrt{R}\sigma\right)$$

$$\leqslant P\left(Q_r - \mathbb{E}(Q_r) \geqslant \frac{R^{1/6}}{\sigma}\sqrt{r}\sigma\right)$$

$$\leqslant \frac{\sigma^2}{R^{1/3}}$$

Here the final inequality follows from Chebyshev's concentration inequality, because $\sqrt{r}\sigma$ is the standard deviation of Q_r. Taking a union bound over all n agents, the probability that any agent spends at least R tokens in the first r rounds is at most $n\sigma^2/R^{1/3} \leqslant n^3/R^{1/3}$.

106

Now suppose agents are limited by R tokens. If some agent runs out of tokens within r rounds in this world, then it must also be the case that some agent spent at least R tokens within r rounds in the unlimited token world. Therefore, the probability that any agent runs out of tokens is at most the probability that some agent spends more than R tokens in the unlimited token world, which is at most $n^3/R^{1/3}$. This approaches 0 as $R \to \infty$. So, with probability going to 1, no agent runs out of tokens before round r.

By the definition of FL, full efficiency is achieved on all rounds for which no agents have their allocation limited by lack of tokens. Therefore, with probability going to 1, FL allocates efficiently for the first r rounds. Therefore, because demands are i.i.d. across rounds, the expected efficiency of the mechanism approaches at least an $r/R = (R - R^{2/3})/R$ fraction of the optimal efficiency. This fraction approaches 1 as $R \to \infty$. ∎

4.6 **T**-Period Mechanism

We have shown that FL satisfies strategy-proofness and a theoretical asymptotic efficiency guarantee. Further, as we show in §4.7, FL exhibits only small efficiency loss in practice in settings where our theoretical guarantee does not apply. However, FL does not achieve (full) sharing incentives. In settings where agents require a strong guarantee in order to participate, it may be desirable to strictly enforce sharing incentives, in which case FL is not a suitable choice. In this section, we introduce the *T-Period mechanism*, which satisfies both SP and SI. While the *T*-Period mechanism does exhibit some gains from sharing (i.e., is more efficient than static allocation), it sacrifices some efficiency relative to FL.

The T-Period mechanism splits the rounds into periods of length $2T$.[1] For the first T rounds of each period, we allow the agents to 'borrow' unwanted resources from others. In the last T rounds of each period, the agents 'pay back' the resources so that their cumulative allocation across the entire period is equal to their endowment, $2Te_i$.

The allocations in the second set of T rounds are independent of reports and determined completely by the allocations in the first set of T rounds. Note that because the number of resources that an agent i can pay back over T rounds is bounded by Te_i, we allow an agent to borrow at most Te_i resources (i.e., receive at most $2Te_i$ resources) over the first T rounds of a period.

In Algorithm 7, each agent i has a borrowing limit, b_i, which is defined to be the maximum amount of resources that agent i can borrow in whatever remains of the first T rounds of each period. For our analysis, we denote the value of b_i at the start of round r by $b_{i,r}$. At the beginning of each period, we set $b_{i,r}$ to be Te_i, because agent i can at most pay back her whole endowment, e_i, at every T 'payback' rounds. We again define $\bar{d}_i$ to be the allocatable demand of agent i at each round of the first T rounds and refer to $\bar{d}'_{i,r}$ as agent i's allocatable demand at round r. At each round r, the allocatable demand of agent i is the minimum of her reported demand $d'_{i,r}$, and her endowment plus her borrowing limit, $e_i + b_{i,r}$.

We illustrate the T-Period mechanism with an example.

Example 7. *Consider the instance from Example 6, where each agent has endowment $e_i = 1$ and demands are given by:*

---

[1] For convenience, we suppose that R is a multiple of $2T$. If this is not the case, we can adapt the mechanism by returning each agent their endowment for any leftover rounds.

ALGORITHM 7: T-Period Mechanism

Input: Agents' reported demands, $\mathbf{d'}$, and their endowments, $\mathbf{e}$
Output: Agents' allocations, $\mathbf{a}$

1 **for** $r \in \{1, \ldots, R\}$ **do**
2   **if** $(r \bmod 2T) = 1$ **then**
3        $\mathbf{b} \leftarrow T\mathbf{e}$;   /* b_i is the amount that i is able to borrow */
4        $\mathbf{y} \leftarrow \mathbf{0}$;   /* resources received so far this period. */
5   **end**
6   **if** $1 \leqslant (r \bmod 2T) \leqslant T$ **then**
7        $\bar{\mathbf{d}} \leftarrow \min(\mathbf{d'_{\cdot,r}}, \mathbf{e} + \mathbf{b})$;   /* $\bar{d}_i$ is i's allocatable demand */
8        $D \leftarrow \sum_{i \in [n]} \bar{d}_i$;
9   **if** $D \geqslant E$ **then**
10           $\mathbf{a_{\cdot,r}} \leftarrow \mathrm{PSWC}(A = E, \mathbf{l} = \bar{\mathbf{d}}, \mathbf{m} = \mathbf{0}, \mathbf{w} = \mathbf{e})$;
11   **else**
12           $\mathbf{a_{\cdot,r}} \leftarrow \mathrm{PSWC}(A = E, \mathbf{l} = \mathbf{e} + \mathbf{b}, \mathbf{m} = \bar{\mathbf{d}}, \mathbf{w} = \mathbf{e})$;
13   **end**
14        $\mathbf{y} \leftarrow \mathbf{y} + \mathbf{a_{\cdot,r}}$;
15        $\mathbf{b} \leftarrow \mathbf{b} - \max(\mathbf{0}, \mathbf{a_{\cdot,r}} - \mathbf{e})$;
16   **else**
17        $\mathbf{a_{\cdot,r}} \leftarrow \frac{1}{T}(2T\mathbf{e} - \mathbf{y})$;
18   **end**
19 **end**

	$d_{i,1}$	$d_{i,2}$	$d_{i,3}$	$d_{i,4}$
$i = 1$	3	1	1	0
$i = 2$	0	2	1	2
$i = 3$	0	0	0	4

When $T = 1$, agents can 'borrow' resources at odd rounds and 'pay back' those resources at even rounds. Therefore, the maximum allocatable demand for each agent and at each round is 2, because the 'payback' period only has one round. The 1-Period (1-P) mechanism allocates resources as follows.

	$a_{i,1}^{1\text{-}P}$	$a_{i,2}^{1\text{-}P}$	$a_{i,3}^{1\text{-}P}$	$a_{i,4}^{1\text{-}P}$
$i = 1$	2	0	1	1
$i = 2$	0.5	1.5	1	1
$i = 3$	0.5	1.5	1	1

At round 1, agent 1 wants 2 extra resources in addition to her endowment. However, under 1-P, she can only afford 1 extra resource. She borrows 0.5 resources from

agent 2 and 0.5 resources from agent 3. At round 2, agent 1 pays back agents 2 and 3 and receives zero resources. When $T = 1$, the mechanism rigidly forces agents to pay back resources right after they borrow them. Agent 1 would prefer to get her high-valued resource at round 2 and delay paying back agents 2 and 3 until the last round where her demand is zero. Note that agent 3 would also prefer to be paid back in the last round, the only round in which she has non-zero demand.

To see how increasing T allows more flexibility, consider $T = 2$ for the same example. The 2-Period (2-P) mechanism allocates resources as follows.

	$a_{i,1}^{2\text{-}P}$	$a_{i,2}^{2\text{-}P}$	$a_{i,3}^{2\text{-}P}$	$a_{i,4}^{2\text{-}P}$
$i = 1$	3	1	0	0
$i = 2$	0	2	1	1
$i = 3$	0	0	2	2

Agent 2 is allowed to borrow 2 extra resources over the first two rounds, whereas, under the 1-P mechanism, she is never allowed to borrow more than one resource per round. She borrows these two resources at the first round from agents 2 and 3, and pays them back at rounds 3 and 4.

Since the T-Period mechanism increases flexibility over the static mechanism, it provides some gains from sharing. We would expect that increasing T, in general, will improve efficiency as it allows for 'borrowed' resources to be spent more flexibly. In the following subsection, we show that these efficiency gains do not harm SI or SP when $T \leq 2$. Many proofs closely follow those in §4.5 and are deferred to the Appendix.

*4.6.2   Axiomatic Properties of **T**-Period Mechanism*

We first state a lemma characterizing the allocations of the T-Period mechanism that is analogous to Lemma 28

Lemma 41. *Let x denote the objective value of a call to PSWC. Suppose that $1 \leq (r$*

mod $2T$) $\leqslant T$. If $D \geqslant E$, then $a_{i,r} = \min(e_i + b_i, d'_{i,r}, xe_i)$. If $D < E$, then $a_{i,r} = \min(e_i + b_i, \max(d'_{i,r}, xe_i))$.

To prove strategy-proofness of the 1-Period and 2-Period mechanisms, we show that no agent has an incentive to report $d'_{i,r} \neq d_{i,r}$ for any round r. We again consider parallel cases, one in which agent i misreports $d'_{i,r}$ and one in which she truthfully reports $d_{i,r}$ with all other reports the same across the two cases. Allocations and borrowing limits in the former case is denoted by a' and b' respectively, and by a and b in the latter case. Let D_r denote the total allocatable demand at a round r in the truthful case and D'_r denote the total allocatable demand at a round r in the misreported case.

Since the T-Period mechanism resets every $2T$ rounds, we can assume without loss of generality that $R = 2T$ for the sake of reasoning about SP and SI. For rounds $r > T$, the allocations depend completely on the allocations at earlier rounds, and not on the agents' reports, so there is clearly no benefit to an agent for misreporting in these rounds. It remains to show that reporting $d'_{i,r} = d_{i,r}$ is optimal for rounds $r \leqslant T$.

Our next lemma is analogous to Lemma 31.

Lemma 42. *Let $a_{i,r}$ and $a'_{i,r}$ denote the allocations of agent i at round r when she reports $d_{i,r}$ and $d'_{i,r}$, respectively, holding fixed the reports of all agents $j \neq i$ and agent i's reports on all rounds other than r. If $d'_{i,r} < d_{i,r}$ then $a'_{i,r} \leqslant a_{i,r}$, and $a'_{j,r} \geqslant a_{j,r}$ for all $j \neq i$.*

Suppose that i reports $d'_{i,r} \neq d_{i,r}$ for some round r, but this misreport does not change i's allocation (that is, $a'_{i,r} = a_{i,r}$). By Lemma 42, $a'_{j,r} = a_{j,r}$ for all $j \neq i$. Therefore, i's misreport has not changed the allocations at round r. Since all future rounds take into account allocations at previous rounds but not reports, i's misreport has had no effect on the allocations in any round. Thus, i did not benefit from this

111

misreport. We therefore assume that $a'_{i,r} \neq a_{i,r}$ for any round r where i reports $d'_{i,r} \neq d_{i,r}$ in the remainder of this section.

The next lemma and corollary are analogous to Lemma 32 and Corollary 33. They say that if i obtains fewer resources from misreporting at round r, then those resources are all high-valued resources.

Lemma 43. *Hold the reports of all agents $j \neq i$ fixed, and the reports of agent i on all rounds other than r fixed. If i reports $d'_{i,r} < d_{i,r}$ and receives $a'_{i,r} < a_{i,r}$, then $a_{i,r} \leqslant d_{i,r}$.*

As a corollary we obtain a formula for the difference between the utility that agent i receives at round r under truthful reporting and misreporting, when i gets fewer resources in the misreported instance.

Corollary 44. *Hold the reports of all agents $j \neq i$ fixed, and the reports of agent i on all rounds other than r fixed. If i reports $d'_{i,r} < d_{i,r}$ and receives $a'_{i,r} < a_{i,r}$, then $u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r}) = H(a_{i,r} - a'_{i,r})$.*

The next lemma and corollary complement Lemma 43 and Corollary 44 in the case where i receives more resources in the misreported instance than the truthful instance at round r.

Lemma 45. *Hold the reports of all agents $j \neq i$ fixed, and the reports of agent i on all rounds other than r fixed. If i reports $d'_{i,r} > d_{i,r}$ and receives $a'_{i,r} > a_{i,r}$, then $a_{i,r} \geqslant d_{i,r}$.*

Corollary 46. *Hold the reports of all agents $j \neq i$ fixed, and the reports of agent i on all rounds other than r fixed. If i reports $d'_{i,r} > d_{i,r}$ and receives $a'_{i,r} > a_{i,r}$, then $u_{i,r}(a'_{i,r}) - u_{i,r}(a_{i,r}) = L(a'_{i,r} - a_{i,r})$.*

We can now show that misreporting in round T is never beneficial to an agent.

Lemma 47. *An agent never improves her utility by reporting $d'_{i,T} \neq d_{i,T}$.*

As a corollary, we immediately have that the 1-Period mechanism is strategy-proof, because misreporting at round $r = 1 = T$ is not beneficial, and misreporting at round $r = 2 > T$ is not beneficial by our earlier argument.

Corollary 48. *The 1-Period mechanism satisfies strategy-proofness.*

Our next lemma is a monotonicity statement for the borrowing limits: if i's borrowing limit at round r increases, and all other agents' borrowing limits decrease, then i's allocation (weakly) increases and all other agents' allocations (weakly) decrease.

Lemma 49. *Suppose that $r \leqslant T$. If $b'_{i,r} \geqslant b_{i,r}$ and $b'_{j,r} \leqslant b_{j,r}$ for all $j \neq i$, and $d'_{k,r} = d_{k,r}$ for all agents k, then $a'_{i,r} \geqslant a_{i,r}$.*

We now show that the 2-Period mechanism is strategy-proof.

Theorem 50. *The 2-Period mechanism satisfies strategy-proofness.*

Proof. By Lemma 47, no agent can benefit by reporting $d'_{i,2} \neq d_{i,2}$. Similarly, no agent can benefit by reporting $d'_{i,r} \neq d_{i,r}$ for $r \in \{3, 4\}$, because the 2-Period mechanism ignores reports for those rounds. We may therefore assume that $d'_{i,r} = d_{i,r}$ for all agents i and all rounds $r \geqslant 2$.

We show that an agent cannot benefit from reporting $d'_{i,1} < d_{i,1}$. The proof that reporting $d'_{i,1} > d_{i,1}$ is not beneficial is very similar. If $a'_{i,1} = a_{i,1}$, then $a'_{j,1} = a_{j,1}$ for all $j \neq i$, by Lemma 42. Therefore, the allocations are unchanged for all rounds i, as the 2-Period mechanism takes into account allocations at earlier rounds, but not reports, and the allocations at round 1 are the same in the truthful and misreported instances. We therefore assume that $a_{i,1} = a'_{i,1} + k$, for some $k > 0$. This implies that $b_{i,2} = b'_{i,2} - k_i$, for some $k_i \leqslant k$. By Corollary 44, i receives kH more utility in round

1 under truthful reporting than under misreporting. For every $j \neq i$, $a_{j,1} \leqslant a'_{j,1}$, and $b_{j,2} = b'_{j,2} + k_j$, where $\sum_{j \neq i} k_j \leqslant k$. By Lemma 49, $a'_{i,2} \geqslant a_{i,2}$. In the following, we show that $a'_{i,2} \leqslant a_{i,2} + k$. Let x and x' denote the objective value in the T-Period mechanism's call to PSWC when i reports $d_{i,r}$ and $d'_{i,r}$, respectively. We consider four cases, corresponding to whether resources in the truthful and misreported instances are over or under demanded at round 2. Suppose first that $D_2 \geqslant E$ and $D'_2 \geqslant E$. First, suppose that $x' < x$. Then, by Lemma 41,

$$a'_{i,2} = \min(e_i + b'_i, d_{i,2}, x'e_i) = \min(e_i + b_{i,2} + k_i, d_{i,2}, x'e_i)$$

$$\leqslant \min(e_i + b_{i,2}, d_{i,2}, x'e_i) + k_i \leqslant \min(e_i + b_{i,2}, d_{i,2}, xe_i) + k_i \leqslant a_{i,2} + k$$

Next, suppose that $x' \geqslant x$. Then for all $j \neq i$,

$$a'_{j,2} = \min(e_j + b'_j, d_{j,2}, x'e_j) = \min(e_j + b_{j,2} - k_j, d_{j,2}, x'e_j)$$

$$\geqslant \min(e_j + b_{j,2}, d_{j,2}, x'e_j) - k_j \geqslant \min(e_j + b_{j,2}, d_{j,2}, xe_j) - k_j = a_{j,2} - k_j$$

Taking the sum over all $j \neq i$ and noting that $\sum_{j \neq i} k_j \leqslant k$, we have that $\sum_{j \neq i} a'_{j,2} \geqslant \sum_{j \neq i} a_{j,2} - k$. Therefore, $a'_{i,2} \leqslant a_{i,2} + k$. Second, suppose that $D_2 \geqslant E$ and $D'_2 < E$. Then, by the definition of the T-Period mechanism, $a_{j,2} \leqslant \min(e_j + b_{j,2}, d_{j,2})$ for all $j \neq i$. Further

$$a'_{j,2} \geqslant \min(e_j + b'_j, d_{j,2}) = \min(e_j + b_{j,2} - k_j, d_{j,2}) \geqslant \min(e_j + b_{j,2}, d_{j,2}) - k_j \geqslant a_{j,2} - k_j$$

By the same argument as in the previous case, this implies that $a'_{i,2} \leqslant a_{i,2} + k$. Third, suppose that $D_2 < E$ and $D'_2 \geqslant E$. Then

$$a'_{i,2} \leqslant \min(e_i + b'_i, d_{i,2}) = \min(e_i + b_{i,2} + k_i, d_{i,2}) \leqslant \min(e_i + b_{i,2}, d_{i,2}) + k_i \leqslant a_{i,r} + k$$

Finally, suppose that $D_2 < E$ and $D'_2 < E$. First, suppose that $x' < x$. Then

$$a'_{i,2} = \min(e_i + b'_i, \max(d_{i,2}, x'e_i)) = \min(e_i + b_{i,2} + k_i, \max(d_{i,2}, x'e_i))$$

$$\leqslant \min(e_i + b_{i,2}, \max(d_{i,2}, x'e_i)) + k_i$$

$$\leqslant \min(e_i + b_{i,2}, \max(d_{i,2}, xe_i)) + k_i \leqslant a_{i,2} + k$$

114

Next, suppose that $x' \geqslant x$. Then for all $j \neq i$,

$$a'_{j,2} = \min(e_j + b'_j, \max(d_{j,2}, x'e_j)) = \min(e_j + b_{j,2} - k_j, \max(d_{j,2}, x'e_j))$$

$$\geqslant \min(e_j + b_{j,2}, \max(d_{j,2}, x'e_j)) - k_j$$

$$\geqslant \min(e_j + b_{j,2}, \max(d_{j,2}, xe_j)) - k_j = a_{j,2} - k_j$$

Again, this implies that $a'_{i,2} \leqslant a_{i,2} + k$.

In all cases, we have that $a_{i,2} \leqslant a'_{i,2} \leqslant a_{i,2} + k$. Therefore, $a'_{i,1} + a'_{i,2} \leqslant a_{i,1} + a_{i,2}$, which means that $a'_{i,3} \geqslant a_{i,3}$ and $a'_{i,4} \geqslant a_{i,4}$. Consider the difference in utility across all four rounds between the truthful and misreported instances.

$$U_{i,4}(\mathbf{a_i}) - U_{i,4}(\mathbf{a'_i}) = \sum_{r=1}^{4} \left(u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r}) \right)$$

$$= kH + \sum_{r=2}^{4} \left(u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r}) \right) \geqslant kH - kH = 0$$

The second transition is by Corollary 44, and the third transition because each $a'_{i,r} \geqslant a_{i,r}$ for all $r \in \{2, 3, 4\}$, $\sum_{r=2}^{4}(a'_{i,r} - a_{i,r}) = k$, and each resource can be worth at most H to agent i. ■

Given that the 1-P and 2-P mechanisms satisfy SP, it is easy to see that they satisfy SI also. By strategy-proofness, the utility that an agent gets from truthfully reporting her demands is at least the utility she gets from reporting $d'_{i,r} = e_i$ for all rounds r. Sharing incentives therefore follows as a corollary of the following proposition.

Proposition 51. *Under the T-Period mechanism, any agent that reports $d'_{i,r} = e_i$ for all rounds r receives $a_{i,r} = e_i$ for all rounds r.*

Corollary 52. *The T-Period mechanism satisfies SI for $T \leqslant 2$.*

One may hope to continue increasing flexibility, and therefore performance, by increasing the length of the 'borrowing' and 'payback' periods, potentially all the

way to having a single borrowing period of length $R/2$ and a single payback period of length $R/2$. Unfortunately, even for periods of length 3, strategy-proofness is violated.

Example 8. *Consider the 3-P mechanism. Suppose that $n = 5$ and $R = 6$. Each agent has endowment $e_i = 1$ (so each agent can borrow a total of three resources over the first three rounds, corresponding to the sum of their endowment across the final three rounds). Truthful demands are given by the following table.*

	$d_{i,1}$	$d_{i,2}$	$d_{i,3}$	$d_{i,4}$	$d_{i,5}$	$d_{i,6}$
$i = 1$	3	3	0	1	1	1
$i = 2$	0	3	3	1	1	1
$i = 3$	0	0	0	0	0	0
$i = 4$	0	0	0	0	0	0
$i = 5$	0	0	0	0	0	0

The corresponding allocations are given by:

	$a_{i,1}^{3-P}$	$a_{i,2}^{3-P}$	$a_{i,3}^{3-P}$	$a_{i,4}^{3-P}$	$a_{i,5}^{3-P}$	$a_{i,6}^{3-P}$
$i = 1$	3	2	0.75	0.08	0.08	0.08
$i = 2$	0.5	3	2	0.17	0.17	0.17
$i = 3$	0.5	0	0.75	1.58	1.58	1.58
$i = 4$	0.5	0	0.75	1.58	1.58	1.58
$i = 5$	0.5	0	0.75	1.58	1.58	1.58

Agent 1's utility is $5.25H + 0.75L$. If agent 1 misreports $d'_{1,1} = 2$, it can be checked that her allocations become 2, 2.5, 0.625, 0.292, 0.292, 0.292. Her utility is then $5.375H + 0.625L$, which is higher than her utility from reporting truthfully.

4.7 Evaluation

In this section, we evaluate different mechanisms using real and synthetic benchmarks. For real benchmarks, we use a Google cluster trace [1, 142], which data collected from a 12.5k-machine cluster over a month-long period in May 2011. All the machines in the cluster share a common cluster manager that allocates agent

tasks to machines.

Agents submit a set of resource demands for each task (e.g., required processors, memory, or disk space). Agent demands are normalized relative to the largest capacity of the resource on any machine in the traces. The cluster manager records any changes in the status of tasks (e.g., being evicted, failed, or killed) during their life cycle in a task event table. We use the task event table to track agents' demands for processors over time. Note that since all demands are scaled by the same factor, we safely use normalized demands as actual demands.

We divide time into 15 min intervals.[2] We define agents' demands for each interval to be the sum of their demands for all tasks they run in that interval. After processing the traces, we remove agents with constant demands or with average demand less than some marginal threshold. We assume that agents' endowments are equal to their average demands.

We observe that, for each agent, demands computed from Google traces have high correlations over time. An agent with high demand at 12am has typically high demand at 12:15am as well. In some deployment scenarios, demands may not be highly correlated. For example, when university cluster machines are allocated to professors and researchers on a daily basis, a researcher may have some jobs today, but may not want to use the cluster tomorrow.

To evaluate mechanisms in scenarios without correlated demands, we use synthetic benchmarks. We create random agent populations and random number of rounds. For each agent, we uniformly and randomly assign an endowment from 1 to 20. Once agents' endowments are set, we uniformly and randomly generate agent demands such that their average is equal to agents' endowments (i.e., $d_{i,r} \sim u[0, 2e_i]$)

Metrics. We report social welfare and Nash welfare, focusing on the number of

---

[2] We have created demands for varying time intervals. Since results do not change significantly for different interval lengths, we only include results on 15-min-long intervals.
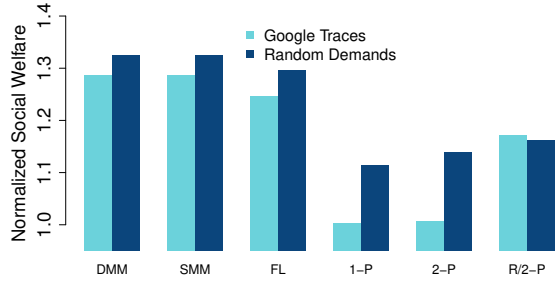
FIGURE 4.3: Social welfare achieved by different dynamic allocation mechanisms normalized to that of static allocations for Google cluster traces and 100 instances of random demands.

high-valued resources that each mechanism allocates. For social welfare, we report the following.

$$\text{Social Welfare} = \sum_i \sum_r \min(d_{i,r}, a_{i,r}).$$

Social welfare is a measure of efficiency but fails to distinguish between fair and unfair outcomes. For instance, suppose agent A with endowment 100 and agent B with endowment 1 both have demand 101. Allocating 100 units to agent A and 1 unit to agent B has the same social welfare as allocating 1 unit to A and 100 units to B. To distinguish between these two allocations, we also report the (weighted) Nash welfare as follows.

$$\text{Nash Welfare} = \sum_i e_i \log(\sum_r \min(d_{i,r}, a_{i,r})).$$

Observe that the Nash welfare metric is higher for the former scenario than the latter, which is in line with our intuition about which allocation is more fair.

4.7.1   Performance Evaluation

Figure 4.3 presents social welfare from varied allocation mechanisms for both Google and random traces normalized to social welfare of static allocations. DMM and
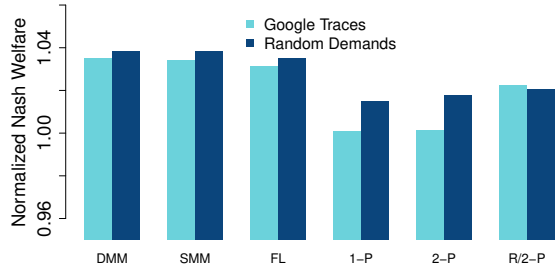
FIGURE 4.4: The Nash welfare achieved by different dynamic allocation mechanisms normalized to that of static allocations for Google cluster traces and 100 instances of random demands.
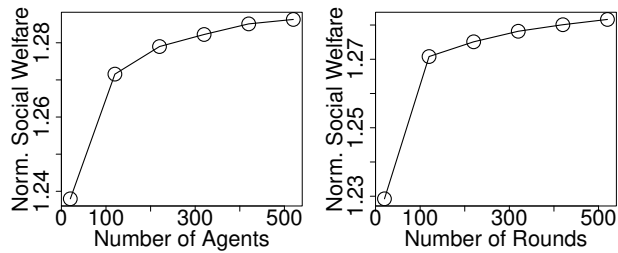
FIGURE 4.5: Social welfare of the flexible lending mechanism normalized to that of static allocations for varying agent population sizes and numbers of rounds. We fix the number of rounds to 50 when we vary the number of agents, and fix the number of agents to 50 when we vary the number of rounds.

SMM produce the same, highest social welfare as they always allocate resources to those agents with high valuations. Note that SMM and DMM both fail to guarantee strategy-proofness when $L > 0$. Therefore, when agents report strategically, for all the mechanism knows, SMM and DMM's allocations could be as inefficient as static allocations. But this is not captured in the figure, which implicitly assumes truthful reporting.

The 1-Period mechanism produces the lowest social welfare. Increasing the period length to 2 slightly improves the welfare of the T-Period mechanism. Note that both mechanisms outperform static allocations. The $R/2$-Period mechanism achieves 87% of SMM welfare for Google traces, but fails to provide strategy-proofness.

The social welfare of FL is competitive with state-of-the-art dynamic allocation
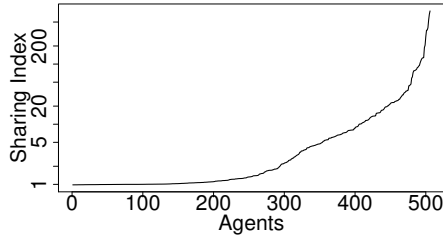
FIGURE 4.6: Sorted Sharing Index for Google Cluster Traces.

mechanisms. FL achieves 97% of SMM's welfare for Google traces and 98% for random demands. In practice, strong game-theoretic desiderata do not come with high welfare costs.

Figure 4.4 compares the normalized Nash welfare from varied mechanisms. Once again, DMM and SMM outperform other mechanisms, but DMM and SMM's outcomes are no longer equal because the number of high-valued resources that each agent receives differs across mechanisms. FL achieves 99.7% of DMM welfare for both Google cluster and random traces. This high Nash welfare could be explained by FL's high social welfare and the fact that FL allocates agents their exact endowment across rounds.

Figure 4.5 shows how social welfare changes when varying population size and number of rounds under FL. As the population size increases, the diversity between agents' demands at each round increases. Agents' complementary demands improve welfare from FL as fewer agents are forced to spend tokens on low-valued resources. Moreover, as the number of rounds increases, agents' flexibility in spending their tokens on high-valued resources increases. We prove in §4.5.5 that, at least when endowments are equal, FL approximates efficiency.

4.7.2 Sharing Incentives

We define the *sharing index* of agent i to be the ratio between the number of high-valued resources agent i receives under FL and under static allocations. In §4.5.4, we
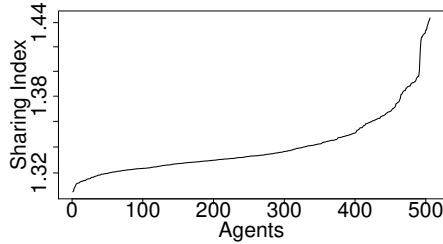
FIGURE 4.7: Sorted Sharing Index for a Random Trace.

show that FL guarantees that the sharing index of each agent is always at least 0.5. In practice, however, our simulations show that the sharing index is much higher.

Figure 4.6 shows the sharing index for all agents in the Google cluster traces, sorted in increasing order and shown on a log scale. The minimum sharing index across all agents is 0.98, and on average agents receive 15x more utility under FL compared to static allocations. As can be seen, there is high variance in sharing index across agents. Agents with high index are those who have zero demand at most of the rounds and very high demand at a few rounds. These agents benefit the most from sharing. When they have zero demand, they do not spend any tokens. Once they have a high demand they spend their tokens to receive the resources they need.

Figure 4.7 shows agents' sharing index for an instance with random demands. Since agents do not have correlated demands, the variance in sharing index is significantly lower compared to the Google cluster traces. Moreover, across all agents over 100 random instances, we do not observe a single violation of SI (i.e., no agent has a sharing index less than 1)

4.8   Related Work

There is a body of work in the *mechanism design without money* literature that is related to our work. Gorokh et al. [96] consider a setting where a single item is to be allocated repeatedly, and extend to more general settings in a follow-up paper

[97]. They do so by endowing each user with a fixed amount of artificial currency and then treating it similarly to if it were real money. They show that, for a large enough number of rounds, incentives to misreport and welfare loss both vanish. However, their notion of strategy-proofness is ex-ante Bayesian, requiring users (and the mechanism) to know the distribution from which other users' demands are drawn and truthful reporting is optimal only in expectation. Our notion of SP is ex-post, meaning that an agent never regrets truthful reporting.

Various other work does not explicitly use artificial currency, but by keeping track of how much utility an agent should receive in the future, achieve guarantees in a way that resembles the use of artificial currency [99, 20, 14]. Again, these results are for a weaker notion of SP.

In a similar setting, Aleksandrov et al. [17, 16] consider a stream of resources arriving one at a time that must be allocated among competing strategic agents. They consider two mechanisms, one of which is similar to SMM and the other similar to DMM. They obtain both positive and negative results for these mechanisms, however their positive results are primarily obtained for the case where agent utilities are 0 or 1, corresponding to our $L = 0$ case. They also consider only the symmetric agent setting, rather than our setting that allows unequal endowments.

There also exists literature on *dynamic fair division* [85, 159, 108], but this work predominantly focuses on agents arriving and departing over time, rather than the preferences themselves being dynamic, as in our work.

In the systems literature, in recent years, there has been a growing body of work on using economic game-theory to allocate resources [89, 169, 160]. These works only consider one-shot allocations and do not study allocations over time. Ghodsi et al. [90] consider dynamic allocations over time but in a completely different alloca-tion setting than ours. Their proposed mechanism allocates resources to packets in a queue. In such a setting, time cannot be divided into fixed intervals, because process-

ing packets take different times, which means a packet could stall all other packets until its processed. As a result, proportional allocations have to be approximated through discrete packet scheduling decisions [128, 67].

In a work that is close to our setting, Tang et al. [154] propose a dynamic allocation policy that resembles DMM. We study the characteristics of DMM in §4.3 and evaluate its performance in §4.7. Another related work in this area is that of Sandholm and Lai [146]. The authors propose a scheduler that allocates resources between users with dynamically changing demands. This work deploys heuristics and does not provide any theoretical guarantees that we study in this paper.

4.9 Discussion

We have considered the problem of designing mechanisms for dynamic proportional sharing in a high-low utility model that both incentivize users to participate and share their resources (sharing incentives), as well as truthfully report their resource requirements to the system (strategy-proofness). We show that while each of these properties is incompatible with full efficiency, it is possible to satisfy both of them and still obtain some efficiency gains from sharing.

The main mechanism that we present, the flexible lending mechanism, is strategy-proof and provides each user a theoretical guarantee of at least half her sharing incentives share. While we do not guarantee full sharing incentives, we show via simulations on both real and synthetic data that in practical situations, no users are significantly worse off by participating in the sharing scheme (and the majority are vastly better off). We show that under certain assumptions, the flexible lending mechanism provides full efficiency in the large round limit, which is supported by our simulation results. By incentivizing truthful reporting, we posit that the flexible lending mechanism will in fact produce significant efficiency gains in settings where agents are strategic.

Many directions for future work remain. The 2-Period mechanism fully satisfies both SP and SI, but remains very inflexible in its allocations. A key challenge is the design of a more flexible mechanism that satisfies both properties (or some upper bound on the efficiency that such mechanisms can achieve). Another direction is to extend the utility model. The high/low model is crucial to the positive strategic results that we obtain because trade-offs are well-defined: swapping an L resource for an H resource is always bad. Even introducing a medium (M) value complicates the situation considerably, and extending to such a setting would represent an exciting step forward.

5

Incentive Compatible and Efficient Wagering: The Double Clinching Auction

5.1   Introduction

In the second part of this thesis, we consider the problem of forecasting. A principal wants to gather information pertaining to the probability of some future event, but has no direct means to do so himself. Therefore, the information must be elicited from a group of self-interested agents. We first turn our attention to a class of elicitation mechanisms known as *wagering mechanisms.*

Wagering mechanisms allow a principal to elicit the beliefs of a group of agents without paying them directly or taking on any risk. Each agent specifies a *belief*, her own subjective estimate of the likelihood of a future event, such as the Democratic nominee winning the 2020 U.S. Presidential election. She also specifies a monetary budget or *wager*, the maximum amount that she is willing to lose. These wagers are then collected by the principal and, after the truth is revealed, redistributed to agents in such a way that agents with more accurate predictions are more highly rewarded. Meanwhile, since agents directly report their beliefs, the principal is able

to leverage the wisdom of crowds to obtain an accurate consensus forecast for the event, for example by computing an average [102], a budget-weighted average [31], a supra-Bayesian inference [121], or another aggregate measure of the forecasts [88].

Lambert et al. [112, 113] introduced the class of *weighted-score wagering mechanisms* (WSWMs), the unique wagering mechanisms to simultaneously satisfy a set of desirable properties including strict budget balance and incentive compatibility. Incentive compatibility is achieved through the use of *strictly proper scoring rules*, reward functions designed to incentivize truthful reports from risk-neutral agents. In particular, each agent's payoff under a WSWM is proportional to the difference between her own score and the budget-weighted score of the other agents. Chen et al. [56] later proposed the class of *no-arbitrage wagering mechanisms* (NAWMs), which are incentive compatible but only weakly budget balanced, allowing the principal to profit off of disagreement among agents. Under an NAWM, an agent's payoff is proportional to her score minus the score of the budget-weighted average belief. To our knowledge, these mechanisms and their derivatives (such as the randomized, private WSWM of Cummings et al. [63]) are the only known incentive compatible wagering mechanisms.

As an artifact of their use of proper scoring rules, these mechanisms have one undesirable property: In general, it is not possible for any agent to lose her full wager, even if all other agents are perfectly informed. In other words, the mechanisms are not *Pareto optimal*, in the sense that agents have significant budget left over even when additional trade would be mutually beneficial. This is a serious concern in practice since agents typically gravitate to venues where they have the opportunity for large gains. If these mechanisms yield badly suboptimal allocations, agents may question the rules or simply go elsewhere. Indeed, all widely deployed wagering mechanisms, including parimutuels, bookmakers, and double auctions, feature Pareto optimality. Additionally, wagers effectively lose their meaning as budgets. This has a surprising

implication on the quality of reports. Because an agent can never lose her full wager, she may be able to artificially inflate her reported budget risk-free. It turns out that when agents misreport their budgets, they can also have incentive to misreport their beliefs (see Example 10).

Motivated by this observation, we ask whether it is possible to design an incentive compatible wagering mechanism that achieves Pareto optimality without sacrificing other key properties. Unfortunately, the answer is no. We prove that no weakly incentive compatible wagering mechanism can achieve Pareto optimality along with individual rationality (meaning agents have incentive to participate) and weak budget balance. If the principal cannot force agents to participate and does not wish to subsidize the market, he must compromise on Pareto optimality. Given that, we seek an incentive compatible mechanism that is near-Pareto-optimal in practice.

Our mechanism is inspired by the observation that the output of a wagering mechanism has a natural interpretation as an allocation of securities. An agent who wins ρ_1 dollars if the Democrat is elected and loses ρ_0 dollars otherwise can equivalently be viewed as paying ρ_0 dollars up front for $\rho_0 + \rho_1$ shares of an Arrow-Debreu security worth \$1 if and only if the Democrat is elected. Thus wagering mechanisms can be viewed as allocating items (the securities) to agents, and it is natural to ask whether techniques from the auctions literature can be used. The *clinching auction* [21] produces VCG allocations and payments for multiple identical items, but VCG-style approaches cannot be applied when agents have budgets. Instead, we build on the *adaptive clinching auction* [68, 32], an extension of the clinching auction that incorporates budget constraints.

Our mechanism, the *double clinching auction* (DCA), is a two-sided version of the adaptive clinching auction. It elicits truthful reports by selling a variable number of securities to the agents via two simultaneous instances of the adaptive clinching auction, one which sells securities that pay off \$1 only if the event of interest happens

127

(`yes` securities), and one which sells securities that pay off \$1 only if it does not (`no` securities). The principal can always sell a pair of `yes` and `no` securities for \$1 or more without risk, since he will owe exactly \$1 to the agents regardless of the outcome. Our key technical contribution is determining the number of security pairs that the principal can sell via adaptive clinching auctions in such a way that he never loses money, without incentivizing agents to misreport their beliefs.

We also show that under the double clinching auction, each agent has at least some risk of losing her entire budget, making the budget declaration risky to inflate and restoring the semantics of the wager as the largest acceptable worst-case loss.

To evaluate the efficiency of the DCA, we run a series of simulations using thousands of probability judgments about hundreds of events, collected from an online forecasting contest called ProbabilitySports [87]. We compare the performance of the DCA with WSWMs, NAWMs, and the *parimutuel consensus mechanism* [72], which is Pareto optimal but not incentive compatible. Our simulations show that the DCA is indeed significantly closer to Pareto optimal than the other incentive compatible mechanisms, sometimes approaching the efficiency of the parimutuel consensus mechanism, which was specifically designed to maximize trade. Given the results, we are optimistic that the DCA can serve as a practical wagering mechanism that both satisfies agent demand and encourages honest revelations.

We follow previous authors [109, 104, 112, 56], assuming that agents have *immutable beliefs* that do not update during wagering. Our agents "agree to disagree", unlike Bayesian agents. While immutable beliefs and perfect Bayesian reasoning are both idealizations, the former is arguably closer to reality. In practice, overconfident opponents, each expecting to gain, trade all the time [105, 69], contradicting the no-trade theorems implied in the Bayesian setting. Other authors have explored incentive properties of wagering mechanisms with Bayesian [113] or boundedly rational [131] agents.

5.2 Wagering Mechanisms

Let X be a binary random variable or *event* with value or *outcome* in $\{0,1\}$. For example, imagine $X = 1$ is the outcome that the Democratic nominee wins the 2020 U.S. Presidential election and $X = 0$ is the outcome that he or she loses. We consider a setting in which a principal is interested in eliciting the beliefs of a set of agents $\mathcal{N}$ about the likelihood that $X = 1$. Following the line of work initiated by Lambert et al. [112], we assume that each agent $i \in \{1, \cdots, N\}$ has a private, subjective, immutable belief p_i about the probability that $X = 1$, and that agents are risk neutral up to some budget limitation. That is, each agent budgets for the largest loss that she is willing to tolerate, then maximizes her expected wealth subject to the budget constraint.

The principal operates a *wagering mechanism* in which each agent i submits a report $\hat{p}_i \in [0, 1]$, capturing her subjective belief about the likelihood that $X = 1$, and a wager $w_i \geqslant 0$, representing the maximum amount that she is prepared to lose. After observing the realized value of X, denoted x, the principal redistributes the agents' wagers, rewarding agents based on their wagers and the accuracy of their reports. We denote by $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}, x)$ the net payoff to agent i under reports $\hat{\mathbf{p}}$ and wagers $\mathbf{w}$ when $X = x$. For a wagering mechanism to be valid, it must be the case that no agent can lose more than her wager (i.e., for all i, $\hat{\mathbf{p}}$, $\mathbf{w}$, and x, we have $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x) \geqslant -w_i$) and an agent can choose not to participate by wagering 0 (i.e., $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x) = 0$ whenever $w_i = 0$). We can therefore let $\mathcal{N} = \mathbb{N}$ without loss of generality, since non-participation is equivalent to a zero wager. We denote by $\hat{\mathbf{p}}_{-\mathbf{i}}$ the predictions of all agents other than i and by $\mathbf{w}_{-\mathbf{i}}$ the wagers of all agents other than i.

There is a close connection between wagering mechanisms and *proper scoring rules* used to elicit truthful predictions from individual agents [147, 92]. A scoring rule s maps a prediction $p \in [0,1]$ and an outcome $x \in \{0,1\}$ to a score or reward in $\mathbb{R} \cup \{-\infty\}$. We say s is *proper* if for all $p, q \in [0,1]$, $ps(p,1) + (1-p)s(p,0) \geqslant ps(q,1) + (1-p)s(q,0)$, and *strictly proper* if this inequality is strict whenever $p \neq q$. An agent who is rewarded for her prediction using a proper scoring rule therefore maximizes her expected reward by reporting her true belief, uniquely if the scoring rule is strictly proper. A common example of a strictly proper scoring rule is the Brier score [40], $s(p,x) = 1 - (x-p)^2$.

For a wagering mechanism to elicit truthful reports about the likelihood of X, it must be the case that, fixing the wagers $\mathbf{w}$ and reports $\hat{\mathbf{p}}_{-\mathbf{i}}$ of other agents, agent i's payoff Π_i is a proper scoring rule. Building on this idea, Lambert et al. [112, 113] introduced the class of *weighted score wagering mechanisms (WSWMs)*. A WSWM has a payoff function of the form

$$\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; x) = w_i \left(s(\hat{p}_i, x) - \frac{\sum_{j \in \mathcal{N}} w_j s(\hat{p}_j, x)}{\sum_{j \in \mathcal{N}} w_j} \right) \tag{5.1}$$

where s is any strictly proper scoring rule bounded in $[0,1]$. WSWMs are the unique wagering mechanisms to simultaneously satisfy a set of desirable axioms that includes strict budget balance (the principal neither makes nor loses money), individual rationality (all agents have incentive to participate), strict incentive compatibility (agents have incentive to truthfully reveal their beliefs about X), anonymity (all agents are treated the same), sybilproofness (agents cannot profit by creating false identities), and a normality property (loosely, if agent i changes her report to improve her own expected payoff, the expected payoffs of other agents can't increase).

Chen et al. [56] pointed out that under a WSWM, it can be possible for an

agent to risklessly profit: there exist reports $\hat{\mathbf{p}}$ and wagers $\mathbf{w}$ such that for some agent i, both $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; 1)$ and $\Pi_i(\hat{\mathbf{p}}; \mathbf{w}; 0)$ are positive. They proposed an alternative class of incentive compatible mechanisms called *no-arbitrage wagering mechanisms (NAWMs)*, in which this extra profit is instead collected by the principal. The payoff to each agent is proportional to the difference between the score of his own prediction and the score of a type of weighted average of the other agents' predictions. We will return to these mechanisms later in the chapter.

5.2.2 Security Interpretation of Wagering Mechanisms

The output of a wagering mechanism has a natural interpretation as an allocation of Arrow-Debreu securities with payoffs that are contingent on the realization of X. We define a `yes` security to be a contract worth \$1 in the outcome $X = 1$ and \$0 if $X = 0$. Similarly, a `no` security is worth \$0 if $X = 1$ and \$1 if $X = 0$. A risk neutral agent with belief p about the likelihood that $X = 1$ would be willing to buy a `yes` security at any price up to p or a `no` security at any price up to $1 - p$. Since such trades reveal information about agents' beliefs, securities of this form are often considered in the context of prediction markets.

Suppose a wagering mechanism would yield a net payoff to agent i of $\rho_1 = \Pi_i(\hat{\mathbf{p}}; \mathbf{w}; 1)$ when $X = 1$ and $\rho_0 = \Pi_i(\hat{\mathbf{p}}; \mathbf{w}; 0)$ when $X = 0$. This is equivalent to the payoff that i would receive if she were sold $y_i = \max\{\rho_1 - \rho_0, 0\}$ `yes` securities and $n_i = \max\{\rho_0 - \rho_1, 0\}$ `no` securities for a total cost of $\sigma_i = \max\{-\rho_0, -\rho_1\}$. For example, if $\rho_0 < \rho_1$, then agent i's participation in the wagering mechanism is equivalent to agent i paying the principal $\sigma_i = -\rho_0$ before X is realized and then receiving $y_i = \rho_1 - \rho_0$ from the principal in the outcome $X = 1$.

Therefore, the output of a wagering mechanism can be completely specified by a triple $(\mathbf{y}, \mathbf{n}, \boldsymbol{\sigma})$, where for each agent i, $y_i \geqslant 0$ is the number of `yes` securities allocated to i, $n_i \geqslant 0$ is the number of `no` securities allocated to i, and σ_i is the

cost paid by i for these securities. To be a valid output, we require that for all i, either $y_i = 0$ or $n_i = 0$ (or both), and $\sigma_i \leqslant w_i$. This requirement is without loss of generality since any (fraction of a) pair of `yes` and `no` securities can be precisely converted into (a fraction of) \$1. We rely on the securities-based interpretation of wagering mechanisms for the remainder of this paper.[1]

5.2.3 Properties of Wagering Mechanisms

Lambert et al. [112] introduced several desirable properties for wagering mechanisms. We focus on three of these properties in our analysis: individual rationality, incentive compatibility, and budget balance. The definitions from Lambert et al. [112] are easily translated into our security-based representation. First, individual rationality requires that agents participate willingly; agents have nothing to lose (in expectation) by participating truthfully.

Definition 10. *A wagering mechanism is* individually rational *if, for any player i and any subjective probability p_i, there exists a report $\hat{p}_i$ such that for all $\hat{\mathbf{p}}_{-\mathbf{i}}, \mathbf{w}$,*

$$p_i y_i(\hat{\mathbf{p}}; \mathbf{w}) + (1 - p_i) n_i(\hat{\mathbf{p}}; \mathbf{w}) \geqslant \sigma_i(\hat{\mathbf{p}}; \mathbf{w}).$$

Incentive compatibility requires that each agent maximizes her expected payoff by reporting truthfully, regardless of the reports and wagers of other agents.

Definition 11. *A wagering mechanism is* weakly incentive compatible *if, for every agent i with belief p_i and all reports $\hat{\mathbf{p}}$ and wagers $\mathbf{w}$,*

$$p_i y_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}) + (1 - p_i) n_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}) - \sigma_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w})$$

$$\geqslant p_i y_i(\hat{\mathbf{p}}; \mathbf{w}) + (1 - p_i) n_i(\hat{\mathbf{p}}; \mathbf{w}) - \sigma_i(\hat{\mathbf{p}}; \mathbf{w}).$$

The mechanism satisfies strict incentive compatibility *if the inequality is strict whenever $p_i \neq \hat{p}_i$.*

---

[1] Note that in the case of WSWMs, the observation of Chen et al. [56] implies that it is possible to have $\sigma_i < 0$ for some i, meaning that i is allocated securities and actually *receives* money from the principal.

Finally, a wagering mechanism is budget balanced if the principal never loses.

Definition 12. *A wagering mechanism is* weakly budget balanced *if, for all* $\hat{\mathbf{p}}$ *and* $\mathbf{w}$,

$$\sum_{i \in \mathcal{N}} y_i(\hat{\mathbf{p}}; \mathbf{w}) \leqslant \sum_{i \in \mathcal{N}} \sigma_i(\hat{\mathbf{p}}; \mathbf{w}) \qquad and \qquad \sum_{i \in \mathcal{N}} n_i(\hat{\mathbf{p}}; \mathbf{w}) \leqslant \sum_{i \in \mathcal{N}} \sigma_i(\hat{\mathbf{p}}; \mathbf{w}).$$

The mechanism is strictly budget balanced *if the inequalities hold with equality for all* $\hat{\mathbf{p}}$ *and* $\mathbf{w}$.

5.3   A Tradeoff Between Efficiency and Incentive Compatibility

Our goal is to design a wagering mechanism not to maximize profit but to maximize the amount of useful and credible information gathered. In this context, both incentive compatibility and Pareto optimality (the formal definition is given in Section 5.3.2) are important. The former literally keeps agents honest, steering them to report their true best estimates and reassuring the principal that probabilities are not tainted by irrelevant strategic play. The latter keeps agents happy, earning them as much utility as possible without inexplicably leaving dollars on the table. Pareto optimality is standard in prediction markets, parimutuel markets, betting exchanges, and financial exchanges. A badly suboptimal allocation may confuse agents, discourage them from playing, or encourage them to inflate their budgets, as we shall see below, which may cause their probability reports to become untruthful too.

5.3.1   Inefficient Allocations and Budget Inflation

In this section, we consider the undesirable effects of Pareto inefficiency. We start with an example.

Example 9. *There are* $N = 4$ *agents with reported beliefs* $\hat{\mathbf{p}} = (0.1, 0.2, 0.5, 0.7)$ *and wagers* $\mathbf{w} = (1, 1, 1, 1)$. *Under the Brier scoring rule WSWM, the outcome is* $(\mathbf{y} = (0, 0, 0.25, 0.65), \mathbf{n} = (0.55, 0.35, 0, 0), \boldsymbol{\sigma} = (0.36, 0.19, 0.05, 0.29))$.

Observe that in Example 9, no agent stands to lose her full wager, regardless of the outcome. Indeed, the closest is agent 1, who risks losing 36% of her wager in the worst case. The *total risk*—the sum of all the agents' worst-case losses—is less than 25% of the total wagers. Thus WSWM is facilitating much less trade than if the agents were left to trade amongst themselves. Further, consider the 0.9 `yes` and 0.9 `no` securities allocated in Example 9. Thinking of these securities as any other commodity, we see that their allocation is not efficient. Some `yes` securities are allocated to agent 3 even though agent 4 has both a higher valuation and leftover budget.

The example above is in no way an edge case or specially manufactured; we will see in Section 4.7 that, if anything, it shows higher-than-average efficiency compared to our real-data simulations. Indeed, the following observation, which was originally made by Lambert et al. [112], shows that under a WSWM, agents who report any uncertainty will never lose their entire wager.

Proposition 53 (Lambert et al. 112). *For any weighted score wagering mechanism, for any $i \in \mathcal{N}$ and any reports $\hat{\mathbf{p}}$ and $\mathbf{w}$, if $\hat{p}_i \in (0, 1)$ and $w_i > 0$, then $\sigma_i(\hat{\mathbf{p}}, \mathbf{w}) < w_i$.*

This observation has an important implication that goes beyond the desire to facilitate as much trade as possible. Because an agent who reports her true budget w_i can never lose it all, she may be able to report a higher budget w_i' such that her maximum loss is still bounded by w_i. It turns out that when agents misreport their budgets, they may also have incentive to misreport their beliefs.

Example 10. *In Example 9, agent 4 derives utility $0.65 \cdot 0.7 - 0.29 = 0.17$ for being allocated 0.65 `yes` securities at a cost of 0.29, since she values each `yes` security at 0.7. However, since this does not exhaust her budget, she could inflate her budget to $w_4' = 2.04$ and instead be allocated 1.05 `yes` securities at a cost of 0.47, deriving utility 0.26. This budget inflation is completely safe in the sense that she never loses*

more than her budget, regardless of the reports and wagers of the other agents, (even if all other agents have arbitrarily large budgets and perfectly predict the outcome).

However, if agent 4 lowers her probability report to $\hat{p}_4 = 0.6$, she is able to inflate her budget even further. Intuitively, this is because 0.6 is a more moderate report than 0.7, so that even if $X = 0$, her loss will be lower. Agent 4 can safely report $w'_4 = 2.78$ along with $\hat{p}_4 = 0.6$ without any risk of spending more than her budget, regardless of the reports and wagers of the other agents. She is then allocated 0.96 yes securities at a cost of 0.38. Her expected utility is now $0.96 \cdot 0.7 - 0.38 = 0.30$, which is higher than she could safely obtain by truthfully reporting $\hat{p}_i = 0.7$.

5.3.2   Pareto Optimality

In this section, we define a natural notion of Pareto optimality for wagering mechanisms. For a fixed number of securities, a Pareto optimal allocation is, as usual, any locally optimal allocation that cannot be improved for one agent without harming others. However, the number of pairs of securities is not fixed: the principal or the agents can always manufacture more yes-no pairs at the cost of $1. Given this, we need a slightly expanded definition of Pareto optimality.

We say that a wagering mechanism is *Pareto optimal* if, treating agents' reports and wagers as their true beliefs and budgets, after all yes and no securities have been allocated and payments for these securities collected by the principal, there is no side bet that agents could make that would strictly benefit one without harming another, even if agents are allowed to create their own securities. We first define the notion of a profitable side bet.

Definition 13. *Given reports $\hat{\mathbf{p}}$, wagers $\mathbf{w}$, allocations $\mathbf{y}$ and $\mathbf{n}$ of yes and no securities, and payments $\boldsymbol{\sigma}$, a triple $(\Delta\mathbf{y}, \Delta\mathbf{n}, \Delta\boldsymbol{\sigma})$ is a* profitable side bet *if the following three conditions hold:*

1. *$\sum_{i \in \mathcal{N}} \Delta y_i = \sum_{i \in \mathcal{N}} \Delta n_i = \sum_{i \in \mathcal{N}} \Delta\sigma_i = 0.$*

135

2. *For all $i \in \mathcal{N}$, $\min\{y_i + \Delta y_i, n_i + \Delta n_i\} - (\sigma_i + \Delta \sigma_i) \geqslant -w_i$.*

3. *For all $i \in \mathcal{N}$, $\hat{p}_i \Delta y_i + (1 - \hat{p}_i) \Delta n_i \geqslant \Delta \sigma_i$, with strict inequality for at least one i.*

Let's examine this definition. The first condition ensures that $(\Delta \mathbf{y}, \Delta \mathbf{n}, \Delta \boldsymbol{\sigma})$ is a valid exchange among the agents, that is, all cash or securities given to one agent must come from other agents. The second condition ensures that no agent's budget is violated. The third guarantees that the exchange is profitable for at least one agent without harming any other agent (assuming truthful reports). We can now formally define Pareto optimality.

Definition 14. *A wagering mechanism is* Pareto optimal *if for all reports $\hat{\mathbf{p}}$ and wagers $\mathbf{w}$, the mechanism's output $(\mathbf{y}(\hat{\mathbf{p}}; \mathbf{w}), \mathbf{n}(\hat{\mathbf{p}}; \mathbf{w}), \boldsymbol{\sigma}(\hat{\mathbf{p}}; \mathbf{w}))$ is such that there exists no profitable side bet.*

This definition is difficult to work with directly. We show that there is an intuitive equivalent characterization of Pareto optimality in terms of allocations and costs: A mechanism is Pareto optimal if and only if there is some threshold price such that all agents with beliefs above the threshold spend their entire budget on `yes` securities while all agents with beliefs below the threshold spend their entire budget on `no` securities. This is formalized in the following theorem.

Theorem 54. *A wagering mechanism is Pareto optimal if and only if for all reports $\hat{\mathbf{p}}$ and $\mathbf{w}$, there exists an agent $j \in \mathcal{N}$ such that*

$$\forall i : \hat{p}_i < \hat{p}_j, \qquad \sigma_i(\hat{\mathbf{p}}; \mathbf{w}) = w_i \quad and \quad y_i(\hat{\mathbf{p}}; \mathbf{w}) = 0,$$

$$\forall i : \hat{p}_i > \hat{p}_j, \qquad \sigma_i(\hat{\mathbf{p}}; \mathbf{w}) = w_i \quad and \quad n_i(\hat{\mathbf{p}}; \mathbf{w}) = 0.$$

The first step of the proof, which appears in the appendix, is to show that any time a profitable side bet exists, there is a profitable side bet with $\Delta\sigma_i = 0$ for all agents i. This is because \$1 in cash is equivalent to a pair of `yes` and `no` securities. Thus we can limit attention to side bets that only involve the exchange of securities. The second step shows that, any time a profitable side bet exists, there exists a profitable side bet involving only two agents. The final step is to show that there is no profitable side bet between two agents if and only if the conditions in Theorem 54 hold.

Eisenberg and Gale [72] defined and analyzed the *parimutuel consensus mechanism* (PCM), a natural Pareto-optimal wagering mechanism. The outcome of the PCM is defined by a price π, such that all agents with $\hat{p}_i > \pi$ exhaust their entire wager buying `yes` securities at price π, and all agents with $\hat{p}_i < \pi$ exhaust their entire wager buying `no` securities at price $1 - \pi$. Any imbalance in demand for `yes` and `no` securities at price π is bridged by agents with report exactly π, who may buy either `yes` or `no` securities at the discretion of the mechanism. We can think of the PCM as a parimutuel mechanism with a proxy agent that switches agents' bets to the outcome most favorable to them, given the price. The PCM satisfies budget balance and individual rationality. However, the PCM does not satisfy incentive compatibility, because an agent may affect the price π in a way that is favorable to them. We explore the parimutuel consensus mechanism in more detail in Chapter 6, and provide in Example 12 an instance where the PCM violates incentive compatibility.

5.3.3 An Impossibility Result

We have shown that WSWM fails to produce Pareto optimal allocations and PCM fails to achieve incentive compatibility. In this section, we show that the tradeoff is unavoidable: no incentive compatible wagering mechanism can achieve Pareto optimality along with two other core properties.

The proof extends the intuition that for any two agents i and j with differing reports $\hat{p}_i < \hat{p}_j$, they must trade according to some intermediate price $p \in [\hat{p}_i, \hat{p}_j]$. It is therefore always in the interests of at least one of the agents to misreport her belief closer to that of the other agent, forcing the price further from her own true belief and thus achieving a higher payoff in expectation.

Theorem 55. *No wagering mechanism simultaneously satisfies individual rationality, weak incentive compatibility, weak budget balance, and Pareto optimality. This holds even if the number of agents is arbitrarily large and all agents wager the same amount of money. Any three of the four properties are simultaneously attainable.*

Proof. We first show the impossibility. Suppose there are $N \geqslant 2$ agents with beliefs **p** and identical wagers $\mathbf{w} = \mathbf{1}$. (It is trivial to extend the proof to $w_i = w$ for all i for any constant w, but complicates notation.) For simplicity, assume that all the p_i are unique.

Assume that we are running a mechanism that satisfies individual rationality, weak incentive compatibility, weak budget balance, and Pareto optimality. We will first show that for any such mechanism, if for all i, $p_i < 1/N$, then the Pareto optimality threshold $p = \max_i p_i$. Throughout the rest of this proof, let j denote the agent i with $p = p_i$.

By individual rationality and incentive compatibility, we have that for all i, $p_i y_i + (1 - p_i) n_i \geqslant \sigma_i$. By Pareto optimality, this implies that for all $i : p_i > p$, $p_i y_i \geqslant 1$, so $y_i \geqslant 1/p_i$. For the special agent j, if $y_j > 0$ then $y_j \geqslant \sigma_j/p_j$. By budget balance, we then have

$$\sum_{i:p_i>p} \frac{1}{p_i} + \mathbb{1}(y_j > 0)\frac{\sigma_j}{p_j} \leqslant \sum_{i=1}^{N} y_i \leqslant \sum_{i=1}^{N} \sigma_i = (N-1) + \sigma_j.$$

Suppose that it were not the case that $j = \arg\max_i p_i$. Then there is at least one

agent i with $p_i > p$, and so

$$\sum_{i:p_i>p} \frac{1}{p_i} + \mathbb{1}(y_j > 0)\frac{\sigma_j}{p_j} \geqslant \sum_{i:p_i>p} \frac{1}{p_i} > N \geqslant (N-1) + \sigma_j.$$

Since this inequality is strict, it contradicts the previous equation.

Now, consider the case in which for all $i < N$, $\hat{p}_i < 1/(N+2)$ and $\hat{p}_N = 1/(N+1)$. We have shown above that if all agents report truthfully then $j = N$. This means that for all $i \neq N$, $\sigma_i = 1$. Furthermore, for any $i \neq N$, this would still be the case even if i changed his report to any other value less than (but arbitrarily close to) $p_N = 1/(N+1)$. By incentive compatibility, such changes in report cannot change n_i, and so by individual rationality, it has to be the case that $(1 - p_N)n_i \geqslant 1$, so $n_i \geqslant 1/(1 - p_N)$ for all $i \neq N$. By budget balance,

$$\frac{N-1}{1-p_N} \leqslant \sum_{i=1}^{N} n_i \leqslant \sum_{i=1}^{N} \sigma_i \leqslant (N-1) + \sigma_N$$

and so

$$\sigma_N \geqslant (N-1)\left(\frac{1}{1-p_N} - 1\right) = \frac{N-1}{N}.$$

By individual rationality, budget balance, and this bound on σ_N, we must have

$$0 \leqslant \frac{1}{N+1}y_N - \sigma_N \leqslant \frac{1}{N+1}\left((N-1) + \sigma_N\right) - \sigma_N = \frac{N-1}{N+1} - \sigma_N\left(\frac{N}{N+1}\right) \leqslant 0.$$

This implies that $\sigma_N = (N-1)/N$, $y_N = N - 1 + \sigma_N$, and the expected utility of agent N is 0.

Suppose that agent N instead reported $p_N = 1/(N+2)$. We would still have $j = N$ since no other reports are as high. By the same argument we made above, it would have to be the case that

$$\sigma_N \geqslant (N-1)\left(\frac{1}{1-p_N} - 1\right) = \frac{N-1}{N+1}.$$

139

Again using a similar argument to the one above, by individual rationality, budget balance, and this bound on σ_N,

$$0 \leqslant \frac{1}{N+2} y_N - \sigma_N \leqslant \frac{1}{N+2} \left((N-1) + \sigma_N \right) - \sigma_N = \frac{N-1}{N+2} - \sigma_N \left(\frac{N+1}{N+2} \right) \leqslant 0.$$

This implies that $\sigma_N = (N-1)/(N+1)$, $y_N = N-1+\sigma_N$, and the expected utility of agent N is

$$\hat{p}_N y_N - \sigma_N = \frac{N-1}{(N+1)^2} > 0.$$

Therefore, agent N would prefer to deviate and the mechanism is not incentive compatible, a contradiction.

It remains to show that any three of the four properties are simultaneously attainable. The parimutuel consensus mechanism achieves individual rationality, weak budget balance, and Pareto optimality (see Chapter 6), and it is known that WSWMs satisfy (strong) incentive compatibility, individual rationality, and (strong) budget balance. To achieve weak incentive compatibility, weak budget balance, and Pareto optimality, we can simply take the entire wager from every agent (that is, let $y_i = n_i = 0$ and $\sigma_i = w_i$ for all i). Finally, to satisfy individual rationality, incentive compatibility and Pareto optimality, we can sell unlimited quantities of yes securities at a per-unit price p (fixed independently of the reports) and no securities at a per-unit price $1 - p$, so that all agents with report $\hat{p}_i \neq p$ fully exhaust their budget buying either yes or no securities. ∎

Individual rationality is hard to imagine giving up: We cannot force agents to participate. Weak incentive compatibility is key to ensuring the credibility of agents' reports. Although untruthful mechanisms like parimutuel wagering flourish in practice and do display an ability to aggregate useful information [18, 133], our goal is to create a mechanism that simplifies reasoning for the agents and principal and

that offers some modicum of assurance that the reports the principal is seeing are accurate to the best abilities of the agents. Some wagering mechanisms, in particular automated market maker algorithms for prediction markets [54], do give up budget balance, subsidizing trade as a reward for information. However, most mechanisms seek profits if anything, not losses. When a subsidy is not possible or desired, we must relax Pareto optimality. In the remainder of this paper, we present and analyze our *double clinching auction* wagering mechanism which maintains individual rationality, weak incentive compatibility, and weak budget balance, while coming close to Pareto optimality in practice.

5.4   The Adaptive Clinching Auction

Since wagering mechanisms can be interpreted as allocating items (securities) to agents, it is natural to ask whether techniques from the auctions literature might be useful. Ausubel's clinching auction [21] produces VCG allocations and payments in the setting in which there are multiple identical items and each agent has a fixed valuation per item. However, VCG-style approaches cannot be applied in our setting since agents have budgets. Instead, we build on the *adaptive clinching auction* of Dobzinski et al. [68], which extends Ausubel's auction to handle budget constraints.

In this section, we review the adaptive clinching auction and state some known results that are used in our analysis. Many details are necessarily omitted. For a full description, we point the reader to Dobzinski et al. [68] and, for the divisible-items version, Bhattacharya et al. [32]. In describing the auction, we use notation that parallels that of the wagering mechanism setting, but the general description in this section is for arbitrary items.

Suppose that there are m identical, indivisible items for sale to a set of agents $\mathcal{N}$. Each agent i has a private value p_i for each item and a budget w_i, which we assume is known to the auctioneer.

The adaptive clinching auction is an ascending price auction. Each agent $i \in \mathcal{N}$ reports a bid $\hat{p}_i$. The price p per item starts at 0 and grows over time. Items are allocated as the price increases. As this happens, the auctioneer keeps track of the number of items $q_i(p)$ that have been allocated to each agent i at prices less than p along with the total cost $c_i(p)$ of those items and the agent's remaining budget $B_i(p) = w_i - c_i(p)$. Define the demand of agent i at price p to be

$$
D_i(p) = \begin{cases} \infty & p = 0, \\ \left\lfloor \frac{B_i(p)}{p} \right\rfloor & 0 < p < \hat{p}_i, \\ 0 & p \geqslant \hat{p}_i \text{ and } p > 0. \end{cases} \tag{5.2}
$$

The adaptive clinching auction allocates items to agent i at price p if the total demand of the other agents falls below the total supply. In particular, let $q(p) = m - \sum_{i \in \mathcal{N}} q_i(p)$ be the total number of items yet to be allocated. At any point, if $D_{-i}(p) = \sum_{j \neq i} D_j(p) < q(p)$, then $q(p) - D_{-i}(p)$ items are allocated to (or "clinched by") agent i at a price of p per item, and the relevant variables are updated accordingly.

The auction ends when the total demand no longer exceeds the total supply, that is, when $\sum_{i \in \mathcal{N}} D_i(p) \leqslant q(p)$. At this point, the price stops ascending and all agents with $D_i(p) > 0$ are allocated their full demand at a per-item price of p. If the total demand at price p is strictly less than the supply (i.e., $\sum_{i \in \mathcal{N}} D_i(p) < q(p)$), then the remaining $q(p) - \sum_{i \in \mathcal{N}} D_i(p)$ items are allocated to agents i with $\hat{p}_i = p$. (We will see below that this is always possible to do.) A worked example is contained in the appendix.

The adaptive clinching auction can be extended to handle divisible items. While this extension is more complicated to write down, conceptually we simply view the auction as a continuous-time process. Bhattacharya et al. [32] give a formal description. We omit the details, but summarize the properties of the auction that we use to derive our results.[2]

---

[2] Both Dobzinski et al. [68] and Bhattacharya et al. [32] describe the divisible-items version in

First, agents have incentive to participate in the auction and to bid truthfully.

Lemma 56 (Dobzinski et al. [68])**.** *The adaptive clinching auction for divisible items is individually rational. When budgets are known to the auctioneer, it is also incentive compatible: Every agent i maximizes expected utility by reporting $\hat{p}_i = p_i$.*

While Dobzinski et al. [68] only state incentive compatibility for the case of indivisible items, their proof carries through for the continuous version, and this fact is used heavily by Bhattacharya et al. [32]. It follows from the observation that the report $\hat{p}_i$ only determines the price at which agent i drops out of the auction. While the price is below p_i, agent i can clinch (portions of) items at a per-item price below her value, thus deriving positive utility. After the price rises above p_i, any items she would clinch would cost more than her value, so she would derive negative utility. Thus, it is optimal to drop out of the auction exactly when the price reaches p_i.

We additionally use the fact that no agent is charged more than her budget.

Lemma 57 (Dobzinski et al. [68], Bhattacharya et al. [32])**.** *The adaptive clinching auction for divisible items never charges an agent more than her budget.*

We also rely heavily on the following facts, which together imply that no agent (or the auctioneer) can be made better off without harming another agent.

Lemma 58 (Dobzinski et al. [68], Bhattacharya et al. [32])**.** *The adaptive clinching auction for divisible items always allocates all m items.*

Lemma 59 (Dobzinski et al. [68], Bhattacharya et al. [32])**.** *If an agent receives a non-zero allocation of items from the adaptive clinching auction for divisible items, then any player with a higher bid exhausts her entire budget.*

terms of a *single* divisible item. For our purposes, it is more convenient to view it as an auction over some number m of divisible items. This is equivalent and simply requires a rescaling of agent values.

Finally, the utility of each agent is (weakly) increasing in the number of items sold.

Lemma 60 (Goel et al. [94])**.** *Fixing $\hat{\mathbf{p}}$ and $\mathbf{w}$, if $\hat{p}_i = p_i$ then i receives weakly greater expected utility from the adaptive clinching auction for divisible items when the number of items m increases.*

5.5   The Double Clinching Auction

In this section, we present the double clinching auction. Motivated by the observation that existing incentive compatible wagering mechanisms do not even allocate securities efficiently, we turn to the adaptive clinching auction as a way to efficiently allocate any fixed number of securities. The principal runs two instances of the adaptive clinching auction for divisible items, deriving the agents' bids from their reports. The first instance, which we refer to as the `yes` auction, sells some number m^* of `yes` securities to the agents, fixing the bid of each agent i to equal her report $\hat{p}_i$. The second instance, which we refer to as the `no` auction, sells m^* `no` securities, fixing the bid of agent i to $1 - \hat{p}_i$. If m^* is chosen such that the payment collected for each pair of `yes` and `no` securities is at least \$1, then the principal never loses money, that is, the mechanism is weakly budget balanced. While many values of m^* balance the budget, we define one particular value of m^*, carefully selected to ensure that agents cannot profit by misreporting their beliefs.

The primary technical contribution of this section is the derivation of m^* and the proof that the resulting auction is indeed (weakly) incentive compatible.

5.5.1   Definition of the Double Clinching Auction

To formally define the double clinching auction, we first describe the selection of m^*, the number of securities to be sold in each of the two instances of the adaptive clinching auction. We start by defining a pair of demand functions. These are similar

to Equation (5.2), but do not take into account items that may have been allocated. Let D_i^y be the demand of agent i for (arbitrarily divisible) yes securities at price p assuming a per-item value of $\hat{p}_i$, and D_i^n her demand for no securities at price p assuming a per-item value of $1 - \hat{p}_i$, that is,

$$
D_i^y(p) = \begin{cases} \infty & p = 0, \\ \frac{w_i}{p} & 0 < p < \hat{p}_i, \\ 0 & p \geqslant \hat{p}_i \text{ and } p > 0, \end{cases} \quad \text{and} \quad D_i^n(p) = \begin{cases} \infty & p = 0, \\ \frac{w_i}{p} & 0 < p < 1 - \hat{p}_i, \\ 0 & p \geqslant \hat{p}_i \text{ and } p > 0. \end{cases}
$$

Let $D^y(p) = \sum_{i \in \mathcal{N}} D_i^y(p)$ be the total demand of all agents for yes securities at price p, and $D_{-i}^y(p) = \sum_{j \neq i} D_j^y(p)$ be the total excluding agent i. Define $D^n(p)$ and $D_{-i}^n(p)$ similarly.

The double clinching auction allocates securities only when there are 4 or more agents with positive wagers. (Agents with wagers of zero can simply be dropped since this is equivalent to not participating.) If there are fewer than 4, then no trade occurs. For the remainder of this section, assume that there are $N \geqslant 4$ agents who submit reports $\hat{p}_1 \leqslant \hat{p}_2 \leqslant \ldots \leqslant \hat{p}_N$ and wagers $\mathbf{w} > \mathbf{0}$.

Fixing the number of securities m, define the *lowest clinching prices* as

$$
c_y(m) = \begin{cases} \inf\{p : \min_{i \in \mathcal{N}} D_{-i}^y(p) < m\} & m > 0, \\ \hat{p}_{N-1} & m = 0, \end{cases}
$$

and

$$
c_n(m) = \begin{cases} \inf\{p : \min_{i \in \mathcal{N}} D_{-i}^n(p) < m\} & m > 0, \\ 1 - \hat{p}_2 & m = 0. \end{cases}
$$

Here $c_y(m)$ can be thought of as the price at which the first (possibly infinitesimal) fraction of a security would be clinched in an adaptive clinching auction for m yes securities, and similarly, $c_n(m)$ the price at which the first fraction of a security would be clinched in an auction for m no securities. The $m = 0$ case is simply a

145

technical definition that is required in our proofs. Both $c_y(m)$ and $c_n(m)$ are well-defined since they each take the infimum of a non-empty set that is bounded below by 0 since $D^y_{-i}(0) = D^n_{-i}(0) = \infty$ for all agents i. It is easy to see that for all m, $c_y(m) \in (0, \hat{p}_{N-1}]$ and $c_n(m) \in (0, 1-\hat{p}_2]$. The following lemma gives additional useful properties of these functions. To show continuity, it is sufficient to show that the functions are surjective (onto), since a surjective, monotonic function is continuous.

Lemma 61. *Fixing reports* $\hat{\mathbf{p}}$ *and wagers* $\mathbf{w}$, c_y *and* c_n *are continuous and weakly decreasing.*

Let $M = \{m : c_y(m) + c_n(m) > 1\}$. For any $m \in M$, auctioning off m `yes` and m `no` securities via two adaptive clinching auctions is guaranteed to collect more than m dollars total, or more than \$1 for each pair, guaranteeing no loss for the principal. We set m^* to be the largest m in M: the most pairs of securities such that every pair, even every fraction of a pair, costs more than \$1 per share (i.e., every ε shares cost more than \$$\varepsilon$). Formally, the number of pairs of securities auctioned is

$$m^* = \begin{cases} \sup M & \hat{p}_2 < \hat{p}_{N-1}, \\ 0 & \hat{p}_2 = \hat{p}_{N-1}. \end{cases} \tag{5.3}$$

The following lemma guarantees that m^* is well-defined. This is clearly the case when $\hat{p}_2 = \hat{p}_{N-1}$. To show that m^* is well-defined when $\hat{p}_2 < \hat{p}_{N-1}$, it is sufficient to show that the set M is non-empty and bounded above, which implies the existence of a unique least upper bound. To show that $m^* > 0$ when $\hat{p}_{N-1} > \hat{p}_2$, we argue that $c_y(0) + c_n(0) > 1$, which implies there must exist some $m' > 0$ such that $c_y(m') + c_n(m') > 1$. This in turn implies that $m' \in M$ and therefore, $m^* = \sup M \geqslant m' > 0$.

Lemma 62. *For any* $\hat{\mathbf{p}}$ *and* $\mathbf{w}$, m^* *is well-defined. Furthermore,* $m^* > 0$ *when* $\hat{p}_{N-1} > \hat{p}_2$.

With these definitions in place, we can formally define the double clinching auction; see Algorithm 8. The principal first sets m^* according to Equation 5.3. He then runs an auction for m^* yes securities (the yes auction) and an auction for m^* no securities (the no auction). A worked example of the double clinching auction on the reports from Example 9 is given in the appendix.

We have already shown that this procedure is well defined. However, to show that the double clinching auction is a valid wagering mechanism, we must also show that no agent ever loses more than her wager; that is, for any $\hat{\mathbf{p}}$ and $\mathbf{w}$, the double clinching auction produces output $(\mathbf{y}, \mathbf{n}, \boldsymbol{\sigma})$ such that for all $i \in \mathcal{N}$, $\min\{y_i, n_i\} = 0$ and $\sigma_i \leqslant w_i$. We show this in the following theorem.

ALGORITHM 8: The Double Clinching Auction. Here $ClinchingAuction(m, \hat{\mathbf{p}}, \mathbf{w})$ denotes the allocation and payments produced by an adaptive clinching auction for m arbitrarily divisible items on bids $\hat{\mathbf{p}}$ and budgets $\mathbf{w}$.

Input: Reports $\hat{\mathbf{p}}$ and wagers $\mathbf{w} > \mathbf{0}$ of N agents
1 **if** $N < 4$ or $\hat{p}_2 = \hat{p}_{N-1}$ **then**
2   Set $(\mathbf{y}, \mathbf{n}, \boldsymbol{\sigma}) = (\mathbf{0}, \mathbf{0}, \mathbf{0})$;
3 **else**
4   Set m^* as in Equation 5.3;
5   Let $(\mathbf{y}, \boldsymbol{\sigma}_y) = ClinchingAuction(m^*, \hat{\mathbf{p}}, \mathbf{w})$;
6   Let $(\mathbf{n}, \boldsymbol{\sigma}_n) = ClinchingAuction(m^*, \mathbf{1} - \hat{\mathbf{p}}, \mathbf{w})$;
7   Let $\boldsymbol{\sigma} = \boldsymbol{\sigma}_y + \boldsymbol{\sigma}_n$;
8 **end**
Output: $(\mathbf{y}, \mathbf{n}, \boldsymbol{\sigma})$

Theorem 63. *The double clinching auction is a valid wagering mechanism.*

From Lemma 57, we know that no agent can lose more than her wager in either the yes auction or the no auction alone. It is therefore sufficient to show that no agent is ever allocated a positive number of securities in both auctions. This follows immediately from the following lemma, taking p to be the report $\hat{p}_i$ of any agent, and the definition of the clinching auction.

Lemma 64. *Fixing any reports $\hat{\mathbf{p}}$ and wagers $\mathbf{w}$, for any $p \in [0, 1]$, either $\min_{i \in \mathcal{N}} D^y_{-i}(p) \geqslant m^*$, $\min_{i \in \mathcal{N}} D^n_{-i}(1 - p) \geqslant m^*$, or both.*

Proof. If $m^* = 0$ then this claim is trivially true, since for all p, $\min_{i \in \mathcal{N}} D^y_{-i}(p) \geqslant 0$ and $\min_{i \in \mathcal{N}} D^n_{-i}(p) \geqslant 0$. So suppose that $m^* > 0$. Suppose that $\min_{i \in \mathcal{N}} D^y_{-i}(p) < m^*$ and $\min_{i \in \mathcal{N}} D^n_{-i}(1-p) < m^*$. Then there exists an $m' < m^*$ such that $\min_{i \in \mathcal{N}} D^y_{-i}(p) < m'$ and $\min_{i \in \mathcal{N}} D^n_{-i}(1 - p) < m'$. Therefore, when m' securities are sold, clinching in the **yes** auction begins at (or before) p, and clinching in the **no** auction begins at (or before) $1-p$. That is, $c_y(m') \leqslant p$ and $c_n(m') \leqslant 1-p$. So $c_y(m')+c_n(m') \leqslant p+1-p = 1$. This implies that m' is a lower upper bound on the set $\{m : c_y(m) + c_n(m) > 1\}$ than m^* is, violating the definition of m^*. ∎

5.5.2 Properties of the Double Clinching Auction

In this section, we discuss some desirable properties of the double clinching auction. We first observe that the double clinching auction is weakly budget balanced and individually rational.

Proposition 65. *The double clinching auction is weakly budget balanced and individually rational.*

Proof. We first prove budget balance. From Lemma 58, we know that all m^* securities are allocated in both the **yes** auction and the **no** auction. Since all **yes** securities are bought for a per-unit price of at least $c_y(m^*)$ and all **no** securities are bought at a per-unit price of at least $c_n(m^*)$, the principal collects payments of at least $m^*(c_y(m^*) + c_n(m^*))$ which equals m^* by Lemma 66. For each pair of **yes** and **no** securities sold, the principal pays out exactly \$1 to the agents, regardless of the outcome. Therefore the principal is guaranteed to collect more than he pays out.

We next prove individual rationality. In particular we show that truthful reporting leads to non-negative expected payoff. From Lemma 56, agents obtain non-

negative utility from truthfully reporting $\hat{p}_i$ and $1 - \hat{p}_i$ in each of the two clinching auctions, regardless of the value of m^*. Since participating truthfully in the double clinching auction is equivalent to participating truthfully in each of the two clinching auctions individually, each agent derives non-negative utility for doing so. So the double clinching auction is individually rational. ∎

Lemma 66. *For any reports* $\hat{\mathbf{p}}$ *and wagers* $\mathbf{w}$, $c_y(m^*) + c_n(m^*) = 1$.

Finally we state our main theoretical result: incentive compatibility of the double clinching auction. The proof is significantly more involved and we develop it in the next subsection.

Theorem 67. *The double clinching auction is weakly incentive compatible.*

5.5.3   Proof of Incentive Compatibility

In this section, we prove Theorem 67, beginning with some useful lemmas. The first states that an agent cannot benefit from misreporting her belief unless it increases the number of securities.

Lemma 68. *For any* $i \in \mathcal{N}$, *fix the wagers* $\mathbf{w}$ *of all agents and reports* $\hat{\mathbf{p}}_{-i}$ *of all agents but* i, *and let* $\hat{p}_i = p_i$. *Agent* i *cannot increase her expected utility under the double clinching auction by reporting any* $\hat{p}_i' \neq p_i$ *unless this report increases the value of* m^*.

Proof. Let m^* denote the number of security pairs allocated by the double clinching auction when i reports $\hat{p}_i = p_i$, and $\hat{m}^*$ the number when i reports $\hat{p}_i'$.

First, observe that agent i cannot benefit from any misreport for which $m^* = \hat{m}^*$. This follows immediately from the incentive compatibility of the adaptive clinching auction (Lemma 56). Agent i maximizes the utility she gains from both the `yes` and `no` auctions individually when her bids in these auctions are truthful. Fixing m^*, the

149

yes and no auctions are run independently, so agent i maximizes her total utility by reporting her true belief. Next, suppose that $\hat{m}^* < m^*$. Agent i's utility for bidding untruthfully for $\hat{m}^*$ securities is weakly less than her utility for bidding truthfully for $\hat{m}^*$ securities, by incentive compatibility of the adaptive clinching auction, which is weakly less than her utility for bidding truthfully for m^* securities, by Lemma 60. Thus she would weakly prefer to bid truthfully than to make any misreport that reduces m^*. ∎

Lemma 69 follows because, when any agent increases her report, the demand for yes (respectively, no) securities at a fixed price does not decrease (respectively, increase).

Lemma 69. *For any $i \in \mathcal{N}$, fix the wagers $\mathbf{w}$ of all agents and reports $\hat{\mathbf{p}}_{-i}$ of all agents but i. For any fixed m, as i's report $\hat{p}_i$ increases, the lowest clinching price $c_y(m)$ weakly increases, while the lowest clinching price $c_n(m)$ weakly decreases.*

The proof of Lemma 70 uses the fact that as the price moves from any value p to a sufficiently close higher value p', no new agent will drop out of the auction, and so demand functions only change by a very small amount.

Lemma 70. *For any reports $\hat{\mathbf{p}}$ and wagers $\mathbf{w}$, and any m, $\min_{i \in \mathcal{N}} D^y_{-i}(c_y(m)) \leqslant m$ and $\min_{i \in \mathcal{N}} D^n_{-i}(c_n(m)) \leqslant m$.*

We are now ready to complete the proof of Theorem 67. We start by observing that no agent can be allocated both yes *and* no securities. We treat two cases. If an agent's misreport does not change the type of security that she is allocated, then it cannot increase the number of securities sold, and so by Lemma 68, cannot be profitable. If her misreport does change the type of security that she is allocated, she may be able to increase the number of securities auctioned. However, in this case, the amount she pays would be higher than her value for the securities she gets.

150

Proof of Theorem 67. Consider an agent i with belief p_i and let $\hat{\mathbf{p}}_{-\mathbf{i}}$ denote the reports of all agents other than i. Let m^* be the number of pairs of securities auctioned if i truthfully reports p_i. Denote by $\hat{c}_y$ and $\hat{c}_n$ the lowest clinching price functions if i misreports $\hat{p}_i \neq p_i$, and by $\hat{D}$ the demand functions in the misreported instance. Let $\hat{m}^*$ denote the number of pairs of securities auctioned in the misreported instance.

Noting that $\hat{p}_i \neq p_i$, we can break the analysis into four cases:

(1) $p_i < c_y(m^*)$ and $\hat{p}_i \leqslant c_y(m^*)$ (3) $p_i \leqslant c_y(m^*)$ and $\hat{p}_i > c_y(m^*)$

(2) $p_i > c_y(m^*)$ and $\hat{p}_i \geqslant c_y(m^*)$ (4) $p_i \geqslant c_y(m^*)$ and $\hat{p}_i < c_y(m^*)$

Case 1 and Case 2 are symmetric, since in Case 2 $1-p_i < c_n(m^*)$ and $1-\hat{p}_i \leqslant c_n(m^*)$, which is equivalent to Case 1 reversing the outcomes `yes` and `no`. Similarly, Case 3 and Case 4 are symmetric. Therefore, it is sufficient to show that i does not benefit from misreporting in Cases 1 or 3.

Case 1: $p_i < c_y(m^*)$ and $\hat{p}_i \leqslant c_y(m^*)$. To show that i can not benefit from this misreport, we prove that she does not change the clinching prices $c_y(m^*)$ and $c_n(m^*)$. We will show that if $\hat{p}_i < c_y(m^*)$ then this is true because the global demand can only change at prices between p_i and $\hat{p}_i$, and this interval does not contain $c_y(m^*)$. When $\hat{p}_i = c_y(m^*)$, some more care is necessary.

If $m^* = 0$ then $c_y(m^*) = p_{N-1}$. Further, we know that $p_{N-1} = p_2$, or else it would be the case that $m^* > 0$, by Lemma 62. And, since we have assumed that $p_i < c_y(m^*)$, we know that i's report is the lowest of all agents. Since $\hat{p}_i \leqslant c_y(m^*) = p_{N-1} = p_2$, $\hat{p}_i$ is still the (equal) lowest report, and therefore both the second highest and second lowest reports are unchanged. In particular, $\hat{p}_2 = \hat{p}_{N-1}$, so $\hat{m}^* = 0$. By Lemma 68, this misreport does not benefit i.

Now suppose that $m^* > 0$. We first show that $c_y(m^*) = \hat{c}_y(m^*)$ and $c_n(m^*) = \hat{c}_n(m^*)$. If $\hat{p}_i < c_y(m^*)$ then the demand locally around $c_y(m^*)$ and $c_n(m^*)$ is unchanged. Therefore, since $c_y(m^*)$ and $c_n(m^*)$ are the prices at which demand

151

drops below m^*, these quantities remain unchanged in the misreported instance. If $\hat{p}_i = c_y(m^*)$ then, by Lemma 69, $c_y(m^*) \leqslant \hat{c}_y(m^*)$, since $p_i < \hat{p}_i = c_y(m^*)$. However, for all $p > c_y(m^*)$, the demand in the misreported instance is exactly the same as that in the truthful instance, and therefore $\min_i \hat{D}^y_{-i}(p) < m^*$ for all $p > c_y(m^*)$, which implies that $\hat{c}_y(m^*) \leqslant c_y(m^*)$. This, together with the earlier statement that $\hat{c}_y(m^*) \geqslant c_y(m^*)$, gives us $c_y(m^*) = \hat{c}_y(m^*)$. By similar reasoning, $c_n(m^*) = \hat{c}_n(m^*)$.

Therefore $\hat{c}_y(m^*) + \hat{c}_n(m^*) = c_y(m^*) + c_n(m^*) = 1$. Since $\hat{c}_y$ and $\hat{c}_n$ are decreasing functions, m^* is an upper bound on the set $M = \{m : \hat{c}_y(m) + \hat{c}_n(m) > 1\}$. Since the double clinching auction sells a number of securities equal to the least upper bound of M, it therefore sells at most m^* securities in the misreported instance. By Corollary 68, agent i does not profit from this misreport.

Case 3: $p_i \leqslant c_y(m^*)$ and $\hat{p}_i > c_y(m^*)$. In this case, i's misreport can increase the number of securities sold. However, we show that to do so, i must be allocated some **yes** securities. But since i's misreport is higher than her true value, it must also be the case that the price for **yes** securities is higher in the misreported instance than the truthful instance. Because all **yes** securities are sold at a price higher than i's valuation in the truthful instance, it must still be the case in the misreported instance. Therefore i does not get any positive utility from the securities she is allocated.

There are two possibilities. First is that $\hat{c}_y(m^*) + \hat{c}_n(m^*) \leqslant 1$, in which case we need to sell (weakly) fewer securities in the misreported instance than the truthful instance. That is, $\hat{m}^* \leqslant m^*$. By Lemma 68, the misreport can not be profitable for i in this case.

Second is that $\hat{c}_y(m^*) + \hat{c}_n(m^*) > 1$. In this case it may be possible to sell more securities, so assume that $\hat{m}^* > m^*$ (otherwise i's misreport is not profitable, by

Lemma 68). By Lemma 70, $\min_j \hat{D}^y_{-j}(\hat{c}_y(m^*)) \leqslant m^* < \hat{m}^*$. So, by Lemma 64,

$$\min_j \hat{D}^n_{-j}(1 - \hat{c}_y(m^*)) \geqslant \hat{m}^*. \tag{5.4}$$

In what remains of the proof, we show that holding the number of securities the same as in the truthful instance, i's misreport cannot result in the clinching price rising all the way *above* $\hat{p}_i$. We can then use the fact that the clinching price decreases as we sell more securities to deduce that $\hat{p}_i \geqslant \hat{c}_y(\hat{m}^*)$, which (after addressing some details) says that i is allocated yes securities, and not no securities. By lower bounding the price of the yes securities by i's true valuation p_i, this says that i can not derive positive utility from this misreport. We now prove this formally.

We first show that $\hat{p}_i \geqslant \hat{c}_y(m^*)$. In the case that $m^* = 0$, this is true because $\hat{p}_i > c_y(m^*) = p_{N-1}$, so therefore $\hat{p}_i$ is one of the two highest reports in the misreported instance. And, since $m^* = 0$, it follows that $\hat{c}_y(m^*) = \hat{p}_{N-1} \leqslant \hat{p}_i$.

For the case that $m^* > 0$, note that the demand is unchanged from the truthful instance at all prices greater than or equal to $\hat{p}_i$. Therefore for all $p \geqslant \hat{p}_i$, we have that $\min_j \hat{D}^y_{-j}(p) = \min_j D^y_{-j}(p) < m^*$, where the inequality holds because $p \geqslant \hat{p}_i > c_y(m^*)$. In particular, $\min_j \hat{D}^y_{-j}(\hat{p}_i) < m^*$, which implies that $\hat{p}_i \geqslant \hat{c}_y(m^*)$.

From $\hat{p}_i \geqslant \hat{c}_y(m^*)$, we have that $1 - \hat{p}_i \leqslant 1 - \hat{c}_y(m^*)$, which implies that $\hat{D}^n_{-i}(1 - \hat{p}_i) \geqslant \hat{D}^n_{-i}(1 - \hat{c}_y(m^*))$. Combining this with Equation 5.4, $\hat{D}^n_{-i}(1 - \hat{p}_i) \geqslant \hat{D}^n_{-i}(1 - \hat{c}_y(m^*)) \geqslant \min_j \hat{D}^n_{-j}(1 - \hat{c}_y(m^*)) \geqslant \hat{m}^*$, which implies that i does not receive no securities in the misreported instance. There are two possibilities remaining. If i also does not receive yes securities, then agent i achieves zero overall payoff after misreporting, which is no better than her payoff from reporting truthfully. Otherwise, the average price paid per yes security is at least $\hat{c}_y(\hat{m}^*) = 1 - \hat{c}_n(\hat{m}^*) \geqslant 1 - \hat{c}_n(m^*) \geqslant 1 - c_n(m^*) = c_y(m^*) \geqslant p_i$, where the first inequality follows from the fact that $\hat{c}_n$ is decreasing and $\hat{m}^* > m^*$, and the second inequality follows from Lemma 69, because

153

$\hat{p}_i > p_i$. Therefore i is paying a price for the securities equal to or greater than they are worth to her, so she obtains non-positive expected payoff, which is no better than her (non-negative) truthful payoff. ∎

5.5.4 Beyond Weak Incentive Compatibility

Theorem 67 proves weak incentive compatibility. Taken at face value, weak incentive compatibility is, well, extremely weak. Indeed, simply paying each agent a constant amount regardless of her report satisfies weak incentive compatibility.

We show that the double clinching auction actually satisfies a stronger property: If agent i makes any misreport $\hat{p}_i \neq p_i$, then, for some set of reports $\hat{\mathbf{p}}_{-i}$ of the other agents, agent i obtains strictly lower expected utility than she would by reporting truthfully. If agent i is sufficiently uncertain about other agents, she is strictly better off reporting her true belief.

Theorem 71. *Fix any set of agents $\mathcal{N}$ with $N \geq 4$ and any wagers $\mathbf{w}$. For any agent i with belief p_i and any report $\hat{p}_i \neq p_i$, there exist reports $\hat{\mathbf{p}}_{-i}$ of the other agents such that under the double clinching auction*

$$p_i y_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}) + (1 - p_i)n_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}) - \sigma_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w})$$

$$> p_i y_i(\hat{\mathbf{p}}; \mathbf{w}) + (1 - p_i)n_i(\hat{\mathbf{p}}; \mathbf{w}) - \sigma_i(\hat{\mathbf{p}}; \mathbf{w}).$$

Proof. Suppose that $\hat{p}_i > p_i$; because of the symmetries in the double clinching auction, this is without loss of generality. For all $j \neq i$, let $\hat{p}_j$ lie in $(p_i, \hat{p}_i)$ and assume there are at least three unique reports from the agents $j \neq i$. This guarantees that $m^* > 0$ by Lemma 62. Since $\hat{p}_i$ is the largest report, $y_i(\hat{\mathbf{p}}, \mathbf{w}) > 0$; this follows from Lemma 59. Furthermore, $c_y(m^*) > p_i$, so the price i pays per `yes` security must be

strictly greater than p_i, so $\sigma_i(\hat{\mathbf{p}}; \mathbf{w}) < p_i y_i(\hat{\mathbf{p}}; \mathbf{w})$. Thus,

$$p_i y_i(\hat{\mathbf{p}}; \mathbf{w}) + (1 - p_i) n_i(\hat{\mathbf{p}}; \mathbf{w}) - \sigma_i(\hat{\mathbf{p}}; \mathbf{w}) = p_i y_i(\hat{\mathbf{p}}; \mathbf{w}) - \sigma_i(\hat{\mathbf{p}}; \mathbf{w})$$

$$< 0 \leqslant p_i y_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}) + (1 - p_i) n_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}) - \sigma_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}),$$

where the final inequality follows from the fact that the double clinching auction is incentive compatible and individually rational. ∎

5.5.5 Budget Inflation Under the Double Clinching Auction

As discussed in Section 5.3, even a wagering mechanism that satisfies incentive compatibility may give an agent incentive to misreport her belief if she can safely inflate her budget. Since the double clinching auction is not Pareto optimal, a bidder with complete knowledge of the reports and wagers of other agents could have incentive to inflate her budget. Further, there exist examples analogous to Example 10 where the potential for an agent to inflate her budget may also affect her incentive to report truthfully.

However, in reality agents operate with only limited knowledge about the reports of other agents. While the budget misreport in Example 10 was safe in the sense that the budget inflation could not lead to the misreporting agent overspending her true budget, we can show that completely safe manipulations are not possible under the double clinching auction. An agent cannot inflate her wager without at least some risk of losing more than her true budget. This is in stark contrast to Proposition 53 for the WSWM.

Theorem 72. *Fix any set of agents $\mathcal{N}$ with $N \geqslant 4$. For any agent i with report $\hat{p}_i$ and wager w_i, there exist reports $\hat{\mathbf{p}}_{-i}$ and wagers $\mathbf{w}_{-i}$ of the other agents such that $\sigma_i = w_i$ under the double clinching auction.*

Proof. Suppose without loss of generality that $\hat{p}_i \geqslant \frac{1}{2}$. Let $\hat{\mathbf{p}}_{-\mathbf{i}} = (p_1, p_2, \ldots, p_{N-2}, p_{N-1}) = (0.1, 0.1, \ldots, 0.1, 0.45)$ and $\mathbf{w}_{-\mathbf{i}} = (2w_i, 2w_i, \ldots, 2w_i, w_i)$.

Consider the allocation of a double clinching auction with $m = \frac{w_i}{0.4}$. Then $c_y(m) = \frac{w_i}{w_i/0.4} = 0.4$, so both agents $N-1$ and $i = N$ begin clinching at $p = 0.4$, and $c_n(m) \geqslant \frac{2w_i}{w_i/0.4} = 0.8$. In particular, since $c_y(m) + c_n(m) > 1$, it must be the case that $m^* > m$ for this instance. Since agent $N-1$ is allocated some non-zero number of yes securities when m pairs of securities are allocated via clinching auction, she is also allocated non-zero yes securities when m^* pairs of securities are auctioned. By Lemma 59, it must be the case that agent $\sigma_i = w_i$. ∎

5.6   Simulations

For a fixed number of yes securities, the adaptive clinching auction is efficient, so we had reason to suspect that the double clinching auction, selling m^* yes and no securities, would be near efficient. In this section, in a series of simulations based on real probability reports, we show that indeed the DCA is much more efficient than the WSWM or the NAWM, in some cases coming remarkably close to Pareto optimality.

We compare the performance of the double clinching auction to the parimutuel consensus mechanism (PCM), the Brier scoring rule version of the weighted score wagering mechanism (WSWM), and the Brier-score no-arbitrage wagering mechanism (NAWM). The PCM is known to be Pareto optimal, serving as the gold standard with respect to the amount of trade generated, though is not incentive compatible. WSWMs and NAWMs provide a natural comparison as the only other known, non-trivial wagering mechanisms that are individually rational, incentive compatible, and budget balanced. We chose the Brier scoring rule since it is commonly used in practice.

We tested each wagering mechanism on a data set consisting of probability reports

collected from an online prediction contest called ProbabilitySports [87].[3,4] The data set consists of probabilistic predictions about the outcomes of 1643 U.S. National Football League games between the start of the 2000 preseason and the end of the 2004 season. For each match, between 64 and 1574 people reported their subjective probability of the home team winning the game. After each game, they earned points in the contest according to a Brier scoring rule.

The ProbabilitySports users provided probabilities but not wager amounts. We simulate wagers in two ways. First, we generate uniform wagers: we fix all wagers at 1, modeling a scenario where agents are equal or cannot vary the default wager amount. Second, we generate wagers according to a Pareto distribution, reflecting the typical distribution of wealth in a population. We used a Pareto distribution with shape parameter 1.16 and scale parameter 1, which is often described as "20% of the population has 80% of the wealth." Each random set of wagers was scaled so that the average wager for any single match is exactly 1, allowing a comparison to the uniform wager case.

5.6.1 Notes on Implementation

A perfectly faithful implementation of the double clinching auction, as defined in Section 5.5, would require running an adaptive clinching auction for arbitrarily divisible goods with continuously increasing price function and allocations. In practice, it is necessary to discretize the price increases, thus computing allocations and prices that approximately match the double clinching auction.

One might be concerned that this discretization could adversely affect the nice properties of the double clinching auction. In particular, it might now be possible

---

[3] We thank Brian Galebach for providing us with this data.

[4] We also conducted simulations with random probability reports, generated both uniformly at random and according to a beta distribution. The results, and in particular the relative performance of the mechanisms, are very similar to those obtained using the ProbabilitySports data set.

for an agent to profit by misreporting her probability. To check whether this was the case, we empirically tested incentive compatibility on the 1643 matches from the ProbabilitySports data set. For each match, we chose a random agent i and ran the double clinching auction 101 times to calculate the expected payoff i would receive reporting each value in the set $\{0, 0.01, 0.02, \ldots, 0.99, 1\}$. We found a single profitable misreport for only a single one of these matches, with the misreporting agent able to increase her expected utility from 5.1611 to 5.1612. This suggests that the mechanism retains (at least approximate) incentive compatibility when discretized.

5.6.2 Results

The results are summarized in Table 5.1. The top table shows various statistics averaged across all 1643 matches, with wagers for each match drawn from a Pareto distribution. The Risk/Wagers column reports the total risk, summed across all agents, divided by the total wager, summed across all agents, or $\sum_{i \in \mathcal{N}} \sigma_i / \sum_{i \in \mathcal{N}} w_i$. A value of 1 means that every agent risks losing her entire wager for one outcome; a value of 0 means that no trade occurs. The %Full Stakes column reports the percentage of agents that risk losing their entire wager under one outcome (i.e., $\sigma_i = w_i$). The #Securities column gives the total number of pairs of securities sold to the agents, or $\sum_{i \in \mathcal{N}} y_i = \sum_{i \in \mathcal{N}} n_i$. The Principal Profit column shows the principal's net profit. Finally, the Agent Utility column gives the sum of the agents' expected utilities, assuming immutable beliefs and truthful reports.

As expected, the PCM facilitates the most trade, in terms of both the risk:budget ratio and the number of securities sold. However, there is a notably large gap in these metrics between the double clinching auction and the NAWM and WSWM, with the double clinching auction selling almost five times as many securities as the NAWM and WSWM. Additionally, under the double clinching auction, over 80% of agents risk their entire wagers, compared with no agents under NAWM and WSWM. This

Table 5.1: Evaluation of wagering mechanisms using reports from the ProbabilityS-ports data set. The top table was generated using Pareto distribution wagers, the bottom with uniform wagers.

	Risk/Wagers	%Full Stakes	#Securities	Profit	Agent Utility
PCM	0.94	97.2	601.1	0	200.1
DCA	0.80	82.6	489.3	28.4	152.1
NAWM	0.20	0	98.4	25.8	27.9
WSWM	0.16	0	101.2	0	53.7
	Risk/Wagers	%Full Stakes	#Securities	Profit	Agent Utility
PCM	0.97	96.9	618.5	0	208.1
DCA	0.97	96.4	616.4	1.3	206.3
NAWM	0.21	0	102.5	28.7	28.8
WSWM	0.17	0	105.7	0	57.4

is further evidence that falsely inflating a wager amount under the double clinching auction, while possibly beneficial in theory, would be extremely risky in practice, with a high chance of the manipulating agent losing more than her true budget.

We note that, while our objective is not to make a profit for the principal, the double clinching auction does yield a reasonable profit without sacrificing agent welfare.

The bottom table reports the same metrics when the agents' wagers are equal. While the three other mechanisms exhibit very similar performance in this case, the double clinching auction displays a marked increase in the amount of trade facilitated, under all metrics, and a drop in profit. For the objective of maximizing trade, this is a particularly compelling argument to use the double clinching auction in cases when equal wagers are natural.

Note that all matches in the ProbabilitySports data set have a relatively large number of agents participating. However, in many cases we are interested in instances with smaller numbers of agents. To investigate this behavior, we generated smaller instances by subsampling reports from the full set of reports for each match. Figure 5.1 plots the ratio of total risk to total budget, $(\sum_{i \in \mathcal{N}} \sigma_i)/(\sum_{i \in \mathcal{N}} w_i)$, for the

FIGURE 5.1: The risk:budget ratio for each of the four wagering mechanisms, plotted as a function of N.

four mechanisms for values of N ranging from 5 to 50, with wagers randomly drawn from a Pareto distribution. We see that while the PCM, NAWM, and WSWM exhibit only minimal change as N increases, the double clinching auction facilitates more trade for larger values of N. However, even for $N = 5$, the double clinching auction facilitates approximately twice the trade as the WSWM and the NAWM, suggesting that the double clinching auction is the best truthful mechanism when maximizing trade is a primary objective.

5.7 Discussion

We have defined and analyzed the double clinching auction, proving that it simultaneously satisfies incentive compatibility, budget balance, and individual rationality. While we showed that no wagering mechanism can simultaneously achieve these three properties along with Pareto optimality, our simulations suggest that the DCA comes close to Pareto optimality in practice, making it the first known incentive compatible wagering mechanism to do so.

It would be valuable, but apparently non-trivial, to extend the DCA to settings with non-binary outcomes. The DCA crucially exploits the fact that agents can be ordered by their reports in one dimension, allowing us to guarantee that no agent is

allocated both `yes` and `no` securities. With larger outcome spaces, this property no longer holds, and designing a mechanism in which the principal auctions off three or more types of securities would require novel techniques.

Even in the binary-outcome setting, a number of interesting problems remain. While our simulations suggest that the DCA comes close to achieving Pareto optimality, we have not established any formal approximation guarantee. An additional particularly compelling question is whether our choice of m^* is the largest number of securities that can be sold via a pair of adaptive clinching auctions while satisfying incentive compatibility and budget balance.

<div align="right">

6

</div>

Efficient Wagering by Relaxing Incentive Compatibility

6.1   Introduction

In Chapter 5, we considered the problem of *low stakes*, where existing wagering mechanisms are not (close to) Pareto optimal. We showed that Pareto optimality is incompatible with the core properties of individual rationality, budget balance, and incentive compatibility. Unwilling to give up any of the core properties, we designed a mechanism that was close to Pareto optimal: the double clinching auction.

In this chapter, we take a different approach. Rather than relaxing Pareto optimality, we relax full incentive compatibility, but achieve a formal approximation to incentive compatibility instead. In this sense, we get three of the four incompatible properties, and come 'close' to the fourth in a precise sense.

To do this, we consider the *parimutuel consensus mechanism (PCM)* [72], which can be seen as the equilibrium of parimutuel betting. In parimutuel betting, each bettor places money on one of several future outcomes—say, horse #1 to win a race. She is allowed to cancel her bet or move her money to a different outcome at any

time, even at the last second before wagering closes and the race begins. After the outcome resolves—say, horse #1 wins—agents who picked the wrong outcome lose their wagers to the agents who picked correctly. Winning agents split the pot in proportion to the size of their wagers.

The PCM is equivalent to parimutuel betting where each agent has a proxy. Each agent's proxy knows her true probabilities for all outcomes. As bets come in, and the prospective payoff per dollar, or *odds*, for each outcome converge, the proxy automatically switches its agent's money to the outcome yielding the highest expected payoff for that agent. In equilibrium, all the proxies are optimizing and none want to switch outcomes. At any point in time, the odds can be interpreted as probabillities, providing a prediction of the outcome of the event.

Eisenberg and Gale discuss one undesirable feature of the equilibrium: it produces odds that sometimes ignore some agents. Manski [118] further explores in detail how the equilibrium of risk-neutral, budget-constrained agents may fail to aggregate beliefs in a sensible way. Additionally, the PCM is not *incentive compatible*, or truthful. An agent may strategically improve her payoff by taking into account what other agents know or what they may do. In the end, her best action may be to report false probabilities to her proxy that differ from her true subjective probabilities. For a principal whose primary goal is information elicitation, this is problematic because some of the reported probabilities may not faithfully reflect the bettors' private information.

Given the potential for bad equilibria and the lack of incentive compatibility, why is the PCM still prevalent? One answer is that, in practice, it often works fine. Parimutuel betting does consistently induce a wisdom-of-crowds effect, producing odds that encode well calibrated and accurate probabilistic forecasts of the outcomes [18, 151, 155], like many prediction markets do [19]. Plott et al. [133] tested parimutuel betting in a laboratory experiment, showing that the mechanism is an

effective vehicle for information aggregation regardless of why it might go wrong in theory. If agents have concave or risk-averse utility for money, the equilibrium of similar mechanisms is stable and induces sensible belief aggregation [31, 166]. In particular, an agent with logarithmic utility does best by betting an amount on each outcome proportional to her probability [62].

In this chapter, we examine another plausible reason why the PCM continues to enjoy usage: the mechanism satisfies a number of desirable axioms for wagering systems. We prove that the PCM is the unique wagering mechanism that is Pareto optimal, individually rational, strictly budget balanced, sybilproof, anonymous, and envy-free, subject to a mild condition on the reports.

To address the lack of incentive compatibility, we show that the PCM is near incentive compatible in some sense. Yes, there are scenarios where agents can gain from lying, but we prove that the PCM is incentive compatible *in the large*, as the number of agents grows. In extensive simulations using real forecasts from an online contest, we show that opportunities for agents to profit from untruthful play are rare, mostly vanishing as the number of agents grows. Our results shed light on the practical success of the PCM. Despite its flaws, identified as early as 1959, it does satisfy six natural and desirable properties of wagering mechanisms and it comes close both theoretically and empirically to obtaining a crucial seventh: incentive compatibility.

6.2 Preliminaries

We work in the same setup as Chapter 5. We have a binary random variable X and a set of agents $\mathcal{N}$, each of whom submits a report $\hat{p}_i$ and wager w_i to the mechanism. We do however impose one additional restriction on the wagers by requiring they be rational. That is, $w_i \in \mathbb{Q}_{\geqslant 0}$. This assumption is required in Section 6.4. The output of the mechanism is specified by a triple $(\mathbf{y}, \mathbf{n}, \boldsymbol{\sigma})$, where y_i is the number of yes

securities allocated to agent i, n_i the number of `no` securities, and σ_i the price paid by i for those securities.

For this chapter, we introduce some more properties of wagering mechanisms. First, anonymity says that the payouts do not depend on the identities of the agents. This is a basic property that all wagering mechanisms proposed in the literature, including the PCM, satisfy.

Second, sybilproofness [112] says that it is not beneficial for agents to participate under multiple fake identities, or for agents reporting the same probability to merge.

Definition 15. *A wagering mechanism is* sybilproof *if for any subset of players S, for any $\hat{\mathbf{p}}$ with $\hat{p}_i = \hat{p}_j$ for $i, j \in S$, for any vectors of wagers $\mathbf{w}, \mathbf{w}'$ with $w_i = w_i'$ for $i \notin S$ and $\sum_{i \in S} w_i = \sum_{i \in S} w_i'$, it is the case that:*

$$\sum_{i \in S}(y_i(\hat{\mathbf{p}}; \mathbf{w}), n_i(\hat{\mathbf{p}}; \mathbf{w}), \sigma_i(\hat{\mathbf{p}}; \mathbf{w})) = \sum_{i \in S}(y_i(\hat{\mathbf{p}}; \mathbf{w}'), n_i(\hat{\mathbf{p}}; \mathbf{w}'), \sigma_i(\hat{\mathbf{p}}; \mathbf{w}'))$$

and for all $i \notin S$,

$$(y_i(\hat{\mathbf{p}}; \mathbf{w}), n_i(\hat{\mathbf{p}}; \mathbf{w}), \sigma_i(\hat{\mathbf{p}}; \mathbf{w})) = (y_i(\hat{\mathbf{p}}; \mathbf{w}'), n_i(\hat{\mathbf{p}}; \mathbf{w}'), \sigma_i(\hat{\mathbf{p}}; \mathbf{w}')).$$

Lastly, we consider the property of envy-freeness [76]. Envy-freeness is a basic fairness property which says that no player should envy the allocation of securities to any other agent.

Definition 16. *Say that agent i* envies *another agent j if $\sigma_j(\hat{\mathbf{p}}; \mathbf{w}) \leqslant w_i$ and*

$$\hat{p}_i y_i(\hat{\mathbf{p}}; \mathbf{w}) + (1 - \hat{p}_i)n_i(\hat{\mathbf{p}}; \mathbf{w}) - \sigma_i(\hat{\mathbf{p}}; \mathbf{w}) < \hat{p}_i y_j(\hat{\mathbf{p}}; \mathbf{w}) + (1 - \hat{p}_i)n_j(\hat{\mathbf{p}}; \mathbf{w}) - \sigma_j(\hat{\mathbf{p}}; \mathbf{w})$$

A wagering mechanism is envy-free if there is no pair of agents (i, j) such that i envies j.

We also note that for this chapter we use a slightly stronger definition of individual rationality than in the previous chapter. Our previous definition required that every

agent had *some* report that guaranteed non-negative expected payoff. The version that we use in this chapter requires that *truthful* reporting guarantees non-negative expected payoff, which is implied by the earlier version of individual rationality in combination with incentive compatibility. We require the stronger version in the proof of Theorem 74.

Definition 17. *A wagering mechanism is* individually rational *if, for any agent i and any belief p_i, for all $\hat{\mathbf{p}}_{-\mathbf{i}}, \mathbf{w}$ it holds that*

$$p_i y_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}) + (1 - p_i) n_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w}) \geqslant \sigma_i(p_i, \hat{\mathbf{p}}_{-i}; \mathbf{w})$$

6.3   The Parimutuel Consensus Mechanism

The Parimutuel Consensus Mechanism (PCM) can be thought of as a direct implementation of the equilibrium of parimutuel betting. The PCM includes the rules of parimutuel betting plus, conceptually, a proxy agent that automatically switches its agent's bet to the outcome with highest expected profit per security. The output of the mechanism is the equilibrium where all proxies are stable. For the binary case of `yes` and `no` outcomes that we consider, the PCM is defined by a price π such that an agent with probability less than π is allocated `no` securities at a price of $1 - \pi$ per security, and an agent with probability more than π is allocated `yes` securities at a price of π per security. The equilibrium condition is

$$\sum_{i:\hat{p}_i < \pi} \frac{w_i}{1 - \pi} + \sum_{i:\hat{p}_i = \pi} \frac{c_1 w_i}{1 - \pi} = \sum_{i:\hat{p}_i > \pi} \frac{w_i}{\pi} + \sum_{i:\hat{p}_i = \pi} \frac{c_2 w_i}{\pi}, \tag{6.1}$$

where c_1 and c_2 lie in the interval $[0, 1]$ and $\min\{c_1, c_2\} = 0$. These represent the possibility of needing agents with $\hat{p}_i = \pi$ to bet (some of) their wager to correctly balance the market prices and allow the market to reach equilibrium, even though they get zero expected profit. At most one of c_1 and c_2 is greater than 0, since it

166

would be redundant to have agents with $\hat{p}_i = \pi$ betting on both yes and no. Note that the left hand side of Equation 6.1 is the total number of no securities allocated, and the right hand side is the total number of yes securities allocated. Eisenberg and Gale [72] show as their main contribution that such a price is both unique and guaranteed to exist. The output of the PCM is defined by

$$y_i(\hat{\mathbf{p}}, \mathbf{w}) = \begin{cases} 0 & \hat{p}_i < \pi \\ \frac{c_2 w_i}{\pi} & \hat{p}_i = \pi \\ \frac{w_i}{\pi} & \hat{p}_i > \pi \end{cases}, \qquad n_i(\hat{\mathbf{p}}, \mathbf{w}) = \begin{cases} \frac{w_i}{1-\pi} & \hat{p}_i < \pi \\ \frac{c_1 w_i}{1-\pi} & \hat{p}_i = \pi \\ 0 & \hat{p}_i > \pi \end{cases}$$

and

$$\sigma_i(\hat{\mathbf{p}}, \mathbf{w}) = \begin{cases} w_i & \hat{p}_i < \pi \\ \max\{c_1, c_2\} w_i & \hat{p}_i = \pi \\ w_i & \hat{p}_i > \pi \end{cases}$$

Example 11. *Suppose that there are four agents, with reports $\hat{\boldsymbol{p}} = (0.3, 0.5, 0.6, 0.8)$ and wagers $\boldsymbol{w} = (1, 1, 3, 6)$. Observe that setting $\pi = 0.6$ and $c_1 = 2/3, c_2 = 0$ satisfies Equation 6.1: each side of the equation has value 10. Further, setting $\pi < 0.6$ results in the right hand side of Equation 6.1 being greater than the left hand side, for any allowed values of c_1 and c_2, and the opposite is true for any $\pi > 0.6$.*

We can now compute the output of the PCM on this instance, according to the formulae above. Agents 1 and 2 are allocated 2.5 yes securities each, for a price of 1, agent 3 is allocated 5 yes securities for a price of 2 (note that this is a $c_1 = 2/3$ fraction of agent 3's budget), and agent 4 is allocated 10 no securities for a price of 6.

Recall that, by Theorem 55, no wagering mechanism can simultaneously satisfy individual rationality, weak incentive compatibility, weak budget balance, and Pareto optimality. Theoretical papers on wagering mechanisms are generally reluctant to give up any of the first three properties, sacrificing Pareto optimality

[56, 82, 112, 113]. However, in practice, Pareto optimality is an important consideration and virtually all real-world wagering mechanisms, including parimutuels, bookmakers, and double auctions, satisfy it. This is because trade drives participation; a mechanism that facilitates little trade is of little use or interest to agents.

Individual rationality seems hard to give up. We cannot force agents to play a game that they expect to lose and, even if we did, they could just wager $w_i = 0$. The center may be willing to pay for the information inherent in the agents' beliefs, subsidizing the mechanism and relaxing budget balance. Market scoring rules [52, 101], for example, do just that, losing a strictly bounded amount of money in service of gaining information. However, a patron will only subsidize events that bear on valuable decisions. Nearly all fielded wagering mechanisms have taxes, not subsidies, yielding profits, not losses.

If we want Pareto optimality, individual rationality, and budget balance, we are forced to give up on incentive compatibility. That's exactly what the PCM does (see Example 12 for a concrete example). In the remainder of this paper, we show that the PCM is the unique wagering mechanism that simultaneously satisfies the other six properties of budget balance, individual rationality, Pareto optimality, anonymity, sybilproofness and envy-freeness, subject to a condition on the reports. We then show that, despite not satisfying incentive compatibility, the PCM is approximately incentive compatibile in two senses. First, we prove that, as the number of agents grows, the mechanism is incentive compatible *in the large*. Second, we show empirically that, across thousands of simulated wagers based on real probability estimates, opportunities to profit from misreports are almost negligible.

6.4   Properties of the PCM

Despite its theoretical flaws, including the possibility of nonsensical information aggregation, the PCM seems well behaved in practice. In this section, we examine

one possible reason for this by providing a theoretical justification for the PCM. We first note that the PCM satisfies six desirable properties for wagering mechanisms. Although incentive compatibility is not one of the six, we know that adding it is impossible without giving something up: no mechanism satisfying even just the first three properties can also be incentive compatible.

Proposition 73. *The parimutuel consensus mechanism satisfies individual rationality, budget balance, Pareto optimality, anonymity, sybilproofness, and envy-freeness.*

That the PCM satisfies the first three properties is noted by Freeman et al. [82].

Proof. For this proof, we assume that $c_2 = 0$ in the equilibrium condition given by Equation 6.1. The proof for the case where $c_1 = 0$ follows via symmetric arguments for all properties.

1. Anonymity Anonymity clearly holds because Equation 6.1 and the allocation of securities do not depend on the identities of the agents.

2. Individual rationality: Consider some agent i. If $p_i < \pi$,

$$p_i y_i(p_i, \hat{\mathbf{p}}_{-\mathbf{i}}; \mathbf{w}) + (1 - p_i) n_i(p_i, \hat{\mathbf{p}}_{-\mathbf{i}}; \mathbf{w}) = (1 - p_i)\frac{w_i}{1 - \pi} > w_i = \sigma_i(p_i, \hat{\mathbf{p}}_{-\mathbf{i}}; \mathbf{w})$$

If $p_i > \pi$,

$$p_i y_i(p_i, \hat{\mathbf{p}}_{-\mathbf{i}}; \mathbf{w}) + (1 - p_i) n_i(p_i, \hat{\mathbf{p}}_{-\mathbf{i}}; \mathbf{w}) = p_i\frac{w_i}{\pi} > w_i = \sigma_i(p_i, \hat{\mathbf{p}}_{-\mathbf{i}}; \mathbf{w})$$

Finally, if $p_i = \pi$,

$$p_i y_i(p_i, \hat{\mathbf{p}}_{-\mathbf{i}}; \mathbf{w}) + (1 - p_i) n_i(p_i, \hat{\mathbf{p}}_{-\mathbf{i}}; \mathbf{w}) = (1 - p_i)c_1\frac{w_i}{1 - \pi} = c_1 w_i = \sigma_i(p_i, \hat{\mathbf{p}}_{-\mathbf{i}}; \mathbf{w})$$

3. Budget balance: First, note that

$$\sum_{i \in \mathcal{N}} y_i(\hat{\mathbf{p}}; \mathbf{w}) = \sum_{\hat{p}_i > \pi} \frac{w_i}{\pi} = \sum_{\hat{p}_i < \pi} \frac{w_i}{1 - \pi} + c_1 \sum_{\hat{p}_i = \pi} \frac{w_1}{1 - \pi} = \sum_{i \in \mathcal{N}} n_i(\hat{\mathbf{p}}; \mathbf{w}),$$

169

where the second transition follows from the definition of π. Next,

$$\sum_{i\in\mathcal{N}} \sigma_i(\hat{\mathbf{p}}; \mathbf{w}) = \sum_{\hat{p}_i>\pi} w_i + \sum_{\hat{p}_i<\pi} w_i + \sum_{\hat{p}_i=\pi} c_1 w_i$$

$$= \pi \sum_{\hat{p}_i>\pi} \frac{w_i}{\pi} + (1-\pi)\left(\sum_{\hat{p}_i<\pi} \frac{w_i}{1-\pi} + c_1 \sum_{\hat{p}_i=\pi} \frac{w_1}{1-\pi}\right)$$

$$= \pi \sum_{\hat{p}_i>\pi} \frac{w_i}{\pi} + (1-\pi) \sum_{\hat{p}_i>\pi} \frac{w_i}{\pi}$$

$$= \sum_{\hat{p}_i>\pi} \frac{w_i}{\pi} = \sum_{i\in\mathcal{N}} y_i(\hat{\mathbf{p}}; \mathbf{w}) = \sum_{i\in\mathcal{N}} n_i(\hat{\mathbf{p}}; \mathbf{w}),$$

Where the third transition is obtained via the definition of π (and noting that $c_2 = 0$, by assumption).

4. Pareto optimality We show that the price π satisfies the condition of the Pareto optimality definition. From the definition of the PCM, if $\hat{p}_i > \pi$ then $\sigma_i(\hat{\mathbf{p}}; \mathbf{w}) = w_i$ and $n_i(\hat{\mathbf{p}}; \mathbf{w}) = 0$, and if $\hat{p}_i < \pi$ then $\sigma_i(\hat{\mathbf{p}}; \mathbf{w}) = w_i$ and $y_i(\hat{\mathbf{p}}; \mathbf{w}) = 0$.

5. Sybilproofness: Consider a set of sybils S such that $\mathbf{w}$ and $\mathbf{w}'$ satisfy the conditions of Definition 15, with corresponding prices π and π' reached by the PCM. By the definition of sybils, the following three conditions hold:

$$\sum_{i:\hat{p}_i<\pi} w_i = \sum_{i:\hat{p}_i<\pi} w_i', \qquad \sum_{i:\hat{p}_i=\pi} w_i = \sum_{i:\hat{p}_i=\pi} w_i', \qquad \sum_{i:\hat{p}_i>\pi} w_i = \sum_{i:\hat{p}_i>\pi} w_i'$$

It follows immediately that

$$\sum_{i:\hat{p}_i<\pi} \frac{w_i}{1-\pi} + c_1 \sum_{i:\hat{p}_i=\pi} \frac{w_i}{1-\pi} = \sum_{i:\hat{p}_i>\pi} \frac{w_i}{\pi}$$

$$\implies \sum_{i:\hat{p}_i<\pi} \frac{w_i'}{1-\pi} + c_1 \sum_{i:\hat{p}_i=\pi} \frac{w_i'}{1-\pi} = \sum_{i:\hat{p}_i>\pi} \frac{w_i'}{\pi},$$

so $\pi = \pi'$, with the same value of c_1 in both cases.

Suppose first that $i \notin S$. If $\hat{p}_i > \pi$ then

$$(y_i(\hat{\mathbf{p}}; \mathbf{w}), n_i(\hat{\mathbf{p}}; \mathbf{w}), \sigma_i(\hat{\mathbf{p}}; \mathbf{w})) = (\frac{w_i}{\pi}, 0, w_i)$$

$$= (\frac{w_i'}{\pi'}, 0, w_i') = (y_i(\hat{\mathbf{p}}; \mathbf{w}'), n_i(\hat{\mathbf{p}}; \mathbf{w}'), \sigma_i(\hat{\mathbf{p}}; \mathbf{w}')).$$

If $\hat{p}_i < \pi$ then

$$(y_i(\hat{\mathbf{p}}; \mathbf{w}), n_i(\hat{\mathbf{p}}; \mathbf{w}), \sigma_i(\hat{\mathbf{p}}; \mathbf{w})) = (0, \frac{w_i}{1 - \pi}, w_i)$$

$$= (0, \frac{w_i'}{1 - \pi'}, w_i') = (y_i(\hat{\mathbf{p}}; \mathbf{w}'), n_i(\hat{\mathbf{p}}; \mathbf{w}'), \sigma_i(\hat{\mathbf{p}}; \mathbf{w}')).$$

Finally, if $\hat{p}_i = \pi$ then

$$(y_i(\hat{\mathbf{p}}; \mathbf{w}), n_i(\hat{\mathbf{p}}; \mathbf{w}), \sigma_i(\hat{\mathbf{p}}; \mathbf{w})) = (0, c_1 \frac{w_i}{1 - \pi}, c_1 w_i)$$

$$= (0, c_1 \frac{w_i'}{1 - \pi'}, c_1 w_i') = (y_i(\hat{\mathbf{p}}; \mathbf{w}'), n_i(\hat{\mathbf{p}}; \mathbf{w}'), \sigma_i(\hat{\mathbf{p}}; \mathbf{w}')).$$

Next, suppose that $i \in S$. If $\hat{p}_i > \pi = \pi'$, then $\hat{p}_j = \hat{p}_i > \pi = \pi'$ for all $j \in S$. We have

$$\sum_{i \in S}(y_i(\hat{\mathbf{p}}; \mathbf{w}), n_i(\hat{\mathbf{p}}; \mathbf{w}), \sigma_i(\hat{\mathbf{p}}; \mathbf{w})) = \left(\sum_{i \in S} \frac{w_i}{\pi}, 0, \sum_{i \in S} w_i \right)$$

$$= \left(\sum_{i \in S} \frac{w_i'}{\pi'}, 0, \sum_{i \in S} w_i' \right) = \sum_{i \in S}(y_i(\hat{\mathbf{p}}; \mathbf{w}'), n_i(\hat{\mathbf{p}}; \mathbf{w}'), \sigma_i(\hat{\mathbf{p}}; \mathbf{w}')).$$

If $\hat{p}_i < \pi = \pi'$, then

$$\sum_{i \in S}(y_i(\hat{\mathbf{p}}; \mathbf{w}), n_i(\hat{\mathbf{p}}; \mathbf{w}), \sigma_i(\hat{\mathbf{p}}; \mathbf{w})) = \left(0, \sum_{i \in S} \frac{w_i}{1 - \pi}, \sum_{i \in S} w_i \right)$$

$$= \left(0, \sum_{i \in S} \frac{w_i'}{1 - \pi'}, \sum_{i \in S} w_i' \right) = \sum_{i \in S}(y_i(\hat{\mathbf{p}}; \mathbf{w}'), n_i(\hat{\mathbf{p}}; \mathbf{w}'), \sigma_i(\hat{\mathbf{p}}; \mathbf{w}')).$$

171

Finally, if $\hat{p}_i = \pi = \pi'$ then

$$\sum_{i \in S}(y_i(\hat{\mathbf{p}}; \mathbf{w}), n_i(\hat{\mathbf{p}}; \mathbf{w}), \sigma_i(\hat{\mathbf{p}}; \mathbf{w})) = \left(0, \sum_{i \in S} \frac{c_1 w_i}{1 - \pi}, \sum_{i \in S} c_1 w_i\right)$$

$$= \left(0, \sum_{i \in S} \frac{c_1 w_i'}{1 - \pi'}, \sum_{i \in S} c_1 w_i'\right) = \sum_{i \in S}(y_i(\hat{\mathbf{p}}; \mathbf{w}'), n_i(\hat{\mathbf{p}}; \mathbf{w}'), \sigma_i(\hat{\mathbf{p}}; \mathbf{w}')).$$

Therefore, the conditions for sybilproofness are satisfied.

6. Envy-freeness: Consider an agent i with $\hat{p}_i < \pi$. Let $j \neq i$. If $\sigma_j(\hat{\mathbf{p}}; \mathbf{w}) > w_i$ then i does not envy j, so suppose that $\sigma_j(\hat{\mathbf{p}}; \mathbf{w}) \leqslant w_i$.

Suppose that $\hat{p}_j > \pi$. Then

$$\hat{p}_i y_j(\hat{\mathbf{p}}; \mathbf{w}) + (1 - \hat{p}_i)n_j(\hat{\mathbf{p}}; \mathbf{w}) - \sigma_j(\hat{\mathbf{p}}; \mathbf{w}) = \hat{p}_i \frac{w_j}{\pi} - w_j$$

$$< 0$$

$$< \hat{p}_i y_i(\hat{\mathbf{p}}; \mathbf{w}) + (1 - \hat{p}_i)n_i(\hat{\mathbf{p}}; \mathbf{w}) - \sigma_i(\hat{\mathbf{p}}; \mathbf{w})$$

Next, suppose that $\hat{p}_j < \pi$. Then

$$\hat{p}_i y_j(\hat{\mathbf{p}}; \mathbf{w}) + (1 - \hat{p}_i)n_j(\hat{\mathbf{p}}; \mathbf{w}) - \sigma_j(\hat{\mathbf{p}}; \mathbf{w}) = (1 - \hat{p}_i)\frac{w_j}{1 - \pi} - w_j$$

$$= w_j\left(\frac{1 - \hat{p}_i}{1 - \pi} - 1\right)$$

$$\leqslant w_i\left(\frac{1 - \hat{p}_i}{1 - \pi} - 1\right)$$

$$= \hat{p}_i y_i(\hat{\mathbf{p}}; \mathbf{w}) + (1 - \hat{p}_i)n_i(\hat{\mathbf{p}}; \mathbf{w}) - \sigma_i(\hat{\mathbf{p}}; \mathbf{w})$$

Finally, suppose that $\hat{p}_j = \pi$. Then

$$\hat{p}_i y_j(\hat{\mathbf{p}}; \mathbf{w}) + (1 - \hat{p}_i)n_j(\hat{\mathbf{p}}; \mathbf{w}) - \sigma_j(\hat{\mathbf{p}}; \mathbf{w}) = (1 - \hat{p}_i)\frac{c_1 w_j}{1 - \pi} - c_1 w_j$$

$$= c_1 w_j\left(\frac{1 - \hat{p}_i}{1 - \pi} - 1\right)$$

$$\leqslant w_i\left(\frac{1 - \hat{p}_i}{1 - \pi} - 1\right)$$

$$= \hat{p}_i y_i(\hat{\mathbf{p}}; \mathbf{w}) + (1 - \hat{p}_i)n_i(\hat{\mathbf{p}}; \mathbf{w}) - \sigma_i(\hat{\mathbf{p}}; \mathbf{w})$$

The cases where $\hat{p}_i = \pi$ and $\hat{p}_i > \pi$ can be proven similarly. ∎

6.4.1 Axiomatic Characterization

We next show that the PCM is the *unique* wagering mechanism satisfying the six properties from Proposition 73, subject to a condition on the reports. Suppose that some non-zero wager is placed on $N > 3$ distinct reports, denoted by $P_1 < P_2 < \ldots < P_N$, and let $W_k = \sum_{i:\hat{p}_i = P_k} w_i$ be the total wager at probability P_k. We say that the *non-extreme* assumption holds if $P_2 W_1 < (1 - P_2) \sum_{i=3}^{N} W_i$ and $(1 - P_{N-1}) W_N < P_{N-1} \sum_{i=1}^{N-2} W_i$. For the data set used in Section 6.5 and wagers generated according to a Pareto($\alpha = 1.16$) distribution (see Section 6.5 for details), the non-extreme assumption held on over 99.97% of instances.

Theorem 74. *Let M be a wagering mechanism satisfying anonymity, individual rationality, budget balance, Pareto optimality, sybilproofness, and envy-freeness. If $N \geqslant 3$, $W_1 < \frac{1 - P_2}{P_2} \sum_{i=3}^{N} W_i$ and $W_N < \frac{P_{N-1}}{1 - P_{N-1}} \sum_{i=1}^{N-2} W_i$, then payoffs defined by M match those defined by the PCM.*

Proof. We first show that any wagering mechanism that satisfies envy-freeness, sybilproofness, and anonymity is defined by fixed prices p_y and p_n for yes and no securities. That is, for all agents i with $y_i(\hat{\mathbf{p}}; \mathbf{w}) > 0$, we have $p_y = \frac{\sigma_i(\hat{\mathbf{p}};\mathbf{w})}{y_i(\hat{\mathbf{p}};\mathbf{w})}$, and for all agents i with $n_i(\hat{\mathbf{p}}; \mathbf{w}) > 0$, we have $p_n = \frac{\sigma_i(\hat{\mathbf{p}};\mathbf{w})}{n_i(\hat{\mathbf{p}};\mathbf{w})}$.

To prove this, suppose otherwise for contradiction. That is, suppose that there exist agents i, j with $y_i(\hat{\mathbf{p}}; \mathbf{w}) > 0$ and $y_j(\hat{\mathbf{p}}; \mathbf{w}) > 0$ such that $\frac{\sigma_i(\hat{\mathbf{p}};\mathbf{w})}{y_i(\hat{\mathbf{p}};\mathbf{w})} > \frac{\sigma_j(\hat{\mathbf{p}};\mathbf{w})}{y_j(\hat{\mathbf{p}};\mathbf{w})}$. Consider a modified instance $(\hat{\mathbf{p}}; \mathbf{w}')$ in which both of i and j participate as sybils, denoted by sets S_i and S_j, instead of their individual identities, such that for all $k, \ell \in S_i \cup S_j$, we have that $\sigma_k = \sigma_\ell$. By sybilproofness and anonymity it must be the case that $\sigma_k(\hat{\mathbf{p}}; \mathbf{w}') = \sigma_i(\hat{\mathbf{p}}; \mathbf{w})/|S_i|$ and $y_k(\hat{\mathbf{p}}; \mathbf{w}') = y_i(\hat{\mathbf{p}}; \mathbf{w})/|S_i|$ for all $k \in S_i$, with the equivalent equalities holding for all $\ell \in S_j$ also. Therefore, for all $k \in S_i$ and

$\ell \in S_j$,

$$\frac{\sigma_k(\hat{\mathbf{p}}; \mathbf{w}')}{y_k(\hat{\mathbf{p}}; \mathbf{w}')} = \frac{\sigma_i(\hat{\mathbf{p}}; \mathbf{w})}{y_i(\hat{\mathbf{p}}; \mathbf{w})} > \frac{\sigma_j(\hat{\mathbf{p}}; \mathbf{w})}{y_j(\hat{\mathbf{p}}; \mathbf{w})} = \frac{\sigma_\ell(\hat{\mathbf{p}}; \mathbf{w}')}{y_\ell(\hat{\mathbf{p}}; \mathbf{w}')}.$$

Because $\sigma_k(\hat{\mathbf{p}}; \mathbf{w}') = \sigma_\ell(\hat{\mathbf{p}}; \mathbf{w}')$, k envies ℓ, violating envy-freeness in the modified instance. An identical argument shows the existence of a fixed price p_n for `no` securities.

By budget balance, the wagering mechanism must sell exactly the same number of `yes` and `no` securities, and it must be the case that each `yes`/`no` pair sells for exactly \$1. Therefore, $p_y + p_n = 1$. By individual rationality, it must be the case that all agents with $\hat{p}_i < p_y$ have $y_i(\hat{\mathbf{p}}; \mathbf{w}) = 0$, and all agents with $\hat{p}_i > p_y$ have $n_i(\hat{\mathbf{p}}; \mathbf{w}) = 0$.

We now use Pareto optimality, along with sybilproofness, anonymity, and envy-freeness, to show that whenever there exist agents i and j, with $\hat{p}_j > \hat{p}_i > p_y$, it must be the case that $\sigma_i(\hat{\mathbf{p}}; \mathbf{w}) = w_i$ and $\sigma_j(\hat{\mathbf{p}}; \mathbf{w}) = w_j$. We know by Pareto optimality that at least one of the equalities must hold; say, $\sigma_i(\hat{\mathbf{p}}; \mathbf{w}) = w_i$. Suppose for contradiction that $\sigma_j(\hat{\mathbf{p}}; \mathbf{w}) < w_j$. Again consider a modified instance $(\hat{\mathbf{p}}; \mathbf{w}')$ in which i and j participate as sybils, denoted by sets S_i and S_j, instead of their individual identities, such that for all $k, \ell \in S_i \cup S_j$, we have that $w'_k = w'_\ell$. By anonymity, we have $\sigma_k(\hat{\mathbf{p}}; \mathbf{w}') = w'_k$ for all $k \in S_i$ and $\sigma_\ell(\hat{\mathbf{p}}; \mathbf{w}') < w'_\ell$ for all $\ell \in S_j$. Now, using that fact that all agents are buying `yes` securities at price p_y, we have

that

$$\hat{p}_\ell y_\ell(\hat{\mathbf{p}}; \mathbf{w}') + (1 - \hat{p}_\ell)n_\ell(\hat{\mathbf{p}}; \mathbf{w}') - \sigma_\ell(\hat{\mathbf{p}}; \mathbf{w}')$$

$$= \hat{p}_\ell \frac{\sigma_\ell(\hat{\mathbf{p}}; \mathbf{w}')}{p_y} - \sigma_\ell(\hat{\mathbf{p}}; \mathbf{w}')$$

$$< \hat{p}_\ell \frac{w'_\ell}{p_y} - w'_\ell$$

$$= \hat{p}_\ell \frac{w'_k}{p_y} - w'_k$$

$$= \hat{p}_\ell \frac{\sigma_k(\hat{\mathbf{p}}; \mathbf{w}')}{p_y} - \sigma_k(\hat{\mathbf{p}}; \mathbf{w}')$$

$$= \hat{p}_\ell y_k(\hat{\mathbf{p}}; \mathbf{w}') + (1 - \hat{p}_\ell)n_k(\hat{\mathbf{p}}; \mathbf{w}') - \sigma_k(\hat{\mathbf{p}}; \mathbf{w}')$$

Therefore, agent $\ell \in S_j$ envies agent $k \in S_i$, violating envy-freeness. A similar argument can be used to show that $\sigma_i(\hat{\mathbf{p}}; \mathbf{w}) = w_i$ and $\sigma_j(\hat{\mathbf{p}}; \mathbf{w}) = w_j$ when $\hat{p}_j < \hat{p}_i < p_y$.

Next, suppose that $\hat{p}_j > \hat{p}_i = p_y$. We show that if $y_i(\hat{\mathbf{p}}; \mathbf{w}) > 0$ then $\sigma_j(\hat{\mathbf{p}}; \mathbf{w}) = w_j$. First, note that if $\sigma_i(\hat{\mathbf{p}}; \mathbf{w}) < w_i$, then Pareto optimality implies that $\sigma_j(\hat{\mathbf{p}}; \mathbf{w}) = w_j$. So suppose that $\sigma_i(\hat{\mathbf{p}}; \mathbf{w}) = w_i$. Then, because we have also assumed that $y_i(\hat{\mathbf{p}}; \mathbf{w}) > 0$, we know that $y_i(\hat{\mathbf{p}}; \mathbf{w}) = \frac{\sigma_i(\hat{\mathbf{p}};\mathbf{w})}{p_y} = \frac{w_i}{p_y}$. We can now use an argument identical to that used in the previous paragraph to argue that if $\sigma_j(\hat{\mathbf{p}}; \mathbf{w}) < w_j$, then we can create the same modified instance $(\hat{\mathbf{p}}; \mathbf{w}')$ so that sybils of j will envy sybils of i. Therefore, $\sigma_j(\hat{\mathbf{p}}; \mathbf{w}) = w_j$.

We now show that, provided the condition on reports in the statement of the theorem holds, $y_i(\hat{\mathbf{p}}; \mathbf{w}) > 0$ for all i with $\hat{p}_i = p_{N-1}$ (note that this, along with individual rationality, implies $p_y \leqslant p_{N-1}$). To see this, suppose otherwise. There are two cases. First, if $p_y < p_{N-1} < p_N$, then by our earlier observation it must be the case that $\sigma_i(\hat{\mathbf{p}}; \mathbf{w}) = w_i$ for all i with $\hat{p}_i = p_{N-1}$ or $\hat{p}_i = p_N$. Therefore, $y_i(\hat{\mathbf{p}}; \mathbf{w}) = \frac{w_i}{p_y} > 0$. Second, if $p_y \geqslant p_{N-1}$ and $y_i(\hat{\mathbf{p}}; \mathbf{w}) = 0$ for some i with $\hat{p}_i = p_{N-1}$,

then we can use sybilproofness and anonymity to argue that $y_i(\hat{\mathbf{p}}; \mathbf{w}) = 0$ for *all* i with $\hat{p}_i = p_{N-1}$. Therefore, the total number of yes securities allocated is strictly less than the total number of no securities allocated:

$$\sum_{i \in \mathcal{N}} y_i(\hat{\mathbf{p}}; \mathbf{w}) \leqslant \frac{W_N}{p_y} \leqslant \frac{W_N}{p_{N-1}} < \frac{1}{1 - p_{N-1}} \sum_{i=1}^{N-2} W_i \leqslant \frac{1}{p_n} \sum_{i=1}^{N-2} W_i \leqslant \sum_{i \in \mathcal{N}} n_i(\hat{\mathbf{p}}; \mathbf{w})$$

This violates budget balance. By a symmetric argument, we can show that $p_n \leqslant 1 - p_2$ and $n_i(\hat{\mathbf{p}}; \mathbf{w}) > 0$ for all i with $\hat{p}_i = p_2$.

In particular, the previous paragraph says that, subject to the conditions of the theorem holding, at least two bettors with distinct reports are allocated yes securities, and at least two bettors with distinct reports are allocated no securities. By the two earlier paragraphs, this implies that for all i with $\hat{p}_i > p_y$, we have $\sigma_i(\hat{\mathbf{p}}; \mathbf{w}) = w_i, y_i(\hat{\mathbf{p}}; \mathbf{w}) = \frac{w_i}{p_y}, n_i(\hat{\mathbf{p}}; \mathbf{w}) = 0$, and for all i with $1 - \hat{p}_i > p_n$, we have $\sigma_i(\hat{\mathbf{p}}; \mathbf{w}) = w_i, y_i(\hat{\mathbf{p}}; \mathbf{w}) = 0, n_i(\hat{\mathbf{p}}; \mathbf{w}) = \frac{w_i}{p_n}$.

Finally, it is easy to see that the only value of p_y/p_n that satisfies this condition while allocating an equal number of yes and no securities and satisfying $p_y = 1 - p_n$ is that defined by $p_y = \pi$ and $p_n = 1 - \pi$, from Equation 6.1. To characterize the allocations and payments of agents with $\hat{p}_i = p_y$, we note that these agents are required to exactly make up the difference between yes and no securities, if such a difference exists. By anonymity and sybilproofness, each of these bettors must be sold a number of securities that is proportional to their wager. This exactly matches the allocations and payments defined by the PCM. ∎

6.4.2 Incentive Properties of the PCM

As a Corollary of Theorem 55 and Proposition 73, we know that the PCM violates incentive compatibility. Intuitively, this is because agents are able to change the price π by changing their reports.

Example 12. *Let* $\mathbf{p} = (0.4, \frac{2}{3}, 0.8)$ *and* $\mathbf{w} = (1, 1, 1)$. *Then the outcome of the PCM is* $(\mathbf{y} = (0, 1.5, 1.5), \mathbf{n} = (3, 0, 0), \boldsymbol{\sigma} = (1, 1, 1))$. *Note that the price* $\pi = \frac{2}{3}$, *so agent 2's utility is 0. However, if agent 2 misreports* $\hat{p}_2 = 0.6$, *then the outcome becomes* $(\mathbf{y} = (0, \frac{5}{6}, \frac{5}{3}), \mathbf{n} = (2.5, 0, 0), \boldsymbol{\sigma} = (1, 0.5, 1))$. *Now the price* $\pi = 0.6$, *so agent 2's utility is* $\frac{5}{6} \cdot \frac{2}{3} - 0.5 = \frac{1}{18} > 0$.

The misreport in Example 12 has a particular form common to all profitable misreports. In order to change the price in a profitable way, a manipulating agent must ensure that her misreport exactly matches the new equilibrium price. The intuition is that the only way an agent can affect the price is to report a probability on the opposite side of the current price as her belief. However, such a misreport is only profitable if she does not 'over-shoot' and end up buying the wrong type of security.

Theorem 75. *Let* $\hat{p}_i \neq p_i$ *be a profitable misreport for agent* i. *Let* π_T *denote the* **yes** *security price when agent* i *reports truthfully, and* π_M *denote the* **yes** *security price in the instance when* i *misreports* $\hat{p}_i$. *Then it must be the case that* $\hat{p}_i = \pi_M$, *and either* $\hat{p}_i < \pi_T \leqslant p_i$ *or* $p_i \leqslant \pi_T < \hat{p}_i$.

Before we give the proof, we first state and prove a monotonicity lemma which states that, all else being equal, if the report of a single agent increases then the security price π also (weakly) increases.

Lemma 76. *Let* $\hat{\mathbf{p}}_{-\mathbf{i}} = \hat{\mathbf{p}}'_{-\mathbf{i}}$. *Let* $\hat{p}'_i < \hat{p}_i$, *and denote by* π' *the equilibrium price under vector of reports* $\hat{\mathbf{p}}'$, *and* π *the equilibrium price under vector of reports* $\hat{\mathbf{p}}$. *Then* $\pi' \leqslant \pi$.

Proof. Consider the equilibrium condition, Equation 6.1:

$$\sum_{j:\hat{p}_j < \pi} \frac{w_j}{1 - \pi} + c_1 \sum_{j:\hat{p}_j = \pi} \frac{w_j}{1 - \pi} = \sum_{j:\hat{p}_j > \pi} \frac{w_j}{\pi} + c_2 \sum_{j:\hat{p}_j = \pi} \frac{w_j}{\pi}$$

177

Suppose that $\hat{p}_i > \pi$ (other cases can be handled similarly). Suppose for contradiction that $\pi' > \pi$. Let c_1, c_2 represent the values of the equilibrium constants in the case that i reports $\hat{p}_i$, and c_1', c_2' represent those values when i reports $\hat{p}_i'$. Then we have

$$\sum_{j:\hat{p}_j>\pi} \frac{w_j}{\pi} + c_2 \sum_{j:\hat{p}_j=\pi} \frac{w_j}{\pi} \geqslant \sum_{j:\hat{p}_j>\pi} \frac{w_j}{\pi} > \sum_{j:\hat{p}_j>\pi} \frac{w_j}{\pi'}$$

$$\geqslant \sum_{j:\hat{p}_j'>\pi'} \frac{w_j}{\pi'} + c_2' \sum_{j:\hat{p}_j'=\pi'} \frac{w_j}{\pi'}$$

$$= \sum_{j:\hat{p}_j'<\pi'} \frac{w_j}{1-\pi'} + c_1' \sum_{j:\hat{p}_j'=\pi'} \frac{w_j}{1-\pi'}$$

$$\geqslant \sum_{j:\hat{p}_j'<\pi'} \frac{w_j}{1-\pi'}$$

$$\geqslant \sum_{j:\hat{p}_j<\pi} \frac{w_j}{1-\pi'} + c_1 \sum_{j:\hat{p}_j=\pi} \frac{w_j}{1-\pi'}$$

$$> \sum_{j:\hat{p}_j<\pi} \frac{w_j}{1-\pi} + c_1 \sum_{j:\hat{p}_j=\pi} \frac{w_j}{1-\pi}$$

where the equality holds by Equation 6.1, and the inequalities all hold due to the assumptions that $\hat{p}_i > \pi$ and that $\pi' > \pi$. Comparing the first and last line contradicts that π is the equilibrium price under reports $\hat{\mathbf{p}}$. Therefore, it must be the case that $\pi' \leqslant \pi$. $\blacksquare$

Proof of Theorem 75. Suppose that $p_i > \pi_T$. The cases $p_i < \pi_T$ and $p_i = \pi_T$ can be proven similarly. Note that if $\pi_M = \pi_T$, then $\hat{p}_i$ cannot be a profitable misreport, because under truthful reporting, i already buys as many **yes** securities as her budget allows, and these are the only securities from which she obtains positive expected profit at the current price π_T. Therefore, to show that any profitable misreport satisfies $\hat{p}_i < \pi_T$, we show that $\pi_M = \pi_T$ whenever $\hat{p}_i \geqslant \pi_T$.

Consider again Equation 6.1. For $\hat{p}_i > \pi_T$, if we set $\pi = \pi_T$ then each term in the equation takes the same value under truthful reporting and misreporting. Therefore,

178

equality holds in the misreported case with $\pi_M = \pi_T$. Next, if $\hat{p}_i = \pi_T < p_i$, then we know that $\pi_M \leqslant \pi_T$, by Lemma 29, since $\hat{p}_i < p_i$. It remains to rule out $\pi_M < \pi_T$. So suppose for contradiction that $\pi_M < \pi_T = \hat{p}_i < p_i$. Let c_1^M, c_2^M denote the equilibrium values of c_1 and c_2 when i misreports $\hat{p}_i$, and c_1^T, c_2^T the equilibrium values when i truthfully reports p_i. Then we have a similar system of inequalities as in the proof of Lemma 29,

$$\sum_{j:\hat{p}_j>\pi_M} \frac{w_j}{\pi_M} + c_2^M \sum_{j:\hat{p}_j=\pi_M} \frac{w_j}{\pi_M} \geqslant \sum_{j:\hat{p}_j>\pi_M} \frac{w_j}{\pi_M}$$

$$> \sum_{j:\hat{p}_j>\pi_M} \frac{w_j}{\pi_T}$$

$$\geqslant \sum_{j:\hat{p}_j>\pi_T} \frac{w_j}{\pi_T} + \frac{w_i}{\pi_T} + c_2^T \sum_{j:\hat{p}_j=\pi_T} \frac{w_j}{\pi_T}$$

$$= \sum_{j:\hat{p}_j<\pi_T} \frac{w_j}{1-\pi_T} + c_1^T \sum_{j:\hat{p}_j=\pi_T} \frac{w_j}{1-\pi_T}$$

$$\geqslant \sum_{j:\hat{p}_j<\pi_T} \frac{w_j}{1-\pi_T}$$

$$\geqslant \sum_{j:\hat{p}_j<\pi_M} \frac{w_j}{1-\pi_T} + c_1^M \sum_{j:\hat{p}_j=\pi_M} \frac{w_j}{1-\pi_T}$$

$$> \sum_{j:\hat{p}_j<\pi_M} \frac{w_j}{1-\pi_M} + c_1^M \sum_{j:\hat{p}_j=\pi_M} \frac{w_j}{1-\pi_M}$$

which contradicts that π_M is the equilibrium price when i reports $\hat{p}_i$. Therefore it is not the case that $\pi_M < \pi_T$, so $\pi_M = \pi_T$ and the misreport $\hat{p}_i \geqslant \pi_T$ is not profitable.

We have shown that $\hat{p}_i < \pi_T < p_i$ must hold for any profitable misreport $\hat{p}_i$. Therefore, by Lemma 29, we know that $\pi_M \leqslant \pi_T$. We now show that $\pi_M = \hat{p}_i$. First, suppose that $\hat{p}_i < \pi_M$. Then i is buying `no` securities at a price $1 - \pi_M \geqslant 1 - \pi_T > 1 - p_i$, where $1 - p_i$ is her value for a `no` security. Therefore, she obtains negative expected profit from this purchase, meaning that $\hat{p}_i$ is not a profitable misreport. Second, suppose that $\hat{p}_i > \pi_M$. In this case, we can argue by setting

$\pi = \pi_M$ in Equation 6.1. It is easy to see that at this equilibrium price, strictly more `yes` securities are sold than in the truthful case, and strictly fewer `no` securities. This violates budget balance, since equal numbers of `yes` and `no` securities are sold in the truthful case. Therefore, $\pi_M = \pi_T$. However, we have already established that if $\pi_M = \pi_T$, then $\hat{p}_i$ is not a profitable misreport, a contradiction. ∎

Incentive Compatibility in the Large. We now show that the PCM satisfies an approximate notion of incentive compatibility known as *incentive compatibility in the large (IC-L)*, introduced by Azevedo and Budish [22].[1] It relaxes incentive compatibility by requiring only that truthful reporting is optimal as the number of agents grows large, and that truthful reporting is only optimal in expectation over the reports, rather than based on the (ex-post) realization of reports, as our definition of incentive compatibility requires.

Conceptually, this section mirrors the work of Azevedo and Budish. Indeed, in cases where only a finite set of reports are allowed, that the PCM satisfies IC-L follows directly from the fact that the PCM satisfies envy-freeness (Azevedo and Budish show that this implies IC-L when the number of possible reports is finite). Since finite sets of reports can provide arbitrary precision, this is usually enough for practical purposes. Most real-world mechanisms only allow reports up to a precision of 1%, and this is also the precision we use in our simulations (see Footnote 3). However, for completeness, we provide an independent proof of IC-L for the case where an infinite number of reports are allowed. The proof is a simple extension of the finite reports case.

Let D denote a probability distribution over $[0, 1]$ with full support. We model each agent as drawing a report $\hat{p}_i$ i.i.d. according to D. So D models the distribution

---

[1] There is a large body of work focusing on other limiting IC criteria, including ε-strategyproofness [143, 71], that we do not focus on here.

of reports, not necessarily beliefs. We will assume that wagers are drawn i.i.d. from some fixed distribution bounded by the interval $[1, W]$ for some $W \geqslant 1$. In particular, the ratio of the wagers of any two agents is bounded by W. Denote the expected value of a randomly drawn wager by $\overline{w}$.

For the remainder of this section, let $(y_i(\hat{p}_i, w_i, D, n), n_i(\hat{p}_i, w_i, D, n), \sigma_i(\hat{p}_i, w_i, D, n))$ denote the expected allocation of securities and payment for an agent reporting $\hat{p}_i$ and wagering $w_i \in [1, W]$ when there are n other agents that draw reports according to D and wagers from the fixed wager distribution. Let $(y_i(\hat{p}_i, w_i, D, \infty), n_i(\hat{p}_i, w_i, D, \infty), \sigma_i(\hat{p}_i, w_i, D, \infty)) = \lim_{n \to \infty}(y_i(\hat{p}_i, w_i, D, n), n_i(\hat{p}_i, w_i, D, n), \sigma_i(\hat{p}_i, w_i, D, n))$. We can now formally define incentive compatibility in the large.

Definition 18. *A wagering mechanism is* incentive compatible in the large *if, for any D with full support over $[0, 1]$, and any $\hat{p}_i$ and w_i,*

$$p_i y_i(p_i, w_i, D, \infty) + (1 - p_i) n_i(p_i, w_i, D, \infty) - \sigma_i(p_i, w_i, D, \infty)$$

$$\geqslant p_i y_i(\hat{p}_i, w_i, D, \infty) + (1 - p_i) n_i(\hat{p}_i, w_i, D, \infty) - \sigma_i(\hat{p}_i, w_i, D, \infty).$$

To show that the PCM satisfies incentive compatibility in the large, we first show that when the number of bettors is large, no single agent can affect the security price π; that is, agents are price-takers in the large market limit. The second step is to show that price takers have no profitable manipulations, which follows immediately from Theorem 75.

Theorem 77. *The parimutuel consensus mechanism satisfies incentive compatibility in the large.*

Proof. Let π^n denote the price defined by the PCM in expectation when there are n agents drawing reports from D, as well as agent i reporting $\hat{p}_i$, and let $\pi^\infty = \lim_{n \to \infty} \pi^n$. We first show that π^∞ exists. For contradiction, suppose otherwise. Fix $\varepsilon > 0$. Then there exist arbitrarily large N_1, N_2 such that $|\pi^{N_1} - \pi^{N_2}| > \varepsilon$ for some

$\varepsilon > 0$. Suppose without loss of generality that $\pi^{N_1} > \pi^{N_2} + \varepsilon$. Note that we can rewrite the equilibrium condition, Equation 6.1,

$$\pi = \frac{\sum_{j:\hat{p}_j > \pi} w_j + c_2 \sum_{j:\hat{p}_j = \pi} w_j}{\sum_{j:\hat{p}_j \neq \pi} w_j + (c_1 + c_2) \sum_{j:\hat{p}_j = \pi} w_j}$$

Therefore, π^{N_1} and π^{N_2} are defined by

$$\pi^{N_1} = \frac{\sum_{j \neq i:\hat{p}_j > \pi^{N_1}} \bar{w} + c_2 \sum_{j \neq i:\hat{p}_j = \pi^{N_1}} \bar{w} + y_1 w_i}{\sum_{j \neq i:\hat{p}_j \neq \pi^{N_1}} \bar{w} + (c_1 + c_2) \sum_{j \neq i:\hat{p}_j = \pi^{N_1}} \bar{w} + w_i} \qquad (6.2)$$

$$\pi^{N_2} = \frac{\sum_{j \neq i:\hat{p}_j > \pi^{N_2}} \bar{w} + c_2 \sum_{j \neq i:\hat{p}_j = \pi^{N_2}} \bar{w} + y_2 w_i}{\sum_{j \neq i:\hat{p}_j \neq \pi^{N_2}} \bar{w} + (c_1 + c_2) \sum_{j \neq i:\hat{p}_j = \pi^{N_2}} \bar{w} + w_i} \qquad (6.3)$$

Where $y_1 = 1$ if $\hat{p}_i > \pi_1^N$ and $y = 0$ if $\hat{p}_i < \pi_1^N$, and similarly for y_2 in Equation 6.3 (for simplicity of notation, we ignore the case where $\hat{p}_i = \pi$, but it can be handled similarly). We can replace wagers w_j by $\bar{w}$ because we are interested in the price in expectation.

Since we can choose N_1 and N_2 to be artbitrarily large, the sum of all wagers w_j becomes large, and the effect of the wager w_i becomes arbitrarily small. Therefore, π^{N_1} and π^{N_2} become arbitrarily close to one another, violating the assumption that they are bounded apart by ε. Thus, π^∞ exists.

To see that π^∞ is independent of $\hat{p}_i$, we divide both the numerator and denominator of Equation 6.2 by N_1 and let $N_1 \to \infty$. The equilibrium condition becomes

$$\pi^\infty = \frac{Pr_{x \sim D}(x > \pi^\infty) + c_2 Pr_{x \sim D}(x = \pi^\infty)}{Pr_{x \sim D}(x \neq \pi^\infty) + (c_1 + c_2) Pr_{x \sim D}(x = \pi^\infty)}$$

Since this equation has no dependence on $\hat{p}_i$ (or w_i), π^∞ is independent of $\hat{p}_i$.

It now follows immediately from Theorem 75 that the PCM satisfies IC-L, since any profitable manipulation must alter the security price. But in the limit as the number of agents goes to ∞, it is impossible for i to affect the price π^∞. ∎

182

6.5 Simulations

We tested the incentive compatibility of the PCM on a data set consisting of probability reports gathered from an online prediction contest called ProbabilitySports [87].[2] The data set consists of probabilistic predictions about the outcome of 1643 National Football League matches from the start of the 2000 NFL preseason until the end of the 2004 season. For each match, between 64 and 1574 players reported their subjective probability of a fixed team (say, the home team) winning the match. Each match was scored according to the Brier scoring rule, with points contributing to a season-long scoreboard.

ProbabilitySports users submitted probabilities but not wagers. We generated wagers from a variety of Pareto distributions. Pareto distributions are a natural choice as they approximately model the distribution of wealth in a population. A Pareto distribution is defined by two parameters: a scale parameter k, which has the effect of multiplicatively scaling the distribution, and a shape parameter α, which affects the size of the distribution's tail. To allow for a fair comparison between distributions and instance sizes, we scaled each set of randomly generated wagers so that the average wager is 1. This means that changing the scale parameter has no effect, as the wagers are rescaled anyway. Therefore, we fix the scale parameter to 1 and vary only the shape parameter.

The first Pareto distribution we use for wager generation has $\alpha = 1.16$, which is often described as "20% of the population has 80% of the wealth," and classically viewed as a realistic distribution of wealth. Second, we use $\alpha = 3$, which produces a more equal distribution of wagers in comparison to $\alpha = 1.16$. Finally, we consider a uniform distribution of wagers (that is, $w_i = 1$ for all agents), corresponding to a situation either where all agents are equal, or where they do not have the

---

[2] We thank Brian Galebach for providing us with this data.

Table 6.1: Profitable misreports under Pareto and uniform wager generation.

	% Agents With Profitable Misreports	Average Profit	Average Wager per Profitable Misreport	Average Misreport Distance
Pareto($\alpha = 1.16$)	0.07	1.55	118.8	0.044
Pareto($\alpha = 3$)	< 0.01	0.03	5.76	0.015
Uniform	0	n/a	n/a	n/a

opportunity to choose their wager (as in the ProbabilitySports competition). Note that the uniform distribution is the limit of the Pareto distribution as $\alpha \to \infty$.

Our first step was to examine the entire dataset. For each of the 1643 matches and each wager distribution, we randomly generated a set of wagers drawn from that distribution. For each set of wagers we chose 50 random agents and simulated 101 reports for them in the range $\{0, 0.01, \ldots, 0.99, 1\}$.[3] For each report, we computed the agent's expected utility, taking their true belief to be their original report p_i. If there exists a misreport $\hat{p}_i \neq p_i$ that yields a higher utility than reporting their true belief, then the agent has a profitable misreport.

The results are summarized in Table 6.1. We report four statistics. The '% Agents With Profitable Misreports' column states the percentage of agents that are able to benefit from misreporting. The 'Average Profit' column gives, out of those agents with a profitable misreport, the average benefit that the agent can gain from misreporting optimally, over and above her utility from reporting truthfully. The 'Average Wager per Profitable Misreport' column gives the average wager of agents with a profitable misreport available. Finally, the 'Average Misreport Distance' gives, for those agents with a profitable misreport, the average distance between the optimal misreport and the true belief.

For wagers generated from a Pareto distribution with $\alpha = 1.16$, we found 55

---

[3] In principle, our setup allows agents to report at a higher precision than this, so there will be some possible misreports that we do not detect. However, we believe that considering reports of only multiples of 0.01 is reasonable, due to limited cognitive capacity of the agents and the practical constraints of many wagering systems.

profitable misreports (out of 82,150 agents that we checked), which means that only around 0.07% of the agents that we checked had a profitable misreport. It is striking to consider the makeup of this small percentage of agents. The average wager of these agents is 118.8 (recall that each set of wagers is scaled so that the average wager is 1). What we are seeing is that only agents with very, very high wagers have sufficient power to change the price π. In contrast, the average profit that these agents obtain by misreporting is only 1.55, suggesting that even these high-wager agents are unable to have too large an effect on the security price π. This average profit is on the order of 1-2% of the misreporting agents' wagers – arguably an insignificant amount. For those agents that do misreport, the optimal misreport only differs from their belief by around 0.04.

As α increases, the number of agents with opportunity to misreport decreases. Indeed, for the uniform wagers, we did not find a single opportunity to profitably misreport. This is not surprising, since when wagers are uniform and there are a large number of agents, no agent will ever be able to significantly affect the price.

So for the full data set, with $64 \leqslant n \leqslant 1574$ agents per match, opportunities to profitably misreport are scarce, as we would expect because the PCM satisfies IC-L. But what about instances with fewer agents? To investigate smaller instances, we subsampled smaller values of n from the complete set of reports and ran the same simulation. For each match, each value of $n \in \{10, 20, 30, 40, 50\}$, and each wager distribution, we randomly sampled n reports and generated wagers. For every instance generated in this way, we tested *every* agent to see whether they had a profitable misreport.

Figure 6.1(a) shows how the percentage of agents that can profitably misreport changes with instance size. Even with only 10 agents per instance, there are relatively few opportunities to profitably misreport, with around 10% of all agents being able to do so. This fraction decreases quickly as n increases – for instances with 50
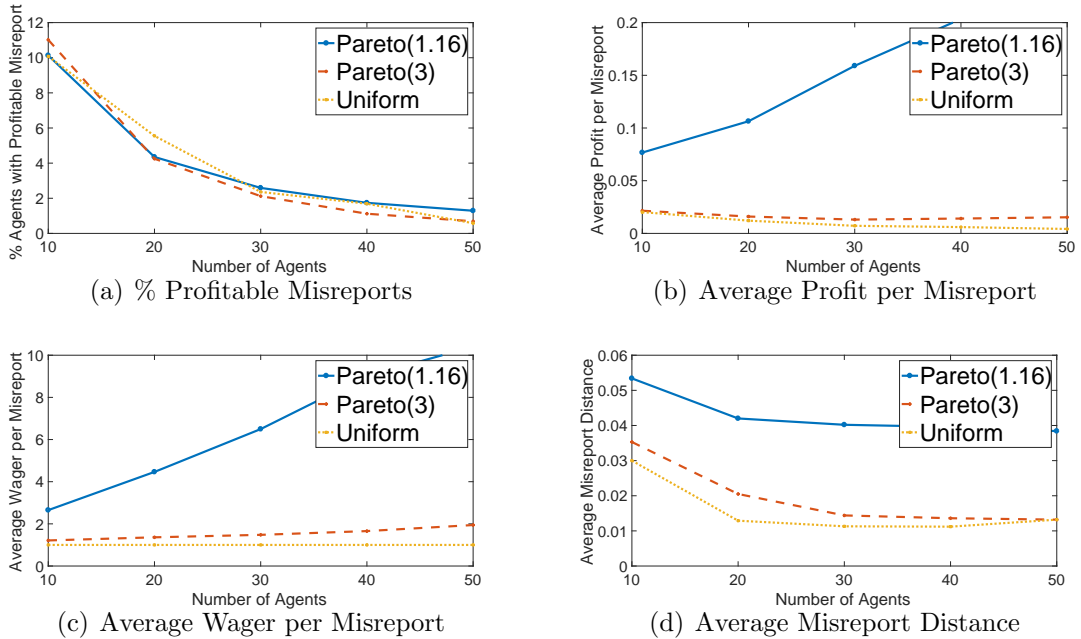
FIGURE 6.1: Profitable misreport characteristics for varying wager distributions and varying numbers of agents, n.

agents, less than 2% of agents are able to profitably misreport. Interestingly, all wager distributions exhibit approximately the same susceptibility to manipulation, in contrast to the full instances. We speculate that this is because, while high-wager agents are more likely to have profitable manipulations available, their existence also prevents low-wager agents from being able to manipulate, thus rendering the existence of high-wager agents something of a wash for small n. For large n, the latter effect disappears, since low-wager agents are unable to profitably misreport, even in the absence of high-wager agents.

Figure 6.1(b) shows how the average value of each profitable misreport changes with n, where the value of a profitable misreport is the difference in expected utility between the agent's optimal misreport and their truthful report. Interestingly, we see three very different trends depending on the wager distribution, all of which are consistent with the results on the full dataset. For $\alpha = 1.16$, the average value of a misreport steadily increases with n, as high-wager agents (who have high-value

misreports) become more and more frequent, while low-value misreports become less frequent. With uniform wagers, the value of a misreport quickly decreases with n. With only 10 agents, a misreporting agent may be able to affect the price quite significantly, however with increasing n, misreports will consist of only being able to make small adjustments to the security price. For $\alpha = 3$, the value of a misreport remains approximately constant as n increases, suggesting some combination of the two previous effects.

Figure 6.1(c) shows how the average wager of agents with a profitable misreport changes with n. For uniform wagers this line is flat, since all agents have wager $w_i = 1$. The other two wager distributions display increasing wagers, which is again explained by increasing frequency of high-wager agents (with this frequency increasing faster for $\alpha = 1.16$ than for $\alpha = 3$), and decreasing frequency of low-wager agents that are actually able to profitably misreport.

Finally, in Figure 6.1(d) we plot the average distance between a profitable misreport and an agent's true belief. In contrast with the other statistics that we consider, this one is actually relatively flat as n increases (with the exception of a significant drop from $n = 10$ to $n = 20$). This tells us that even for small numbers of forecasters, misreporting is limited to agents with beliefs fairly close to the price π and does not significantly affect the equilibrium price.

We note that we have considered an omniscient setting where manipulating agents have precise knowledge of the reports of other agents. In practice, of course, the manipulating agent has uncertainty about her opponents. A misreport is risky, involving some possibility of being forced to buy securities at a price favorable to her misreport but not her true belief. High-budget agents have the most opportunities to misreport but also the most to lose if they miscalculate.

6.6   Discussion

We have provided an axiomatic justification of the parimutuel consensus mechanism. While no wagering mechanism can satisfy anonymity, individual rationality, budget balance, Pareto optimality, sybilproofness, envy-freeness and incentive compatibility, we show that the PCM comes very close in that it satisfies all of the first six properties, and a relaxation of the seventh: incentive compatibility in the large. Subject to a mild condition on the reports, the PCM is the only wagering mechanism that satisfies all six properties. Via comprehensive simulations based on real contest data, we have shown that on large instances, opportunities to profitably manipulate are extremely rare. Even on small instances, the vast majority of agents cannot gain from misreporting.

A particularly interesting future direction would be to study the relative quality of forecasts generated by wagering mechanisms in practice, when bettors have several mechanisms to choose from. We would expect that the PCM would induce higher participation than non-Pareto optimal mechanisms, leading to a more accurate forecast, but that some agents may lie about their belief, leading to a less accurate forecast. Determining which of these pressures is more significant in practice would shed light on the most suitable wagering mechanism to implement for real forecasting applications.

7

Crowdsourced Outcome Determination in Decentralized Prediction Markets

7.1   Introduction

In this chapter, we move from wagering mechanisms to another common mechanism for probabilistic information elicitation and aggregation: prediction markets. The market operates by allowing participants to buy and sell securities which pay off according to the outcome of the event, and participants with an informational edge are able to place profitable trades when the market price disagrees with their own forecast. Through this trading process, the market price can be construed as a consensus forecast of the underlying event probability. Prediction markets have proven effective at forecasting events in a variety of domains, including business and politics [152, 30].

A key challenge in implementing and scaling prediction markets is the question of outcome determination (i.e., closing markets for events). Traditional prediction markets are centralized, in the sense that there exists a trusted center that creates markets, oversees transactions, and closes the market appropriately. The trusted

189

center is a bottleneck for defining and closing markets, limiting the scope of what can be predicted. However, there has recently been a rise of interest in *decentralized* prediction markets, where any user may create a market for an event. The markets are closed by consensus among a group of arbiters rather than by a center.

A decentralized platform removes the requirement for a highly trusted center, but it also allows each arbiter to directly influence the outcome of the market, in much the same way that agents may deliberately manipulate an event due to their own stake in the market; this is known as *outcome manipulation* [148, 30, 50]. Additionally, by allowing anyone to create a market, there is no longer any guarantee that all questions will be sensible, or even have a well-defined outcome. In this paper, we propose a specific prediction market mechanism with crowdsourced outcome determination that addresses several challenges faced by decentralized markets of this sort.

First is the issue of *outcome ambiguity*. At the time the market closes, it might be unreasonable to assign a binary value to the event outcome due to lack of clarity in the outcome. In a centralized market, it may be possible to postpone the market closing date to allow for rare cases of ambiguity, but it is not clear who should make such decisions in a decentralized marketplace. Therefore, it may be more fitting to allow outcomes to be non-binary, reflecting some level of disagreement. Outcomes in our mechanism are determined by the fraction of arbiters that report an event to have occurred. This also guarantees that every market is well-defined, by having traders explicitly trade on their expectations of the arbiter reports, not the actual event.

Second is the problem of *peer prediction*. For the credibility of the market, it is essential that arbiters are incentivized to truthfully report their opinion as to the realized outcome. If, for instance, we reward arbiters for agreeing with the majority opinion, then they are forced to anticipate the reports of other arbiters, not their independent opinion. We address this problem by making a technical change to an

existing peer prediction mechanism, the 1/prior mechanism.

Third is the inherent *conflict of interest* that arises by combining prediction markets and peer prediction mechanisms. Even if arbiters can be incentivized to report truthfully in isolation, there is no way to prevent them also having a stake in the market. An arbiter holding securities that pay off in a particular event will be incentivized to report that the event has occurred, even if they do not genuinely believe it to be the case, as long as they have a chance of affecting the market outcome. We tackle this issue by charging a trading fee that is later used to pay the arbiters. We show that, as long as each agent is responsible for a limited fraction of trading, and questions are clear enough, realistic trading fees can fully subsidize truthful reporting on the part of the arbiters.

Related Work. This work is inspired by decentralized prediction markets based on crypto-currencies, including Truthcoin, Gnosis, and especially Augur [132]. As in Augur, our mechanism consists of a prediction market stage and an arbitration stage, with trading fees from the market stage subsidizing the arbitration. The details of the mechanisms differ in both stages, however, and Augur includes additional layers of complexity such as a reputation system. While this complexity does provide additional protection against attack, Peterson and Krug [132] do not obtain any theoretical guarantees or justification for their chosen parameters. Clark et al. [57] also discuss outcome determination in crypto-based prediction markets, among several other implementation aspects.

Our work is most closely related to that of Chakraborty and Das [50], who consider a model where two agents participate in a prediction market whose outcome is determined by a vote among the two agents. Our model extends theirs by allowing an arbitrary number of traders, and not requiring that all traders are arbiters. Further, we take a mechanism design approach, obtaining an incentive compatible

mechanism, while Chakraborty and Das [50] focus on analyzing the equilibrium of a simple (to play) trading-voting game, with no peer prediction mechanism in the voting phase to incentivize truthful voting. Recent work by Witkowski et al. [164] also looks at a combination of forecasting and peer prediction, although the forecasts in their paper are elicited via proper scoring rules, rather than prediction markets.

The work of Bacon et al. [27] is similar in spirit to ours, as is the literature on outcome manipulation mentioned previously, but in all cases the concrete setting is quite different. We also draw heavily on existing literature in prediction markets [101, 52, 55] and peer prediction [119, 136, 163]; Chen and Pennock [53] survey these topics.

7.2   Preliminaries

Let N be a set of agents and let $A \subset N$ be a small set of distinct and verifiable (i.e., their identities are known to the mechanism) *arbiters*. Let $m = |A|$ denote the number of arbiters. The agents are anonymous in the sense that one cannot verify whether a trade is placed by an arbiter or non-arbiter, and whether several trades are placed by the same agent. Let X be a binary event with some realized outcome in $\{0, 1\}$. We are interested in setting up a prediction market for the outcome of X, with the resolution of the market decided upon by the arbiters.

Prediction markets.   We consider prediction markets implemented via a *Market Scoring Rule (MSR)*, where the underlying scoring rule is strictly proper [101, 52]. Every strictly proper MSR can be implemented as a market maker based on a *convex cost function*. Under this market structure, agents trade shares of a security with a centralized market maker, who commits to quoting a buy and sell price for the security at any time. The security payout is contingent on the outcome of X. In the usual implementation, one share of the security pays out \$1 in the event that $X = 1$, and \$0 otherwise.

Let q denote the total number of outstanding shares owned by the agents (note that q can be negative, in the case that more shares have been sold than bought). The market maker charges trades according to a convex, differentiable, and monotonically increasing function C. An agent wishing to buy $q' - q$ securities pays $C(q') - C(q)$. Negative payments encode a transaction where securities are sold back to the market maker. The instantaneous price of the security is given by $p = \frac{dC}{dq}$. Because the market maker always commits to trading, it may run a loss once the outcome is realized and the securities pay out, but the loss is bounded.

In practice, the cost function is also tuned using a *liquidity parameter* b, via the transformation $C_b(q) \equiv b\,C(q/b)$. A higher setting of b results in lower price responsiveness, in the sense that the price will change less for a fixed dollar trading amount. It also results in a higher worst-case loss bound for the market maker. Unless otherwise stated, our results assume that each agent participates in the market only once. The mechanism and results extend to situations in which agents can participate more than once, and we highlight these extensions where relevant throughout the paper.

Peer prediction. Peer prediction mechanisms are designed to truthfully elicit private information from a pool of agents via a reward structure that takes advantage of information correlation across agents. After the realization of X, each arbiter i receives either a positive or negative signal x_i, which we denote by $x_i = 1$ and $x_i = 0$ respectively. Let μ be the prior probability that an agent receives a positive signal. Let μ_1 (resp. μ_0) be the probability that, given agent i receives a positive (resp. negative) signal, another randomly chosen agent also receives a positive signal.[1] In a standard peer prediction belief model, the quantities μ_1 and μ_0 can be calculated

---

[1] Our analysis will assume that μ_1 and μ_0 are common across agents, but this is not a strict requirement. If we allow each agent to have distinct updates μ_1^i, μ_0^i, then we can let $\mu_1 = \min_i \mu_1^i$, corresponding to the minimum update given $\hat{x}_i = 1$, and similarly $\mu_0 = \max_i \mu_0^i$.

given μ and the signal beliefs $P(x_i = 1|X = 1)$ and $P(x_i = 1|X = 0)$; Witkowski [162] provides an overview. Assuming common information is not always reasonable, but it is natural in our setting to assume that agents take the closing price of the prediction market as their prior probability of receiving a positive signal (if not, then they can profit in expectation by participating in the market). The probabilities μ_1 and μ_0 are specific to the nature of the event X.

The peer prediction mechanism of interest in this work is the 1/prior ("one over prior") mechanism [162, 106, 107]. The 1/prior mechanism first asks each agent for their signal report $\hat{x}_i$. Then, every agent i is randomly paired with a peer agent $j \neq i$, and paid

$$
u(\hat{x}_i, \hat{x}_j) = \begin{cases} k\mu & \text{if } \hat{x}_i = \hat{x}_j = 0 \\ k(1 - \mu) & \text{if } \hat{x}_i = \hat{x}_j = 1 \\ 0 & \text{if } \hat{x}_i \neq \hat{x}_j, \end{cases}
$$

where k is some positive constant that can be freely adjusted to scale the payments received by the arbiters. Truthfully reporting $\hat{x}_i = x_i$ is an equilibrium provided that $\mu_1 \geqslant \mu \geqslant \mu_0$ [86]. This is a natural condition that we will assume throughout the paper—receiving signal $x_i = 1$ should not *decrease* i's estimate that another agent j also receives signal $\hat{x}_j = 1$. We also assume that at least one of the inequalities is strict, so that $\mu_1 > \mu_0$; this condition is known as *stochastic relevance*. Via a simple adaptation of the corresponding proof for the 1/prior mechanism, it can be shown that truthful reporting remains an equilibrium if μ is replaced by some other constant c with $\mu_0 < c < \mu_1$ in the payments; we will exploit this fact to adapt the 1/prior mechanism for our purpose.

We call the quantity $\delta = \mu_1 - \mu_0$ the *update strength*. This quantity is specific to the instance and, roughly speaking, measures how strongly positively correlated the signals are across arbiters. The update strength is high if, after receiving a positive (negative) signal, an arbiter believes that another given arbiter receives a

1. Market stage.

    (a) A prediction market is set up for an event X using a market scoring rule.
    (b) Agents trade in the market. For a security bought at price p, a trading fee of fp is charged, and for a security sold at price p, a fee of $f(1-p)$ is charged.
    (c) The market closes, trading stops.

2. Arbitration stage.

    (a) Each arbiter i receives a signal $x_i \in \{0, 1\}$ and reports an outcome $\hat{x}_i \in \{0, 1\}$.
    (b) Each arbiter i is assigned a peer arbiter $j \neq i$ and paid according to the 1/prior with midpoint mechanism.
    (c) The outcome of the market, and the payoff of each share sold, is set to the fraction of arbiters that report $\hat{x}_i = 1$.

FIGURE 7.1: Prediction Market with Outcome Determined Using Peer Prediction.

positive (negative) signal with high probability. For instance, if event X is "Will the Cleveland Cavaliers win the 2016 NBA playoffs?" then we would expect $\delta \approx 1$, since any arbiter reaching a conclusion about the outcome of the series (by watching it live, reading in the news, etc.) would strongly expect any other arbiter to reach the same conclusion. On the other hand, a question like "Will a Presidential candidate tell a lie in the televised debate?" is considerably more open to interpretation, and we would expect it to have a smaller value of δ. If an arbiter believes a candidate to have lied, it is not necessarily the case that another arbiter believes the same.

7.3 Mechanism

A step by step description of our mechanism is given in Figure 7.1. The mechanism runs a prediction market where the outcome is determined by a vote among arbiters. The arbiters' signals should be interpreted as the information they receive regarding

195

the outcome of X: checking news sources, observing events, their own opinions, etc. To ensure that arbiters truthfully report their information, we incentivize them via a peer prediction mechanism.[2] In both stages we implement non-standard versions of existing mechanisms, which we detail in the following.

Market stage

We make use of an MSR with non-binary outcome. The outcome takes a value $\hat{X} \in [0, 1]$ corresponding to the fraction of arbiters that report $\hat{x}_i = 1$. Each share sold pays off $\hat{X}$. Observe that this fundamentally changes the value of a security to a market participant: in a standard prediction market, an agent's value for a security is his subjective probability that event X occurs, while in our market his value is the fraction of arbiters that he expects to report $\hat{x}_i = 1$. However, given the agent's valuation for a security, his incentives in both markets are similar. A risk-neutral, non-arbiter agent will trade shares until the market price matches the security's expected payoff, or the agent's budget is exhausted.

This change to the payoff structure has two advantages. First, it ensures that any question has a well-defined and unambiguous outcome, avoiding problems with badly worded questions. This is important in any situation where users are allowed to generate markets. Second, any market with a binary outcome that relies on arbitration must have a point of 'discontinuity', where a change in report from a single arbiter results in the value of a security changing by \$1.[3] There will therefore always be situations where, given the reports of the other arbiters, a single arbiter completely controls the market outcome. If this arbiter also has a significant stake in the market, this creates a very large incentive problem. By utilizing non-binary

_____

[2] Each arbiter makes his report without knowledge of the report of any other arbiter; for instance, the reports could be made simultaneously.

[3] To see this, consider the case where all arbiters report $\hat{x}_i = 1$, and flip one report at a time to $\hat{x}_i = 0$. One of these flips must change the outcome from $\hat{X} = 1$ to $\hat{X} = 0$.

outcomes, a single arbiter can only change the value of each security by at most $1/m$.

Our mechanism imposes trading fees. Theoretical models of prediction markets do not typically incorporate trading fees (an exception is the work of Othman et al., where a fee in the form of a bid-ask spread is used to allow liquidity to increase over time), but they are standard in real-world implementations. To understand how the fee is implemented, it is important to distinguish between transactions (buy or sell) where an agent increases its position (in terms of risk), and transactions where it liquidates its position. The trading fee that we implement can be seen as a fee on the worst-case loss incurred by an agent: the fee is on p when a new security is bought, and $1 - p$ when a security is sold short (because it may pay out $1). However, no fee is levied when an agent sells back a share that it holds, or buys back a share that was previously sold short—these are liquidation transactions.

The trading fee serves two distinct purposes in our mechanism. First, it allows us to raise funds which can then be used to pay arbiters. Even assuming that arbiters behave honestly (in the absence of a sophisticated peer prediction mechanism), they still need to be compensated for the time spent looking up the outcome of X and reporting it to the mechanism. This can, in principle, be achieved by any of a number of fee structures.

Second, the fee provides natural bounds on the value of any given security. Even if an event is certain to occur, with a fee of $f = 2\%$ an agent who moves the market price to (say) 99c actually pays a marginal cost of $\$0.99 \cdot 1.02 > \$1$ (see Lemma 79 for an exact bound). The multiplicative fee effectively bounds the price of the security away from 0 and 1. Thus, it is impossible for an agent to buy securities at an arbitrarily cheap price, which allows us to bound the number of securities, and therefore maximum payout, of any agent with a fixed budget B. We note that there are other reasonable fee structures which do not provide such a lower bound on the

price. For example, if the agents only pay a fee on any profit they gain from the market, then the price of an event that is certain to happen will still converge to 1.

Arbitration stage

The main challenge in our setting is to incentivize arbiters to truthfully report their signal regarding the realized value of X. In the absence of any conflict of interest, this is a simple peer prediction problem. Since the closing price of the market gives us a natural common prior on the probability that a given arbiter receives signal $x_i = 1$, it is natural to use the 1/prior mechanism. For prior signal probability μ, the 1/prior mechanism uses the fact that $\mu_1 \geqslant \mu \geqslant \mu_0$ to guarantee that truthful reporting achieves higher payoff than misreporting. However, as μ_1 approaches μ, the payoff for truthfully reporting signal $\hat{x}_i = 1$ approaches the payoff for misreporting $\hat{x}_i = 0$. In isolation, there is still no reason to misreport, but if the arbiter has some stake in the market then it may be worthwhile to incur a small misreporting loss to achieve other gains. The following example illustrates this issue.

Example 13. *Consider a prediction market for the event "Will the Democratic presidential candidate be leading the Republican presidential candidate in the polls at the end of the month?" Suppose it is known that 10% of arbiters only check conservative news sources, which always report that the Republican candidate is ahead, and another 10% only check liberal news sources, which always report the opposite. Suppose the market closes at $\mu = 0.89$. Consider an arbiter j who checks a (moderate) news source and finds that the Democratic candidate is ahead (i.e., $x_j = 1$). Since it is still the case that 10% of the arbiters will certainly receive signal $x_i = 0$, the updated belief μ_1 remains no higher than 0.9. That is, the update is very small, and the expected profit from reporting $\hat{x}_j = 1$ is also small. If j has bet against the outcome (i.e., sold some securities to the market maker), it could be in his interest to lie and report $\hat{x}_j = 0$.*

However, if the moderate news site had reported that the Republican candidate was leading (i.e., $\hat{x}_j = 0$), the updated belief μ_0 could be quite small, even in the range of 0.1 (since most arbiters check moderate sources). Now j has a lot to gain from reporting $\hat{x}_j = 0$. Therefore, j would have to hold a relatively large number of shares for misreporting to outweigh the expected profit from the 1/prior mechanism.

Example 13 stems from an asymmetry in update strength, leading to potentially different incentives for arbiters depending on which signal they receive. We modify the mechanism, making the update strength symmetric. Given that we know the updated beliefs μ_1 and μ_0, we can pay arbiters according to the 1/prior mechanism but use the value $(\mu_1 + \mu_0)/2$ instead of the prior, μ. We call this the *1/prior with midpoint* mechanism. Using the midpoint guarantees that the incentives for arbiters are the same regardless of the signal they receive. For the arbiter with the greatest incentive to misreport, using the 1/prior with midpoint mechanism (weakly) decreases his incentive to misreport over the standard 1/prior mechanism, allowing us to achieve better bounds in our worst-case analysis.

Analysis

In this section, we derive conditions for truthful reporting ($\hat{x}_i = x_i$) to be a best response, given that all other arbiters report truthfully. The main restriction we require is an upper bound B on the total budget any given arbiter spends in the market—without such a bound, an arbiter could have an arbitrarily large incentive to manipulate the market's outcome. Thus, B appears as a parameter in our analysis.

Arguably, an arbiter confident in their ability to manipulate a market outcome could procure enough funds as to have a very large budget, especially relative to a small market. However, in current decentralized prediction markets, each arbiter arbitrates only a small fraction of markets. As long as the assignment of arbiters to markets is done *after* the market closes, there is no way for manipulators to target

199

a specific market. For this reason, we believe that manipulations are most likely to be of a form where arbiters participate honestly in the first stage, but, if they happen to be assigned to arbitrate a market that they also participated in, may be able to gain by not reporting truthfully, rather than arbiters mounting deliberate high-budget attacks in the market stage. Of course, our analysis is not specific to that particular interpretation, but we do consider it a compelling argument in favor of using a budget bound in our analysis.

Intuitively, we need to scale the payments made to arbiters in the arbitration stage by a sufficiently large k so that the increased payoff for truthful reporting in this stage overwhelms the gains from manipulating the outcome.

Lemma 78. *Let n_i be the number of securities held by arbiter i. Then truthfully reporting $\hat{x}_i = x_i$ is a best response for arbiter i, given that all other arbiters report truthfully, if and only if*

$$k \geqslant \frac{2|n_i|}{m\delta}.$$

Proof. We prove the case where $n_i > 0$; the case for $n_i < 0$ is symmetric. The total payoff for arbiter i is the sum of the payoffs from the market phase and the arbitration phase. Fixing the reports of the other arbiters, the market payout for i is higher when i reports $\hat{x}_i = 1$. And, in expectation, the payoff for i in the arbitration phase is higher for truthful reporting than for lying. Thus, the only problematic case is when $x_i = 0$, but i may wish to report $\hat{x}_i = 1$.

So suppose that $x_i = 0$. The expected payoff for truthfully reporting $\hat{x}_i = 0$, assuming all other arbiters truthfully report their signal, is

$$n_i\mu_0 \frac{m-1}{m} + (1 - \mu_0)k \frac{\mu_0 + \mu_1}{2}. \tag{7.1}$$

Here $\mu_0(m-1)$ is the expected number of arbiters that report signal 1, and therefore $n_i\mu_0(m-1)/m$ is i's expected payoff from the market, while the remaining term is

$1 - \mu_0$, the probability of i's peer agent also reporting 0, multiplied by the payment i receives in this case. On the other hand, the expected payoff for misreporting $\hat{x}_i = 1$ is

$$n_i \left(\mu_0 \frac{m-1}{m} + \frac{1}{m} \right) + \mu_0 \, k \left(1 - \frac{\mu_0 + \mu_1}{2} \right), \tag{7.2}$$

where the extra $1/m$ in the first term is due to the additional market payoff from i reporting $\hat{x}_i = 1$, and the latter term is now the probability of i's peer agent reporting 1, multiplied by the payoff i receives when this happens.

We require that the expected payoff for reporting $\hat{x}_i = 1$ be at most the expected payoff for truthfully reporting $\hat{x}_i = 0$. Setting term (7.2) to be at most term (7.1) and simplifying yields the result. ∎

This characterization requires an upper bound on the number of securities that any single agent owns. In itself this is an unsatisfying restriction; however, we can think about it in terms of the size of the fee, f, and the amount of money that any single arbiter spends in the market, B. For fixed fee f, let q^- and q^+ be the number of outstanding securities such that the market prices are $p(q^-) = f/(1 + f)$ and $p(q^+) = 1/(1 + f)$ respectively. Note that these quantities depend on the liquidity parameter b used in the cost function.

Lemma 79. *For fixed percentage fee f, the number of outstanding securities lies in the interval $[q^-, q^+]$.*

Proof. Suppose that some agent sells a security when there are already q^- outstanding. Then the marginal price is exactly $f/(1 + f)$. When selling a security at this price, the agent receives $f/(1 + f)$ from the mechanism but must pay a trading fee of

$$f \left(1 - \frac{f}{1 + f} \right) = \frac{f}{1 + f}.$$

Thus the agent's net revenue from the sale is 0 (and the possibility remains that he must pay the mechanism in the event that X occurs). Therefore no agent makes such a sale, and the number of outstanding securities never drops below q^-.

A similar argument shows that q never exceeds q^+. To buy a security when there are already q^+ outstanding, an agent must pay a price of at least \$1, when the fee is included. ∎

Lemma 79 provides us with the minimum and maximum number of outstanding securities at any time. As a corollary, we can derive the maximum number of securities that a single agent with budget B is able to purchase or short sell. We interpret the budget as an upper bound on the worst-case loss that the agent is able to incur. When buying a security for price p, the worst-case loss is p, under outcome $X = 0$. When selling a security for price p, the worst-case loss is $1 - p$, under outcome $X = 1$. Let $\phi_b^+(B) = C_b^{-1}(B + C_b(q^-)) - q^-$. Define q' implicitly by $B + C_b(q^+) - C_b(q') = q^+ - q'$, and let $\phi_b^-(B) = q' - q^+$.

Corollary 80. *At the end of the market stage, an agent i with budget B holds $n_i \in \left[\phi_b^-(B), \phi_b^+(B)\right]$ securities.*

Proof. We first show the upper bound. Given an existing number of outstanding securities, q, an agent is able to increase the number of outstanding securities to q', where

$$C_b(q') - C_b(q) = B. \tag{7.3}$$

For a fixed budget, the maximum number of securities that can be bought in a single transaction is in the case that q is as small as possible; in our case, $q = q^-$. Substituting into (7.3) gives

$$q' = C_b^{-1}(B + C_b(q^-)).$$

The number of securities held by i is $q' - q = q' - q^-$, which gives the upper bound in the statement of the corollary.

We now show the lower bound. Given an existing number of securities q, an agent is able to decrease the number of outstanding securities to q', where

$$B + C_b(q) - C_b(q') = q - q'. \tag{7.4}$$

The right hand side of (7.4) is the number of securities sold by the agent to the mechanism, and therefore the amount that he may be required to pay the mechanism in the case that $X = 1$. The left hand side is exactly the funds that the agent is able to use to reimburse the mechanism: his budget, B, plus the amount paid to the agent by the mechanism for the securities, $C(q) - C(q')$. The maximum number of securities that can be sold in a single transaction is in the case that $q = q^+$; making this substitution in (7.4) yields the implicit formula for q' in the definition of $\phi_b^-(B)$. The number of securities sold by i is $q' - q = q' - q^+$, which gives the lower bound in the statement of the corollary. ∎

An interesting special case is the limit as $b \to \infty$. This corresponds to the market having zero price responsiveness, meaning that all securities are purchased at a fixed price. Conceptually, it is equivalent to the situation where agents participate in the market more than once. In that setting, an agent could wait until the market price reaches $\frac{f}{1+f}$, buy a small number of securities, then wait again until the price drops. An agent spending all their budget in this way can, in the extreme case, buy as if the market has infinite liquidity.

Corollary 81. *For an agent that spends at most B dollars in a market with trading fee f and infinite liquidity, n_i lies in the range $\left[-\frac{B(1+f)}{f}, \frac{B(1+f)}{f} \right]$.*

Proof. The minimum price for a single security is $\frac{f}{1+f}$, by Lemma 79 and the definition of q^-. Therefore, an agent with budget B can buy at most $\frac{B(1+f)}{f}$, the upper

bound in the corollary statement.

The maximum price for a single security is $\frac{1}{1+f}$, by Lemma 79 and the definition of q^+. Thus, an agent selling a security has worst case loss at least $1 - \frac{1}{1+f} = \frac{f}{1+f}$. So, an agent with budget B can sell at most $\frac{B(1+f)}{f}$ securities, which yields the lower bound. ∎

If every agent has budget at most B in the market stage, we can combine the bounds from Corollaries 80 and 81 and Lemma 78 to determine the minimum payment that guarantees truthful reporting in the arbitration phase.

Theorem 82. *Given that all other arbiters report truthfully, truthful reporting is a best response for arbiter i if*

$$k \geqslant \frac{2 \max\{|\phi_b^-(B)|, |\phi_b^+(B)|\}}{m\delta}.$$

In the case that agents may participate in the market many times, truthful reporting requires that

$$k \geqslant \frac{2B(1+f)}{fm\delta}.$$

Proof. The theorem follows directly from substituting the lower bound on n_i from Corollary 80 and Corollary 81 into the inequality from Lemma 78. ∎

Therefore, fixing an agent budget B and a trading fee f, we know how large one needs to make the payments in the arbitration phase in order to incentivize truthful reporting. We now take a global view, and examine the total funds required to incentivize all arbiters to report truthfully.

Lemma 83. *The total payment made to the arbiters is at most mk. We can implement a truthful equilibrium with total payment at most*

$$\frac{2 \max\{|\phi_b^-(B)|, |\phi_b^+(B)|\}}{\delta}.$$

204

In the case that agents may participate in the market many times, we require total payment at most

$$\frac{2B(1+f)}{f\delta}.$$

Proof. As $0 \leqslant \mu_0, \mu_1 \leqslant 1$, their mean also lies between 0 and 1, and therefore each arbiter's payment in the 1/prior with midpoint mechanism is at most k. Thus the total payment to the arbiters is at most mk, which proves the first part. Combining this with the bounds on k from Theorem 82 yields the second part. ∎

Now that we have an expression for the total amount needed to pay the arbiters, we can determine a suitable value for the fee f so that the mechanism does not need any outside subsidy to finance these payments. Let c_i denote the total cost paid by agent i to the mechanism (so c_i is negative if agent i sells securities). Define M by

$$M = \sum_{i:n_i>0} c_i + \sum_{i:n_i<0} (n_i + c_i).$$

M can be interpreted as the sum of the worst-case losses of the agents. By definition, the total fee revenue collected by the mechanism is fM. The mechanism is guaranteed to generate enough fees to incentivize truthful reporting if the revenue is at least as large as the total payment required for the arbiters. We state this result as a theorem.

Theorem 84. *The mechanism generates enough fee revenue to pay the arbiters without requiring any outside subsidy if*

$$fM \geqslant \frac{2 \max\{|\phi_b^-(B)|, |\phi_b^+(B)|\}}{\delta}. \tag{7.5}$$

If agents may participate in the market many times, then we require that

$$fM \geqslant \frac{2B(1+f)}{f\delta}. \tag{7.6}$$

Observe that inequality (7.6) aligns with intuition. An increase in total trader spend M, or the trading fee f, makes it easier to incentivize the arbiters to report truthfully since the market collects more revenue. Likewise, an increase in δ helps us satisfy the inequality, since a large update strength increases the incentive for arbiters to report truthfully to the peer prediction mechanism. However, a large value of B increases the stake that any single arbiter can have in the market, which in turn increases their payoff for misreporting.

An interesting feature of inequalities (7.5) and (7.6) is the lack of any dependence on the number of arbiters m. One might expect that increasing the number of arbiters would be beneficial, since this reduces the influence that any one of them has on the market outcome. However, this is canceled out by the fact that as we add arbiters, the payment per arbiter decreases, so that we cannot incentivize them as strongly.

7.4   Parameter Calibration

(a) Multiple Entry Case   (b) Single Entry Case
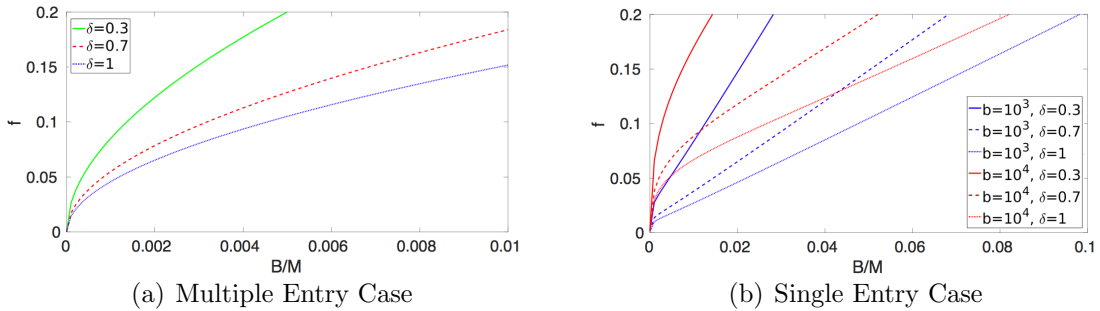
FIGURE 7.2: Minimum fee f required to adequately incentivize arbiters, plotted as a function of $\frac{B}{M}$. In both cases, $M = 10^6$ is fixed. Relationships are shown for selected values of update strength δ and, in the right-hand plot, liquidity b.

In this section we investigate the constraints imposed by inequalities (7.5) and (7.6). The purpose of the exercise is to illustrate how Theorem 84 can be used to inform the choice of fee f, and to confirm that realistic fees could be charged in practice to subsidize truthful arbitration. We consider the logarithmic market scoring rule

(LMSR), which is the most common MSR used in practice. For the LMSR, the cost
and price functions are

$$C_b(q) = b\log(1 + e^{q/b}), \qquad p(q) = \frac{e^{q/b}}{1 + e^{q/b}}.$$

By the symmetry of LMSR, $q^- = -q^+$ and $\phi_b^-(B) = -\phi_b^+(B)$. We will therefore
solve for $\phi_b^+(B)$. To find q^-, we set $p(q) = f/(1 + f)$ and solve for q, which gives
$q^- = b\log f$. Now, substituting the relevant components into the expression $\phi_b^+(B) =$
$C_b^{-1}(B + C_b(q^-)) - q^-$ leads to the following expression for inequality (7.5):

$$fM \geqslant \frac{2b(\log((1 + f)e^{B/b} - 1) - \log f)}{\delta}. \qquad (7.7)$$

In the case where we allow agents to participate multiple times, inequality (7.6)
remains unchanged.

We plot (7.5) and (7.6) in Figure 7.2, considering their tight versions as equalities.
First consider Figure 7.2(a), which represents the worst-case scenario in which agents
can enter multiple times and potentially spend their entire budget buying securities
at minimum price p^-. Suppose that some entity is creating a prediction market for
event X. Having decided on a question, the main decision is what value to set for f,
typically in the 2-5% range. To do so, the market creator needs to first estimate a
value for δ, which will be determined by question clarity, whether the arbiters have
reliable sources to check the outcome, and other such factors. Each line in the graph
represents a specific value of δ. With δ fixed, the market creator can estimate a
value for $\frac{B}{M}$. This is the maximum proportion of money that any single arbiter will
contribute to the market. We would expect $\frac{B}{M}$ to be small for markets that generate
a lot of interest, while niche markets would be vulnerable to having a single agent
contribute a large percentage of the total trade. Given these values, the creator can
arrive at the smallest f that is guaranteed to subsidize truthful reporting. From

the graph, we see that in the case of a question where $\delta = 1$ and $\frac{B}{M} = 0.001$, we can subsidize the arbiter payment with a fee of approximately 4%. This may seem large for a clear question with high participation, but we stress that this fee is based on a severe worst case where an agent is able to spend its entire budget purchasing securities at the minimum price.

Now consider Figure 7.2(b), which returns to the case where an agent only enters once, where liquidity now plays a role and we have to consider different values for parameter b. Figure 7.2(b) includes two reasonable values for b, as well as three different values for δ. We note that the situation looks considerably better for the market creator; indeed, the horizontal axis is now ten times larger indicating that we can now handle much smaller markets. When $\delta = 1$, we can handle situations where a single agent can contribute as much as 2% of the total trade with a fee of less than 5%. Even for questions with δ as low as 0.3, in a market with $b = 1000$ and $\frac{B}{M} = 0.005$ the fee can be set to approximately 5%.

7.5   Discussion

This paper proposed and analyzed a mechanism where the outcome of an MSR prediction market is determined via a peer prediction mechanism among a set of arbiters. The mechanism relies on two key adaptations to incentivize truthful arbitration: market shares pay out according to the proportion of arbiters who vote affirmatively, instead of a binary payout, and peer prediction payments are based on the midpoint of the two possible posteriors, rather than the prior. We showed that, with this combination of adaptations, it is possible to charge a trading fee that fully subsidizes truthful arbitration. Calibration based on plausible values of question clarity and trading volume suggests that realistic fees of 5% should be sufficient in practice.

While we have derived conditions under which truthful reporting is an equilib-

rium, there remains the possibility of the arbiters reporting according to uninformative equilibria that achieve higher payoff. This problem has recently been addressed in the peer prediction literature in situations where reporters complete several tasks instead of just one [65, 150]; it may be worthwhile to apply these techniques to our setting. In practice, arbiters vote on many questions across time, which opens the possibility of using a reputation system to incentivize them to vote truthfully and accurately [132]. The interplay of the incentives from all these mechanisms is fertile ground for future research.

8

Conclusion

In this thesis, we have considered algorithm design to allow for better decision making by harnessing information provided by a set of agents. We have considered two classes of problem: fairly allocating shared resources, and probabilistic forecasting.

For the allocation of shared resources, we have proposed and analyzed a model of public decision making that generalizes private goods allocation. We proposed novel fairness guarantees and analyzed algorithms for computing fair allocations in both the offline and online setting. We also considered an online, private goods allocation problem and proposed the flexible lending mechanism, which achieves strategy proofness and a 0.5 approximation to sharing incentives, as well as empirically high efficiency.

Many open problems remain; see the individual chapters for details. A particularly interesting theme for future research in this space is to investigate the limits of incentive compatibility. For private goods, there are many settings in which reasonable incentive guarantees can be obtained, while for public goods, the problem seems hopeless due to the famous free-rider problem. But there are intermediate settings, such as excludable, non-rival (i.e., club) goods, where there is evidence that incentive

guarantees are still possible [139, 122, 66, 123]. Even in the public goods setting, it may be possible to consider restricted utility models to obtain some incentive guarantees.

In the second part of the thesis, we considered the problem of eliciting and aggregating probabilistic forecasts. We considered the tradeoff between Pareto optimality and incentive compatibility in wagering mechanisms and proposed two possible solutions to the problem. We also considered prediction market design when there is no trusted center to close the markets, forcing us to rely on a vote among a group of untrusted entities with potential conflict of interest.

Again, there are many open questions. For many forecasting mechanisms, it is still unclear how they perform in practice both in an absolute sense and relative to one another. It would be interesting to examine, for instance, whether user satisfaction is higher with a Pareto optimal wagering mechanism than with a wagering mechanism that suffers from low stakes. It is also unclear whether, in real situations, participants have access to enough information to allow them to profitably misreport.

Appendix A

Omitted Proofs and Results for Chapter 2

A.1 Relationships Among Fairness Axioms

In this section, we analyze the relationship between the fairness properties we introduce in this paper, namely RRS, PPS, and Prop1. First, it is easy to show that Prop1 does not give any approximation to RRS or PPS, both for public decisions and for private goods, because it is easy to construct examples where a player receives zero utility, still satisfies Prop1, but has non-zero PPS share.

In the other direction, for public decisions, we showed that RRS implies 1/2-Prop1 (Lemma 3). For private goods, we can refine this result a bit further.

Theorem 85. *For private goods division, RRS implies Prop1 if and only if $m \leqslant 4n - 2$.*

Proof. First, let us assume $m > 4n - 2$. Consider the following fair division instance. Player 1's values, in the descending order, are as follows.

$$
u_{max}^{j}(1) = \begin{cases} (k-1) \cdot n + 1 & \text{if } j = 1, \\ n & \text{if } j \in \{2, \ldots, k \cdot n - 1\}, \\ 1 & \text{if } j \geqslant k \cdot n. \end{cases}
$$

It is easy to check that player 1's RRS share is exactly $u_{max}^1(1)$. Consider an allocation in which player 1 receives only his most valuable good, and the remaining goods are partitioned among the other players arbitrarily. For the sake of completeness, let each other player have value 1 for each good he receives under this allocation, and 0 for the remaining goods. Hence, the allocation satisfies RRS.

Now, player 1's proportional share is given by

$$\frac{(k-1) \cdot n + 1 + (kn-2) \cdot n + (m-kn+1) \cdot 1}{n} = \frac{m + kn^2 - 3n + 2}{n}$$
$$> \frac{n + kn^2}{n}$$
$$= kn + 1,$$

where the second transition follows because $m > 4n - 2$.

The highest value that player 1 can achieve by adding one more good to his allocation is $(k-1) \cdot n + 1 + n = kn + 1$, which falls short of the proportional share. Hence, the allocation is not Prop1.

Now, let us assume that $m \leqslant 4n - 2$. Hence, $k \leqslant 3$. Take a fair division instance, and let us focus on a player i. For the sake of notational convenience, we define $u_{max}^j(i) = 0$ for $j \in \{m+1, \ldots, 4n-2\}$. Note that this affects neither his RRS share nor his satisfaction of Prop1.

We now show that if player i receives at least as much value as his RRS share $u_{max}^n(1) + u_{max}^{2n}(1) + u_{max}^{3n}(1)$, then we can make player i receive his proportional share by adding a single good to his allocation.

If player i does not receive his most valuable good, then this can be accomplished

by adding his most valuable good to his allocation because

$$u_{max}^1(i) + u_{max}^n(1) + u_{max}^{2n}(1) + u_{max}^{3n}(1)$$

$$\geqslant \frac{\sum_{j=1}^n u_{max}^j(i)}{n} + \frac{\sum_{j=n+1}^{2n} u_{max}^j(i)}{n} + \frac{\sum_{j=2n+1}^{3n} u_{max}^j(i)}{n} + \frac{\sum_{j=3n+1}^{4n-2} u_{max}^j(i)}{n}$$

$$\geqslant \frac{\sum_{j=1}^m u_{max}^j(i)}{n}.$$

The first transition follows because $u_{max}^j(i) \geqslant v_{max}^{j+1}(i)$ for all $j \in [4n-3]$.

Suppose player i receives his most valuable good. Let t be the smallest index such that player i does not receive his t^{th} most valuable good. Hence, $t \geqslant 2$. Let u_i denote the utility to player i under the current allocation. Then, we have that

$$u_i \geqslant \sum_{j=1}^{t-1} u_{max}^j(i). \tag{A.1}$$

$$u_i \geqslant v_{max}^n(i) + u_{max}^{2n}(i) + v_{max}^{3n}(i). \tag{A.2}$$

Multiplying Equation (A.1) by $1/n$ and Equation (A.2) by $(n-1)/n$, and adding, we get

$$u_i \geqslant \frac{\sum_{j=1}^{t-1} u_{max}^j(i)}{n} + \frac{n-1}{n} \cdot (u_{max}^n(i) + u_{max}^{2n}(i) + u_{max}^{3n}(i))$$

$$\geqslant \frac{\sum_{j=1}^{t-1} u_{max}^j(i) + \sum_{j=n+1}^{2n-1} u_{max}^j(i) + \sum_{j=2n+1}^{3n-1} u_{max}^j(i) + \sum_{j=3n}^{4n-2} u_{max}^j(i)}{n}$$

$$= \frac{\sum_{j=1}^{t-1} u_{max}^j(i) + \sum_{j=n+1}^{4n-2} u_{max}^j(i) - u_{max}^{2n}(i)}{n}$$

$$\geqslant \frac{\sum_{j=1}^{t-1} u_{max}^j(i) + \sum_{j=n+2}^{4n-2} u_{max}^j(i)}{n}.$$

If $t \geqslant n+2$, player i already receives his proportional share. Otherwise, let us now add player i's t^{th} most valuable good to his allocation. His utility increases to

$$u_i + u_{max}^t(i) \geqslant \frac{\sum_{j=1}^{t-1} u_{max}^j(i) + \sum_{j=n+2}^{4n-2} u_{max}^j(i)}{n} + \frac{\sum_{j=t}^{n+1} u_{max}^j(i)}{n} = \frac{\sum_{j=1}^m u_{max}^j(i)}{n},$$

214

where the first transition follows because $t \geqslant 2$. Thus, player i receives his proportional share after adding a single good to his allocation. Because player i was chosen arbitrarily, we have that the allocation satisfies Prop1. ∎

A.2 Proof of Lemma 5

We will say that a set $\{x_1, \ldots, x_n\}$ of n non-negative real numbers is *feasible* for a given $0 < \delta < 1$ if $\sum_{k=1}^n \max\{0, 1 - x_k\} \leqslant \delta$.

Let $X = \{x_1, \ldots, x_n\}$ be a feasible set and let $i = \arg\min_{k \in [n]}\{x_k\}$, so that x_i is the minimum value in X. Suppose that $x_i > 1 - \delta$. We will show that there necessarily exists another feasible set $X' = \{x'_1, \ldots, x'_n\}$ with $\prod_{i=1}^n x'_i < \prod_{i=1}^n x_i$.

To that end, let $j = \arg\min_{k \in [n] \setminus \{i\}}\{x_k\}$, so that x_j is the second-lowest value in X. If $x_j \geqslant 1$ then set $x'_i = 1 - \delta$ and $x'_k = x_k$ for all $k \in [n] \setminus \{i\}$. Clearly this set X' is feasible since

$$\sum_{k=1}^n \max\{0, 1 - x_k\} = 1 - x_i = \delta,$$

and

$$\prod_{k=1}^n x'_k = x'_i \prod_{k \neq i} x'_k < x_i \prod_{k \neq i} x_k = \prod_{k=1}^n x_k.$$

Now consider the case where $x_j < 1$. Let $\varepsilon = \frac{1}{2}(1 - x_j)$ and define $x'_i = x_i - \varepsilon$,

$x'_j = x_j + \varepsilon$, and $x'_k = x_k$ for all $k \in [n]\backslash\{i,j\}$.. For feasibility of X',

$$\sum_{k=1}^{n} \max\{0, 1 - x'_k\} = \sum_{k \in [n]\backslash\{i,j\}} \max\{0, 1 - x'_k\} + \max\{0, 1 - x_i + \varepsilon\} + \max\{0, 1 - x_i + \varepsilon\}$$

$$= \sum_{k \in [n]\backslash\{i,j\}} \max\{0, 1 - x_k\} + (1 - x_i + \varepsilon) + (1 - x_j - \varepsilon)$$

$$= \sum_{k \in [n]\backslash\{i,j\}} \max\{0, 1 - x_k\} + (1 - x_i) + (1 - x_j)$$

$$= \sum_{k=1}^{n} \max\{0, 1 - x_k\},$$

and

$$\prod_{k=1}^{n} x'_k = x'_i x'_j \prod_{k \neq i,j} x'_k = (x_i - \varepsilon)(x_j + \varepsilon) \prod_{k \neq i,j} x_k = (x_i x_j - (x_j - x_i)\varepsilon - \varepsilon^2)$$

$$< x_i x_j \prod_{k \neq i,j} x_k = \prod_{k=1}^{n} x_k.$$

To complete the proof, it remains to show that any feasible set with smallest entry $x_i = 1 - \delta$ has $\prod_{k=1}^{n} x_k \geq 1 - \delta$. Note that if $x_i = 1 - \delta$ then for X to be feasible it can not be the case that $x_j < 1$. Therefore,

$$\prod_{k=1}^{n} x_k = (1 - \delta) \prod_{k \neq i} x_k \geq (1 - \delta).$$

∎

A.3   Proof of Theorem 14

First, we prove that Algorithm 1 produces an allocation satisfying PPS and PO. In particular, we prove that at the end, allocation $\mathbf{A}$ satisfies (A) $|A_i| \geq p$ for every player $i \in N$, and (B) $\mathbf{A}$ maximizes the weighted social welfare $\sum_{i \in N} w_i \cdot u_i(A_i)$. Note

216

that property (A) implies that **A** satisfies PPS, and property (B) implies that **A** is PO.

Before we prove these claims, we note that **A** maximizes the weighted social welfare if and only if it allocates each good $g \in M$ to a player i maximizing $w_i \cdot u_i(g)$.

First, we prove claim (A). Because the outer loop (Lines 3-22) explicitly runs until each player receives at least p goods, we simply need to prove that the loop terminates. For this, we need to analyze the first inner loop (Lines 8-14) and the second inner loop (Lines 16-21).

Each iteration of the first inner loop reduces the weights of players in DEC, and adds a player $j^* \in N \backslash DEC$ to DEC. Because $|DEC|$ increases by one in each iteration, the first inner loop terminates after $O(n)$ iterations.

The second inner loop starts from the player $j^* \in LS$ that was added to DEC at the end of the first inner loop, and traces back to the player i^* that was about to lose good g^* to j^* when j^* was added to DEC. The good is explicitly transferred, and if player i^* had exactly p goods initially, the algorithm continues to find another good to give back to player i^* by tracing back to the conditions under which player i^* was added to DEC. This way, the transfers add a good to a player who was in LS, maintain the number of goods of the players who were in EQ, and remove a good from a player who was in GT. Because in each iteration, player i^* had to be present in DEC before player j^* was added, this loop cannot continue indefinitely, and must terminate in $O(n)$ iterations as well.

Thus, both inner loops terminate in $O(n)$ iterations, and in each iteration of the outer loop, a player in LS receives an additional good without any new players being added to LS. This monotonically reduces the metric $\sum_{i \in LS} p - |A_i|$ by at least 1 in each iteration. Because this metric can be at most $p \cdot n \leqslant m$ to begin with, the outer loop executes $O(m)$ times.

We now prove claim (B). That is, we want to show that allocation **A** remains a

maximizer of the weighted social welfare according to weight vector $\mathbf{w}$ at during the execution of the algorithm. Because we specifically select $\mathbf{A}$ to be a weighted social welfare maximizer in Line 2, we simply show that neither the weight update in the first inner loop nor the changes to the allocation in the second inner loop violate this property.

In the first inner loop, because the weights of players in DEC are reduced by the same multiplicative factor, goods can only transfer from players in DEC to players in $N\backslash DEC$. However, the choice of r in Line 10 ensures that the weight reduction stops when the first such potential transfer creates a tie, preserving allocation $\mathbf{A}$ as a weighted welfare maximizer. Alterations to allocation $\mathbf{A}$ during the second inner loop also do not violate this property because this loop only transfers a good g^* from player i^* to player j^* when the two players were anyway tied to receive the good.

This concludes our claim that the algorithm terminates, and correctly produces an allocation satisfying PPS and PO. We already established that the outer loop executes $O(m)$ times, and the two inner loops execute $O(n)$ times. The bottleneck within the inner loops is the $\arg\min$ computation in Line 9, which requires $O(n \cdot m)$ time to find the minimum across all goods owned by players in DEC and all players outside DEC. Consequently, the asymptotic running time complexity of the algorithm is $O(m \cdot n \cdot nm) = O(n^2 \cdot m^2)$. $\blacksquare$

Appendix B

Omitted Proofs and Results for Chapter 4

B.1   Omitted Proofs

B.1.1   Proof of Lemma 29

We use the characterization of the FL mechanism allocations from Lemma 28. We consider four cases, corresponding to whether or not supply exceeds demand in the truthful and misreported instances. Let x' denote the objective value in the FL mechanism's call to PSWC in the misreported instance, and x in the truthful instance. Suppose first that $D_r \geqslant E$ and $D'_r \geqslant E$. Suppose that $x' \leqslant x$. Then, for all $j \neq i$,

$$a'_{j,r} = \min(x'e_j, d_{j,r}, t'_{j,r}) \leqslant \min(xe_j, d_{j,r}, t_{j,r}) = a_{j,r},$$

which implies that $a'_{i,r} \geqslant a_{i,r}$, since $\sum_{k \in [n]} a_{k,r'} = \sum_{k \in [n]} a'_{k,r'}$. On the other hand, if $x' > x$, then

$$a'_{i,r} = \min(x'e_i, d_{i,r}, t'_{i,r}) \geqslant \min(xe_i, d_{i,r}, t_{i,r}) = a_{i,r}.$$

Second, suppose that $D_r \geqslant E$ and $D'_r < E$. Then

$$a'_{i,r} \geqslant \min(d_{i,r}, t'_{i,r}) \geqslant \min(d_{i,r}, t_{i,r}) \geqslant a_{i,r}.$$

Third, suppose that $D_r < E$ and $D'_r \geqslant E$. Then

$$a'_{j,r} \leqslant \min(d_{j,r}, t'_{j,r}) \leqslant \min(d_{j,r}, t_{j,r}) \leqslant a_{j,r}$$

for all $j \neq i$, which implies that $a'_{i,r} \geqslant a_{i,r}$. Finally, suppose that $D_r < E$ and $D'_r < E$. If $x' \leqslant x$, then for all $j \neq i$, we have that

$$a'_{j,r} = \min(t'_{j,r}, \max(d_{j,r}, x'e_j)) \leqslant \min(t_{j,r}, \max(d_{j,r}, xe_j)) = a_{j,r},$$

which implies that $a'_{i,r} \geqslant a_{i,r}$. If $x' > x$, then

$$a'_{i,r} = \min(t'_{i,r}, \max(d_{i,r}, x'e_i)) \geqslant \min(t_{i,r}, \max(d_{i,r}, xe_i)) = a_{i,r}.$$

Thus, the lemma holds in all cases.

B.1.2 Proof of Lemma 41

If $D \geqslant E$, substituting the relevant terms into Lemma 27 gives us the following.

$$a_{i,r} = \max(0, \min(\min(d'_{i,r}, e_i + b_i), xe_i)) = \min(e_i + b_i, d'_{i,r}, xe_i).$$

If $D < E$, then again by substituting into Lemma 27 we have the following.

$$a_{i,r} = \max(\min(e_i + b_i, d'_{i,r}), \min(e_i + b_i, xe_i)) = \min(e_i + b_i, \max(d'_{i,r}, xe_i)).$$

The final equality, $\max(\min(A, B), \min(A, C)) = \min(A, \max(B, C))$, can easily be checked to hold case by case for any relative ordering of A, B, and C.

B.1.3 Proof of Lemma 42

If $r > T$, then the allocation of agent i is independent of her reported demand, thus $a_{i,r} = a'_{i,r}$. Now suppose that $r \leqslant T$. Let $\bar{d}_{i,r} = \min(d_{i,r}, e_i + b_{i,r})$ and $\bar{d}'_{i,r} = \min(d'_{i,r}, e_i + b_{i,r})$. Also, let x and x' denote the objective value in the T-period mechanism's call to PSWC when i reports $d_{i,r}$ and $d'_{i,r}$, respectively. Observe first that $D' = \bar{d}'_{i,r} + \sum_{j \neq i} \bar{d}_{j,r} \leqslant \bar{d}_{i,r} + \sum_{j \neq i} \bar{d}_{j,r} = D$.

Suppose first that $E \leqslant D' \leqslant D$. Let $a_{j,r}$ and $a'_{j,r}$ denote the allocations of $j/not = i$ when i reports $d_{i,r}$ and $d'_{i,r}$, respectively. If $x' \geqslant x$, then for all $j \neq i$, by Lemma 41 we have:

$$a'_{j,r} = \min(e_j + b_{j,r}, d_{j,r}, x'e_j) \geqslant \min(e_j + b_{j,r}, d_{j,r}, xe_j) = a_{j,r}.$$

This immediately implies that $a'_{i,r} \leqslant a_{i,r}$, because $\sum_{k \in [n]} a_{k,r} = \sum_{k \in [n]} a'_{k,r} = E$. If $x' < x$, then again by Lemma 41 we have the following:

$$a'_{i,r} = \min(e_i + b_{i,r}, d'_{i,r}, x'e_i) \leqslant \min(e_i + b_{i,r}, d_{i,r}, xe_i) = a_{i,r}.$$

By the same lemma, for all $j \neq i$, we also have:

$$a'_{j,r} = \min(e_j + b_{j,r}, d_{j,r}, x'e_j) \leqslant \min(e_j + b_{j,r}, d_{j,r}, xe_j) = a_{j,r}.$$

Therefore, for all $k \in [n]$, $a_{k,r} \geqslant a'_{k,r}$. However, since $\sum_{k \in [n]} a_{k,r} = \sum_{k \in [n]} a'_{k,r} = E$, it has to be the case that $a_{k,r} = a'_{k,r}$ for all k.

Next, suppose that $D' < E \leqslant D$. By the definition of the T-period mechanism, for all $j \neq i$, $a'_{j,r} \geqslant \bar{d}_{j,r}$, and $a_{j,r} \leqslant \bar{d}_{j,r}$. Therefore, $a'_{j,r} \geqslant a_{j,r}$ which implies that $a'_{i,r} \leqslant a_{i,r}$.

Finally, suppose that $D' \leqslant D < E$. If $x' \geqslant x$, then by Lemma 41, for all $j \neq i$, we have:

$$a'_{j,r} = \min(e_j + b_{j,r}, \max(d_{j,r}, x'e_j)) \geqslant \min(e_j + b_j, \max(d_{j,r}, xe_j)) = a_{j,r}.$$

This implies $a'_{i,r} \leqslant a_{i,r}$. If $x' < x$, then, by Lemma 41 we have:

$$a'_{i,r} = \min(e_i + b_{i,r}, \max(d'_{i,r}, x'e_i)) \leqslant \min(e_i + b_{i,r}, \max(d_{i,r}, xe_i)) = a_{i,r}$$

By the same lemma, for all $j \neq i$, we also have:

$$a'_{j,r} = \min(e_j + b_{j,r}, \max(d_{j,r}, x'e_j)) \leqslant \min(e_j + b_{j,r}, \max(d_{j,r}, xe_j)) = a_{j,r}.$$

Therefore, for all $k \in [n]$, $a'_{k,r} \leqslant a_{k,r}$. However, since $\sum_{k \in [n]} a_{k,r} = \sum_{k \in [n]} a'_{k,r} = E$, it has to be the case that $a_{k,r} = a'_{k,r}$ for all k.

B.1.4 Proof of Lemma 43

Note that $D' \leqslant D$, since $d'_{i,r} < d_{i,r}$ and $d'_{j,r} = d_{j,r}$ for all agents $j \neq i$. If $E \leqslant D$, then by the definition of the T-period mechanism we have:

$$a_{i,r} \leqslant \bar{d}_{i,r} = \min(e_i + b_i, d_{i,r}) \leqslant d_{i,r}.$$

Next, assume that $D' \leqslant D < E$. Then $a'_{i,r} < a_{i,r}$ implies that there is at least one agent j with $a'_{j,r} > a_{j,r}$. In the proof of Lemma 42 we show that if $x' < x$, then $a'_{k,r} = a_{k,r}$ for all k. Therefore, it has to be the case that $x' \geqslant x$. By Lemma 41, $a_{i,r} = \min(e_i + b_{i,r}, \max(d_{i,r}, xe_i))$ and $a'_{i,r} = \min(e_i + b_i, \max(d'_{i,r}, x'e_i))$. It is easy to see that if $d_{i,r} < xe_i$, then $a'_{i,r} \geqslant a_{i,r}$, which contradicts the assumption in the lemma statement. Therefore, we have:

$$a_{i,r} = \min(e_i + b_i, \max(d_{i,r}, xe_i)) = \min(e_i + b_i, d_{i,r}) \leqslant d_{i,r}.$$

B.1.5 Proof of Corollary 44

Because $a'_{i,r} < a_{i,r} \leqslant d_{i,r}$, we can substitute the utility values from Equation (4.2),

$$u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r}) = a_{i,r}H - a'_{i,r}H = H(a_{i,r} - a'_{i,r}).$$

B.1.6 Proof of Lemma 45

Note that $D' \geqslant D$, since $d'_{i,r} > d_{i,r}$ and $d'_{j,r} = d_{j,r}$ for all agents $j \neq i$. If $D < E$, then $a_{i,r} \geqslant \bar{d}_{i,r} = \min(e_i + b_{i,r}, d_{i,r})$. We show that $e_i + b_{i,r} \geqslant d_{i,r}$, and therefore, $a_{i,r} \geqslant d_{i,r}$. Suppose for contradiction that $e_i + b_{i,r} < d_{i,r} < d'_{i,r}$, which means $\bar{d}_{i,r} = \bar{d}'_{i,r} = e_i + b_{i,r}$. By definition of the T-period mechanism, $\bar{d}_{i,r} \leqslant a_{i,r} \leqslant e_i + b_{i,r}$, which implies $a_{i,r} = e_i + b_{i,r}$. Also, by the definition of the mechanism, $a'_{i,r} \leqslant \bar{d}'_{i,r} = e_i + b_{i,r}$ if $D' \geqslant E$, and $a'_{i,r} \leqslant e_i + b'_{i,r} = e_i + b_{i,r}$ if $D' < E$. In both cases, $a'_{i,r} \leqslant e_i + b_{i,r} = a_{i,r}$, a contradiction to the assumption in the lemma statement.

If $D' \geqslant D \geqslant E$, then $a'_{i,r} > a_{i,r}$ implies that there is at least an agent j with $a'_{j,r} < a_{j,r}$. In the proof of Lemma 42 we show that if $x < x'$, then $a'_{k,r} = a_{k,r}$

for all k. Therefore, it has to be the case that $x \geqslant x'$. By Lemma 41, $a_{i,r} = \min(e_i + b_{i,r}, d_{i,r}, xe_i)$ and $a'_{i,r} = \min(e_i + b_{i,r}, d'_{i,r}, x'e_i)$. It is easy to see that if $a_{i,r}$ is xe_i or $e_i + b_{i,r}$, then $a'_{i,r} \leqslant a_{i,r}$. Therefore, $a_{i,r} = d_{i,r}$, which means the lemma holds.

B.1.7 Proof of Corollary 46

Because $d_{i,r} \leqslant a_{i,r} < a'_{i,r}$, we can substitute the utility values from Equation (4.2),

$$u_{i,r}(a'_{i,r}) - u_{i,r}(a_{i,r}) = d_{i,r}H + (a'_{i,r} - d_{i,r})L - (d_{i,r}H + (a_{i,r} - d_{i,r})L) = L(a'_{i,r} - a_{i,r}).$$

B.1.8 Proof of Lemma 47

Suppose first that agent i reports $d'_{i,T} < d_{i,T}$. Then, by Lemma 42, $a'_{i,T} \leqslant a_{i,T}$. If $a'_{i,T} = a_{i,T}$, then the misreport has had no effect on the allocations, since the allocation at rounds $r \leqslant T$ is unchanged, and the allocations at rounds $r > T$ depend only on the allocations at rounds $r \leqslant T$, not the reports. So assume that $a'_{i,T} = a_{i,T} - k$ for some $k > 0$. By the definition of the T-Period mechanism, i's allocation increases by $\frac{k}{T}$ for each of rounds $T + 1, \ldots, 2T$. The difference between her utility from truthfully reporting at round T and from misreporting at round T is given by

$$U_{i,R}(\mathbf{a_i}) - U_{i,R}(\mathbf{a'_i}) = \sum_{r=1}^{R} \left(u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r}) \right)$$

$$= u_{i,T}(a_{i,T}) - u_{i,T}(a'_{i,T}) + \sum_{r=T+1}^{2T} \left(u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r}) \right)$$

$$= kH + \sum_{r=T+1}^{2T} \left(u_{i,r}(a_{i,r}) - u_{i,r}(a_{i,r} + \frac{k}{T}) \right) \geqslant kH - kH = 0$$

where the second transition follows because $d'_{i,r} = d_{i,r}$ for all rounds $r < T$, the third transition from Corollary 44, and the final transition because each of the extra resources received in the misreported case for rounds $r > T$ can each be worth at most H to i.

223

Next suppose that agent i reports $d'_{i,T} > d_{i,T}$. Then, by Lemma 42, $a'_{i,T} \geqslant a_{i,T}$. As before, assume that $a'_{i,T} \neq a_{i,T}$. That is, $a'_{i,T} = a_{i,T} + k$ for some $k > 0$. By the definition of the T-Period mechanism, i's allocation decreases by $\frac{k}{T}$ for each of rounds $T + 1, \ldots, 2T$. The difference between her utility from truthfully reporting at round T and from misreporting at round T is given by

$$U_{i,R}(\mathbf{a_i}) - U_{i,R}(\mathbf{a'_i}) = \sum_{r=1}^{R} \left(u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r}) \right)$$

$$= u_{i,T}(a_{i,T}) - u_{i,T}(a'_{i,T}) + \sum_{r=T+1}^{2T} \left(u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r}) \right)$$

$$= -kL + \sum_{r=T+1}^{2T} \left(u_{i,r}(a_{i,r}) - u_{i,r}(a_{i,r} - \frac{k}{T}) \right) \geqslant -kL + kL = 0$$

where the second transition follows because $d'_{i,r} = d_{i,r}$ for all rounds $r < T$, the third transition from Corollary 46, and the final transition because each of the extra resources received in the truthful case for rounds $r > T$ are each worth at least L to i.

B.1.9 Proof of Lemma 49

We treat four cases, corresponding to whether or not supply exceeds demand in the truthful and misreported instances. Let x' denote the objective value in the T-Period mechanism's call to PSWC in the misreported instance, and x in the truthful instance. All cases rely heavily on the characterization of the allocation from Lemma 41.

Suppose first that $D_r \geqslant E$ and $D'_r \geqslant E$. Suppose that $x' \leqslant x$. Then, for all $j \neq i$, $a'_{j,r} = \min(x'e_j, d_{j,r}, e_j + b'_{j,r}) \leqslant \min(xe_j, d_{j,r}, e_j + b_{j,r}) = a_{j,r}$, which implies that $a'_{i,r} \geqslant a_{i,r}$, since $\sum_{k \in [n]} a_{k,r'} = \sum_{k \in [n]} a'_{k,r'}$. On the other hand, if $x' > x$, then $a'_{i,r} = \min(x'e_i, d_{i,r}, e_i + b'_{i,r}) \geqslant \min(xe_i, d_{i,r}, e_i + b_{i,r}) = a_{i,r}$. Second, suppose that $D_r \geqslant E$ and $D'_r < E$. Then $a'_{i,r} \geqslant \min(d_{i,r}, e_i + b'_{i,r}) \geqslant \min(d_{i,r}, e_i + b_{i,r}) \geqslant a_{i,r}$.

Third, suppose that $D_r < E$ and $D'_r \geqslant E$. Then $a'_{j,r} \leqslant \min(d_{j,r}, e_j + b'_{j,r}) \leqslant \min(d_{j,r}, e_j + b_{j,r}) \leqslant a_{j,r}$ for all $j \neq i$, which implies that $a'_{i,r} \geqslant a_{i,r}$.

Finally, suppose that $D_r < E$ and $D'_r < E$. If $x' \leqslant x$, then for all $j \neq i$, we have that $a'_{j,r} = \min(e_j + b'_{j,r}, \max(d_{j,r}, x'e_j)) \leqslant \min(e_j + b_{j,r}, \max(d_{j,r}, xe_j)) = a_{j,r}$, which implies that $a'_{i,r} \geqslant a_{i,r}$. If $x' > x$, then $a'_{i,r} = \min(e_i + b'_{i,r}, \max(d_{i,r}, x'e_i)) \geqslant \min(e_i + b_{i,r}, \max(d_{i,r}, xe_i)) = a_{i,r}$. Thus, the lemma holds in all cases.

B.1.10 Proof of Proposition 51

Let $r \leqslant T$. First suppose that $D < E$. Then i's minimum allocation is $\bar{d}_{i,r} = \min(d'_{i,r}, e_i + b_{i,r}) = e_i$. So we know that $a_{i,r} \geqslant e_i$. Suppose for contradiction that $a_{i,r} > e_i$. Then there must be some agent $j \neq i$ with $a_{j,r} \leqslant e_j$. But now we could obtain a smaller value of x in the PSWC program by assigning slightly higher allocation to j, and slightly lower allocation to any agent with $a_{k,r}/e_k = x$ (we know that j is not one of these agents since $a_{j,r}/e_j < 1 < a_{i,r}/e_i \leqslant x$). This contradicts optimality of the PSWC program, therefore $a_{i,r} = e_i$.

Next, suppose that $D \geqslant E$. Then i's limit allocation is $\bar{d}_{i,r} = \min(d'_{i,r}, e_i + b_{i,r}) = e_i$. So we know that $a_{i,r} \leqslant e_i$. Suppose for contradiction that $a_{i,r} < e_i$. Then there must exist some agent j with $a_{j,r} > e_j$. But now the objective value x of the call to PSWC could be improved by transferring some small amount of allocation to i from all agents k with $a_{k,r}/e_k = x$ (we know that i is not one of these agents since $a_{i,r}/e_i < 1 < a_{j,r}/e_j \leqslant x$). This contradicts optimality of the PSWC program, therefore $a_{i,r} = e_i$.

B.2 Over-reporting Demand is not Advantageous

In this section we assume that $d'_{i,r'} > d_{i,r'}$. The setup otherwise mirrors that of §4.5.3.

Lemma 86. *For all agents $j \neq i$, we have that $a'_{j,r'} \leqslant a_{j,r'}$. Further, $a'_{i,r'} \geqslant a_{i,r'}$.*

Proof. We prove the statement for all $j \neq i$. The statement for i follows immediately because the total number of resources to allocate is fixed.

Observe first that

$$D_{r'} = \sum_{k \in [n]} \min(d_{k,r'}, t_{k,r'}) \leqslant \sum_{k \in [n]} \min(d'_{k,r'}, t_{k,r'}) = D'_{r'},$$

since i's demand increases in the misreported instances but all other demands and token counts stay the same. Let x' denote the objective value in FL's call to PSWC in the misreported instance, and x in the truthful instance.

Suppose that $E \leqslant D_{r'} \leqslant D'_{r'}$. Suppose first that $x' < x$. Then, by Lemma 28,

$$a_{j,r'} = \min(xe_j, d_{j,r'}, t_{j,r'}) \geqslant \min(x'e_j, d_{j,r'}, t_{j,r'}) = a'_{j,r'}$$

for all $j \neq i$. Next, suppose that $x' \geqslant x$. Then, again by Lemma 28 and the fact that $d'_{i,r'} > d_{i,r'}$,

$$a'_{i,r'} = \min(x'e_i, d'_{i,r'}, t_{i,r'}) \geqslant \min(xe_i, d_{i,r'}, t_{i,r'}) = a_{i,r'}.$$

And, for all $j \neq i$,

$$a'_{j,r'} = \min(x'e_j, d_{j,r'}, t_{j,r'}) \geqslant \min(xe_j, d_{j,r'}, t_{j,r'}) = a_{j,r'}.$$

Because $a'_{k,r'} \geqslant a_{k,r'}$ for all users k, and $\sum_{k \in [n]} a_{k,r'} = \sum_{k \in [n]} a'_{k,r'}$, it must be the case that $a'_{k,r'} = a_{k,r'}$ for all k, which satisfies the statement of the lemma.

Next, suppose that $D_{r'} < E \leqslant D'_{r'}$. By the definition of FL, $a_{k,r'} \geqslant \min(d_{k,r'}, t_{k,r'})$ for all k, and $a'_{k,r'} \leqslant \min(d'_{k,r'}, t_{k,r'})$ for all k. Since $\min(d'_{j,r'}, t_{j,r'}) = \min(d_{j,r'}, t_{j,r'})$ for all $j \neq i$, we have that $a_{j,r'} \geqslant a'_{j,r'}$, implying also that $a_{i,r'} \leqslant a'_{i,r'}$.

Finally, suppose that $D_{r'} \leqslant D'_{r'} < E$. Suppose first that $x \leqslant x'$. Then, by Lemma 28 and the assumption that $d_{i,r'} < d'_{i,r'}$, we have

$$a_{i,r'} = \min(t_{i,r'}, \max(xe_i, d_{i,r'})) \leqslant \min(t_{i,r'}, \max(x'e_i, d'_{i,r'})) = a'_{i,r'}$$

226

and

$$a_{j,r'} = \min(t_{j,r'}, \max(xe_j, d_{j,r'})) \leqslant \min(t_{j,r'}, \max(x'e_j, d_{j,r'})) = a'_{j,r'}$$

for all $j \neq i$. Because $a_{k,r'} \leqslant a'_{k,r'}$ for all users k, and $\sum_{k \in [n]} a'_{k,r'} = \sum_{k \in [n]} a_{k,r'}$, it must be the case that $a_{k,r'} = a'_{k,r'}$ for all k, which satisfies the lemma statement. Next, suppose that $x > x'$. Then, again by Lemma 28, for all $j \neq i$, we have

$$a_{j,r'} = \min(t_{j,r'}, \max(xe_j, d_{j,r'})) \geqslant \min(t_{j,r'}, \max(x'e_j, d_{j,r'})) = a'_{j,r'}.$$

∎

If it is the case that $a'_{i,r'} = a_{i,r'}$, then it must also be that $a'_{j,r'} = a_{j,r'}$ for all $j \neq i$. So allocations at round r' are the same in the misreported instance as the truthful instance. Therefore, for all rounds $r \leqslant r'$, allocations in both universes are the same. In all rounds $r > r'$, reports in both universes are the same. Together, these imply that allocations for all rounds $r > r'$ are the same in both universes. In particular, i does not profit from her misreport and could weakly improve her utility by reporting $d'_{i,r'} = d_{i,r'}$. So, for the remainder of this section, we assume that $a'_{i,r'} > a_{i,r'}$.

Our next lemma says that the additional resources that i receives in round r' are *low* valued resources for her. The intuition is that if it were the case that i was receiving only high-valued resources under truthful reporting, then she will not receive any extra resources by misreporting (since no agent donates any additional resources for i to receive).

Lemma 87. *If $a'_{i,r'} > a_{i,r'}$, then $a_{i,r'} \geqslant d_{i,r'}$.*

Proof. Suppose for contradiction that $a_{i,r'} < d_{i,r'}$. We also know that $a_{i,r'} < a'_{i,r'} \leqslant t'_{i,r'} = t_{i,r'}$, where the equality holds because allocations before round r' are identical in the truthful and misreported instances. It must therefore be the case that $D'_{r'} \geqslant D_{r'} > E$, where the first inequality holds because $d'_{j,r'} = d_{j,r'}$ for all $j \neq i$ and

227

$d'_{i,r'} > d_{i,r'}$, and the second because $a_{i,r'} < \min(t_{i,r'}, d_{i,r'})$. Let x denote the objective value of FL's call to PSWC in the truthful instance, and x' in the misreported instance. Suppose that $x \leqslant x'$. Then, by Lemma 28 and the assumption that $d_{i,r'} < d'_{i,r'}$,

$$a_{i,r'} = \min(t_{i,r'}, xe_i, d_{i,r'}) \leqslant \min(t_{i,r'}, x'e_i, d'_{i,r'}) = a'_{i,r'},$$

and for all $j \neq i$

$$a_{j,r'} = \min(t_{j,r'}, xe_j, d_{j,r'}) \leqslant \min(t_{j,r'}, x'e_j, d_{j,r'}) = a'_{j,r'}.$$

Because $a_{k,r'} \leqslant a'_{k,r'}$ for all agents k, and $\sum_{k\in[n]} a_{k,r'} = \sum_{k\in[n]} a'_{k,r'}$, it must be the case that $a'_{k,r'} = a_{k,r'}$ for all k. This contradicts the assumption that $a_{i,r'} < a'_{i,r'}$.

Now suppose that $x > x'$. Note that $xe_i < d_{i,r'} < d'_{i,r'}$, where the first inequality holds because $a_{i,r'} < \min(t_{i,r'}, d_{i,r'})$. Then, again by Lemma 28 and the previous observation, we have

$$a'_{i,r'} = \min(t_{i,r'}, x'e_i, d'_{i,r'}) \leqslant \min(t_{i,r'}, xe_i, d_{i,r'}) = a_{i,r'},$$

which contradicts that $a_{i,r'} < a'_{i,r'}$.

Since we arrive at a contradiction in all cases, the lemma statement must be true.

■

As a corollary, we can write the difference in utility between the truthful and misreported instances that i derives from round r'.

Corollary 88. $u_{i,r'}(a'_{i,r'}) - u_{i,r'}(a_{i,r'}) = L(a'_{i,r'} - a_{i,r'})$.

Proof. Because $d_{i,r'} \leqslant a_{i,r'} < a'_{i,r'}$, we can substitute the utility values from Equation (4.2):

$$u_{i,r'}(a'_{i,r'}) - u_{i,r'}(a_{i,r'}) = d_{i,r'}H + (a'_{i,r'} - d_{i,r'})L - d_{i,r'}H - (a'_{i,r} - d_{i,r'})L = L(a'_{i,r'} - a_{i,r'}).$$

■

228

For a fixed agent k, denote by r'_k the round at which agent k runs out of tokens in the misreported universe. That is, r'_k is the first (and only) round with $a'_{r_k} = t'_{k,r_k} > 0$. Note that $r'_i \geqslant r'$, since $a'_{i,r'} > 0$. Given this, our next lemma states that, under certain conditions, the effect of i's misreport, $d'_{i,r} > d_{i,r}$, is to decrease the objective value of FL's call to PSWC.

Lemma 89. *Let $r < r'_i$ (i.e., $a'_{i,r} < t'_{i,r}$). Suppose $t_{j,r} \leqslant t'_{j,r}$ for all agents $j \neq i$. Suppose that either $\min(D_r, D'_r) \geqslant E$ or $\max(D_r, D'_r) < E$. Then $x' \leqslant x$, where x' denotes the objective value of FL's call to PSWC in the misreported instance and x in the truthful instance.*

Proof. First, suppose that $\min(D_r, D'_r) \geqslant E$. Suppose for contradiction that $x < x'$. By Lemma 28,

$$a_{j,r} = \min(xe_j, d_{j,r}, t_{j,r}) \leqslant \min(x'e_j, d_{j,r}, t'_{j,r}) = a'_{j,r}$$

for all $j \neq i$, where the inequality follows from the assumption that $x < x'$ and that $t_{j,r} \leqslant t'_{j,r}$. Further,

$$a_{i,r} = \min(xe_i, d_{i,r}, t_{i,r}) \leqslant \min(xe_i, d_{i,r}) \leqslant \min(x'e_i, d_{i,r})$$
$$= \min(x'e_i, d_{i,r}, t'_{i,r}) = a'_{i,r},$$

where the second inequality follows from the assumption that $x < x'$ and the second to last equality from the assumption $a'_{i,r} < t'_{i,r}$.

Therefore, $a_{k,r} \leqslant a'_{k,r}$ for all agents k. Since $\sum a'_{k,r} = \sum a_{k,r}$, it must be the case that $a'_{k,r} = a_{k,r}$ for all agents k. Therefore, by the definition of FL, $a'_{k,r}/e_k \leqslant x < x'$ for all agents k with $a'_{k,r} > m_k = 0$. Therefore x' is not the optimal objective value of the PSWC program in the misreported instance, a contradiction. Thus, $x \geqslant x'$.

Next, suppose that $\max(D_r, D'_r) < E$. Suppose for contradiction that $x < x'$. By Lemma 28,

$$a_{j,r} = \min(t_{j,r}, \max(xe_j, d_{j,r})) \leqslant \min(t'_{j,r}, \max(x'e_j, d_{j,r})) = a'_{j,r}$$

for all $j \neq i$, where the inequality follows from the assumption that $x < x'$ and that $t_{j,r} \leqslant t'_{j,r}$. Further,

$$a_{i,r} = \min(t_{i,r}, \max(xe_i, d_{i,r})) \leqslant \max(xe_i, d_{i,r}) \leqslant \max(x'e_i, d_{i,r})$$
$$= \min(t'_{i,r}, \max(x'e_i, d_{i,r})) = a'_{i,r},$$

where the second inequality follows from the assumption that $x < x'$ and the second to last equality from the assumption $a'_{i,r} < t'_{i,r}$.

Therefore, $a_{k,r} \leqslant a'_{k,r}$ for all agents k. Since $\sum a'_{k,r} = \sum a_{k,r}$, it must be the case that $a'_{k,r} = a_{k,r}$ for all agents k. Consider all agents with $\min(d_{k,r}, t'_{k,r}) < a'_{k,r}$ (that is, those agents for which the first constraint in the PSWC program binds in the misreported instance). For all such agents, we have

$$\min(d_{k,r}, t'_{k,r}) < a'_{k,r} \implies d_{k,r} < a'_{k,r} \leqslant t'_{k,r} \implies d_{k,r} < a_{k,r} \leqslant t_{k,r} \implies \min(d_{k,r}, t_{k,r}) < a_{k,r},$$

so the constraints bind in the truthful instance as well. Therefore, $a_{k,r}/e_k \leqslant x < x'$ for all agents k for which the first constraint binds in the misreported instance. Therefore x' is not the optimal objective value of the PSWC program in the misreported instance, a contradiction. Thus, $x \geqslant x'$. $\blacksquare$

Using Lemma 89, we show our main lemma. It allows us to make an inductive argument that, after gaining some extra resources in round r', i's allocation is (weakly) smaller for all other rounds in the mireported instance than the truthful instance.

Lemma 90. *Let $r' < r < r'_i$ (that is, $a'_{i,r} < t'_{i,r}$). Suppose that $t_{j,r} \leqslant t'_{j,r}$ for all agents $j \neq i$. Then for all $j \neq i$, either: (1) $a_{j,r} = t_{j,r}$, or (2) $a_{j,r} \geqslant a'_{j,r}$.*

Proof. Note that $t_{j,r} \leqslant t'_{j,r}$ for all $j \neq i$ implies that $t_{i,r} \geqslant t'_{i,r}$, which we use in the proof. Also, because $r' < r$, we know that $d'_{i,r} = d_{i,r}$, as r' is the last round for which $d'_{i,r} \neq d_{i,r}$. We assume that condition 1) from the lemma statement is false (i.e., $a_{j,r} < t_{j,r}$) and show that condition 2) must hold. Suppose first that

230

$D_r' < E$. Then, because $a_{i,r}' < t_{i,r}'$, we know that $d_{i,r}' \leqslant t_{i,r}' \leqslant t_{i,r}$. This implies that $\min(d_{i,r}, t_{i,r}') = \min(d_{i,r}, t_{i,r}) = d_{i,r}$. Let $j \neq i$. Since $t_{j,r} \leqslant t_{j,r}'$, we have $\min(d_{j,r}, t_{j,r}) \leqslant \min(d_{j,r}, t_{j,r}')$. Therefore, it is the case that $D_r \leqslant D_r' < E$. By Lemma 28 and the assumption that $a_{j,r} < t_{j,r}$, it must be the case that $a_{j,r} = \max(d_{j,r}, xe_j)$. Further, by Lemma 89, we know that $x \geqslant x'$. Therefore, we have

$$a_{j,r}' = \max(d_{j,r}, x'e_j) \leqslant \max(d_{j,r}, xe_j) = a_{j,r}.$$

That is, condition (2) from the lemma statement holds.

Now suppose that $D_r' \geqslant E$. Then, from the definition of the mechanism, we have that $a_{j,r}' \leqslant \min(d_{j,r}, t_{j,r}') \leqslant d_{j,r}$. If it is the case that $D_r < E$ then we have that $a_{j,r} \geqslant \min(d_{j,r}, t_{j,r}) = d_{j,r}$, where the equality holds because otherwise we would have $a_{j,r} \geqslant \min(d_{j,r}, t_{j,r}) = t_{j,r}$, violating the assumption that $a_{j,r} < t_{j,r}$. Using these inequalities, we have $a_{j,r} \geqslant d_{j,r} \geqslant a_{j,r}'$, so condition (2) from the statement of the lemma holds. Finally, it may be the case that $D_r' \geqslant M$ and $D_r \geqslant M$. By Lemma 28 and the assumption that $a_{j,r} < t_{j,r}$, we have

$$a_{j,r} = \min(d_{j,r}, xe_k) \geqslant \min(d_{j,r}, x'e_k) = a_{j,r}',$$

where the inequality follows from Lemma 89. Thus, condition (2) of the lemma statement holds. ∎

We now show an analogous result to Lemma 29.

Lemma 91. *Suppose that $t_{j,r} \leqslant t_{j,r}'$ for all $j \neq i$, and $d_{k,r} = d_{k,r}'$ for all $k \in [n]$. Then $a_{i,r} \geqslant a_{i,r}'$.*

Proof. Note that the condition that $t_{j,r} \leqslant t_{j,r}'$ for all $j \neq i$ implies that $t_{i,r} \geqslant t_{i,r}'$. We use these assumptions, along with the characterization of the FL mechanism allocations from Lemma 28, to prove the lemma.

We treat four cases, corresponding to whether or not supply exceeds demand in the truthful and misreported instances. Let x' denote the objective value in the FL mechanism's call to PSWC in the misreported instance, and x in the truthful instance. Suppose first that $D'_r \geqslant E$ and $D_r \geqslant E$. Suppose that $x \leqslant x'$. Then, for all $j \neq i$, $a_{j,r} = \min(xe_j, d_{j,r}, t_{j,r}) \leqslant \min(x'e_jd_{j,r}, t'_{j,r}) = a'_{j,r}$, which implies that $a_{i,r} \geqslant a'_{i,r}$, since $\sum_{k \in [n]} a'_{k,r'} = \sum_{k \in [n]} a_{k,r'}$. On the other hand, if $x > x'$, then $a_{i,r} = \min(xe_i, d_{i,r}, t_{i,r}) \geqslant \min(x'e_i, d_{i,r}, t'_{i,r}) = a'_{i,r}$. Second, suppose that $D'_r \geqslant E$ and $D_r < E$. Then $a_{i,r} \geqslant \min(d_{i,r}, t_{i,r}) \geqslant \min(d_{i,r}, t'_{i,r}) \geqslant a'_{i,r})$. Third, suppose that $D'_r < E$ and $D_r \geqslant E$. Then $a_{j,r} \leqslant \min(d_{j,r}, t_{j,r}) \leqslant \min(d_{j,r}, t'_{j,r}) \leqslant a'_{j,r}$ for all $j \neq i$, which implies that $a_{i,r} \geqslant a'_{i,r}$.

Finally, suppose that $D'_r < E$ and $D_r < E$. If $x \leqslant x'$, then for all $j \neq i$, we have that $a_{j,r} = \min(t_{j,r}, \max(d_{j,r}, xe_j)) \leqslant \min(t'_{j,r}, \max(d_{j,r}, x'e_j)) = a'_{j,r}$, which implies that $a_{i,r} \geqslant a'_{i,r}$. If $x > x'$, then $a_{i,r} = \min(t_{i,r}, \max(d_{i,r}, xe_i)) \geqslant \min(t'_{i,r}, \max(d_{i,r}, x'e_i)) = a'_{i,r}$. Thus, the lemma holds in all cases. ∎

Finally, we show that the mechanism is strategy-proof.

Theorem 92. *Agent i never benefits from reporting $d_{i,r'} > d_{i,r'}$.*

Proof. We first observe that for every $r \leqslant r'_i$, $t_{j,r} \leqslant t'_{j,r}$ for every $j \neq i$. This is true for every $r \leqslant r'$ because $a'_{j,r} = a_{j,r}$ for $r < r'$, by Lemma 30. For $r = r' + 1$, it follows from Lemma 86, which says that $a_{j,r'} \geqslant a'_{j,r'}$. For all subsequent rounds, up to and including $r = r'_i$, it follows inductively from Lemma 90: $t_{j,r} \leqslant t'_{j,r}$ implies that either $a_{j,r} = t_{j,r}$ (in which case $t_{j,r+1} = 0 \leqslant t'_{j,r+1}$), or $a_{j,r} \geqslant a'_{j,r}$ (in which case $t_{j,r+1} = t_{j,r} - a_{j,r} \leqslant t'_{j,r} - a'_{j,r} = t'_{j,r+1}$).

Consider an arbitrary round $r \neq r'$, with $r \leqslant r'_i$. By the above argument, we know that $t_{j,r} \leqslant t'_{j,r}$ for all $j \neq i$. Further, because reports in the truthful and misreported instances are identical on all rounds $r \neq r'$, we have that $d_{k,r} = d'_{k,r}$ for

all $k \in [n]$. Therefore, by Lemma 91, $a_{i,r} \geqslant a'_{i,r}$. For rounds $r > r'_i$, it is also true that $a_{i,r} \geqslant a'_{i,r}$, since $a'_{i,r} = 0$ for these rounds by the definition of r'_i.

Finally,

$$U_{i,R}(\mathbf{a_i}) - U_{i,R}(\mathbf{a'_i}) = \sum_{r=1}^{R} (u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r}))$$

$$= \sum_{r \neq r'} (u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r})) + (u_{i,r'}(a_{i,r'}) - u_{i,r'}(a'_{i,r'}))$$

$$= \sum_{r \neq r'} (u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r})) - L(a'_{i,r'} - a_{i,r'})$$

$$\geqslant L(a'_{i,r'} - a_{i,r'}) - L(a'_{i,r'} - a_{i,r'}) = 0$$

Where the third transition follows from Corollary 88, and the final transition because $\sum_{r \neq r'}(a'_{i,r} - a_{i,r}) = a_{i,r'} - a'_{i,r'}$, and every term in the sum is positive. $\blacksquare$

Appendix C

Omitted Material from Chapter 5

C.1 Proof of Theorem 54

In this section, we provide a proof of Theorem 54. The first step of the proof is to show that any time a profitable side bet exists, there is a profitable side bet with $\Delta \sigma_i = 0$ for all agents i. This is because \$1 in cash is equivalent to a pair of `yes` and `no` securities. Thus we can limit attention to side bets that only involve the exchange of securities. The second step uses this fact to show that any time a profitable side bet exists, there exists a profitable side bet involving only a pair of agents. The final step is to show that there is no profitable side bet between any pair of agents if and only if the conditions in Theorem 54 hold.

Lemma 93. *For a given set of reports $\hat{\mathbf{p}}$, wagers $\mathbf{w}$, allocations $\mathbf{y}$ and $\mathbf{n}$ of `yes` and `no` securities, and payments $\boldsymbol{\sigma}$, if there exists a profitable side bet $(\Delta \mathbf{y}, \Delta \mathbf{n}, \Delta \boldsymbol{\sigma})$, then there exists a profitable set bet $(\Delta \mathbf{y}', \Delta \mathbf{n}', \Delta \boldsymbol{\sigma}')$ with $\Delta \sigma_i' = 0$ for all $i \in \mathcal{N}$.*

Proof. For all $i \in \mathcal{N}$, let

$$\Delta y_i' = \Delta y_i - \Delta \sigma_i, \qquad \Delta n_i' = \Delta n_i - \Delta \sigma_i, \qquad \Delta \sigma_i' = 0.$$

234

We show that the three conditions of a profitable side bet are met for $(\Delta \mathbf{y}', \Delta \mathbf{n}', \Delta \boldsymbol{\sigma}')$.

1. First, we have

$$\sum_{i \in \mathcal{N}} \Delta y_i' = \sum_{i \in \mathcal{N}} \Delta y_i - \sum_{i \in \mathcal{N}} \Delta \sigma_i = 0,$$

$$\sum_{i \in \mathcal{N}} \Delta n_i' = \sum_{i \in \mathcal{N}} \Delta n_i - \sum_{i \in \mathcal{N}} \Delta \sigma_i = 0,$$

$$\sum_{i \in \mathcal{N}} \Delta \sigma_i' = 0.$$

2. Next, we have for all $i \in \mathcal{N}$,

$$\min\{y_i + \Delta y_i', n_i + \Delta n_i'\} - (\sigma_i + \Delta \sigma_i')$$

$$= \min\{y_i + \Delta y_i - \Delta \sigma_i, n_i + \Delta n_i - \Delta \sigma_i\} - \sigma_i$$

$$= \min\{y_i + \Delta y_i, n_i + \Delta n_i\} - (\sigma_i + \Delta \sigma_i) \geqslant -w_i.$$

3. Finally, for all $i \in \mathcal{N}$,

$$\hat{p}_i \Delta y_i' + (1 - \hat{p}_i) \Delta n_i' = \hat{p}_i (\Delta y_i - \Delta \sigma_i) + (1 - \hat{p}_i)(\Delta n_i - \Delta \sigma_i)$$

$$= \hat{p}_i \Delta y_i + (1 - \hat{p}_i) \Delta n_i - \Delta \sigma_i$$

$$\geqslant 0 = \Delta \sigma_i'.$$

The inequality must be strict for at least one i since $(\Delta \mathbf{y}, \Delta \mathbf{n}, \Delta \boldsymbol{\sigma})$ is a profitable side bet.

∎

Lemma 94. *For a given set of reports $\hat{\mathbf{p}}$, wagers $\mathbf{w}$, allocations $\mathbf{y}$ and $\mathbf{n}$ of* yes *and* no *securities, and payments $\boldsymbol{\sigma}$, if there exists a profitable side bet $(\Delta \mathbf{y}, \Delta \mathbf{n}, \Delta \boldsymbol{\sigma})$, then there exists a profitable set bet $(\Delta \mathbf{y}', \Delta \mathbf{n}', \Delta \boldsymbol{\sigma}')$ and pair of agents j and k such that $\Delta y_i' = 0$ for all i but j and k, $\Delta n_i' = 0$ for all i but j and k, and $\Delta \sigma_i' = 0$ for all i.*

Proof. From Lemma 93, we can assume without loss of generality that $\Delta\sigma_i = 0$ for all $i \in \mathcal{N}$.

Let $S_y = \{i : \Delta y_i > 0 > \Delta n_i\}$ and $S_n = \{i : \Delta n_i > 0 > \Delta y_i\}$. We first show that these sets are not empty. First note that in order for $(\Delta\mathbf{y}, \Delta\mathbf{n}, \Delta\boldsymbol{\sigma})$ to be a profitable trade, one agent's utility must strictly increase, implying that there must be some i with either $\Delta y_i > 0$ or $\Delta n_i > 0$. If $\Delta n_i > 0$ then there must be some j with $\Delta n_j < 0$ (since $\sum_{i \in \mathcal{N}} \Delta n_i = 0$), which implies that $\Delta y_j > 0$ or agent j would not find the side bet (weakly) profitable. But then there must be some k with $\Delta y_k < 0$ (since $\sum_{i \in \mathcal{N}} \Delta y_i = 0$), and by a similar argument, $\Delta n_k > 0$. The same type of argument can be made starting with $\Delta y_i > 0$.

We next show that there must exist some $j \in S_y$ and $k \in S_n$ such that $\hat{p}_j > \hat{p}_k$. Suppose this were not the case. Then there is some p such that $\hat{p}_j \leqslant p$ for all $j \in S_y$ and $\hat{p}_k \geqslant p$ for all $k \in S_n$. We have already argued that there cannot exist any agent i with $\Delta y_i < 0$ and $\Delta n_i < 0$, which implies that

$$\sum_{i \in S_y \cup S_n} \Delta y_i \leqslant \sum_{i \in \mathcal{N}} \Delta y_i = 0, \qquad \sum_{i \in S_y \cup S_n} \Delta n_i \leqslant \sum_{i \in \mathcal{N}} \Delta n_i = 0, \qquad \text{(C.1)}$$

where both inequalities simultaneously hold with equality only if $\Delta n_i = \Delta y_i = 0$ for all agents $i \notin S_y \cup S_n$. Therefore,

$$\sum_{j \in S_y} (\hat{p}_j \Delta y_j + (1 - \hat{p}_j)\Delta n_j) + \sum_{k \in S_n} (\hat{p}_k \Delta y_k + (1 - \hat{p}_k)\Delta n_k)$$

$$\leqslant p \sum_{j \in S_y} \Delta y_j + (1 - p) \sum_{j \in S_y} \Delta n_j + p \sum_{k \in S_n} \Delta y_k + (1 - p) \sum_{k \in S_n} \Delta n_k \leqslant 0.$$

This shows that the total utility of the set of agents in $S_y \cup S_n$ weakly decreases as a result of the side bet, meaning that either the utility of some agent $i \in S_y \cup S_n$ strictly decreases (in which case the bet cannot be profitable), or the utility of all agents in $S_y \cup S_n$ is unchanged, which means that both inequalities from Equation C.1 are in fact equalites. This in turn implies that $\Delta n_i = \Delta y_i = 0$ for all agents $i \notin S_y \cup S_n$,

so no agent is strictly better off as a result of the side bet (and therefore it is not a profitable side bet).

It remains to show that the agents $j \in S_y$ and $k \in S_n$ with $\hat{p}_j > \hat{p}_k$ can form a profitable trade with each other without violating their budgets. Choose any $p \in [\hat{p}_k, \hat{p}_j]$. For some $\delta > 0$, let

$$\Delta y_j' = \delta \qquad \Delta n_j' = -\frac{p}{1-p}\delta \qquad \Delta \sigma_j' = 0$$

$$\Delta y_k' = -\delta \qquad \Delta n_k' = \frac{p}{1-p}\delta \qquad \Delta \sigma_k' = 0$$

and $\Delta y_i' = \Delta n_i' = \Delta \sigma_i' = 0$ for all other $i \in \mathcal{N}$. We have

$$\hat{p}_j \Delta y_j' + (1 - \hat{p}_j)\Delta n_j' = \hat{p}_j \delta - (1 - \hat{p}_j)\frac{p}{1-p}\delta > 0$$

and

$$\hat{p}_k \Delta y_k' + (1 - \hat{p}_k)\Delta n_k' = -\hat{p}_k \delta + (1 - \hat{p}_k)\frac{p}{1-p}\delta > 0,$$

so the side bet is strictly profitable for both j and k. Finally, since the initial allocation $(\mathbf{y}, \mathbf{n}, \boldsymbol{\sigma})$ and the side bet $(\Delta \mathbf{y}, \Delta \mathbf{n}, \Delta \boldsymbol{\sigma})$ were both feasible, we know that

$$\min\{y_j, y_j + \Delta y_j, n_j, n_j + \Delta n_j\} - \sigma_j \geqslant -w_j$$

$$\min\{y_k, y_k + \Delta y_k, n_k, n_k + \Delta n_k\} - \sigma_k \geqslant -w_k$$

and as long as $\delta \leqslant \min\{-\Delta y_k, -\Delta n_j(1 - p)/p\}$, budgets are not violated. ∎

With these lemmas in place, we are ready to complete the proof.

Proof of Theorem 54. From Lemma 94, we know that a wagering mechanism is Pareto optimal if and only if for all reports $\hat{\mathbf{p}}$ and wagers $\mathbf{w}$, the mechanism's output $(\mathbf{y}(\hat{\mathbf{p}}, \mathbf{w}), \mathbf{n}(\hat{\mathbf{p}}, \mathbf{w}), \boldsymbol{\sigma}(\hat{\mathbf{p}}, \mathbf{w}))$ is such that there exists no profitable side between any pair of agents.

Suppose that for reports $\hat{\mathbf{p}}$ and wagers $\mathbf{w}$, there exists a profitable side bet between agents j and k. Suppose for contradiction that there exists an agent $\ell \in \mathcal{N}$ such that

$$\forall i : \hat{p}_i < \hat{p}_\ell, \qquad \sigma_i(\hat{\mathbf{p}}, \mathbf{w}) = w_i \quad \text{and} \quad y_i(\hat{\mathbf{p}}, \mathbf{w}) = 0, \tag{C.2}$$

$$\forall i : \hat{p}_i > \hat{p}_\ell, \qquad \sigma_i(\hat{\mathbf{p}}, \mathbf{w}) = w_i \quad \text{and} \quad n_i(\hat{\mathbf{p}}, \mathbf{w}) = 0. \tag{C.3}$$

Note that if $\Delta y_j = 0$ then $\Delta n_j \geqslant 0$ (or else j would not find the trade profitable), which implies that $\Delta n_k \leqslant 0$ and $\Delta y_k = 0$ (since $\sum_{i \in \mathcal{N}} n_i = 0$ and $\sum_{i \in \mathcal{N}} y_i = 0$). But then k does not find the trade profitable.

So it must be the case that $\Delta y_j \neq 0$; suppose without loss of generality that $\Delta y_j > 0$. By similar reasoning to above, this implies that $\Delta y_k < 0$, $\Delta n_k > 0$, and $\Delta n_j < 0$. By the definition of a profitable side bet, we know that

$$\hat{p}_j \Delta y_j + (1 - \hat{p}_j)\Delta n_j \geqslant 0 \implies \frac{\Delta y_j}{\Delta n_j} \leqslant \frac{\hat{p}_j - 1}{\hat{p}_j}$$

and

$$\hat{p}_k \Delta y_k + (1 - \hat{p}_k)\Delta n_k \geqslant 0 \implies \frac{\Delta y_k}{\Delta n_k} \geqslant \frac{\hat{p}_k - 1}{\hat{p}_k},$$

with at least one of these inequalities being strict. And, since $\Delta y_k = -\Delta y_j$ and $\Delta n_k = -\Delta n_j$,

$$\frac{\hat{p}_k - 1}{\hat{p}_k} \leqslant \frac{\Delta y_k}{\Delta n_k} \leqslant \frac{\hat{p}_j - 1}{\hat{p}_j},$$

with one of the inequalities being strict. We can now deduce that

$$\frac{\hat{p}_k - 1}{\hat{p}_k} < \frac{\hat{p}_j - 1}{\hat{p}_j} \implies \hat{p}_j > \hat{p}_k.$$

Since $\Delta y_k < 0$, it must be the case that either $y_k(\hat{\mathbf{p}}, \mathbf{w}) > 0$ or $\sigma_k(\hat{\mathbf{p}}, \mathbf{w}) < w_i$ (or both), or else $\min\{y_k + \Delta y_k, n_k + \Delta n_k\} - (\sigma_k + \Delta \sigma_k) < -w_i$, violating the definition of a profitable side bet. By statement C.2, $\hat{p}_k \geqslant p_\ell$. By similar reasoning, $\hat{p}_j \leqslant p_\ell$. But this contradicts $\hat{p}_j > \hat{p}_k$.

For the converse, suppose that for reports $\hat{\mathbf{p}}$ and wagers $\mathbf{w}$, there does not exist an agent $\ell \in \mathcal{N}$ such that

$$\forall i : \hat{p}_i < \hat{p}_\ell, \qquad \sigma_i(\hat{\mathbf{p}}, \mathbf{w}) = w_i \quad \text{and} \quad y_i(\hat{\mathbf{p}}, \mathbf{w}) = 0,$$

$$\forall i : \hat{p}_i > \hat{p}_\ell, \qquad \sigma_i(\hat{\mathbf{p}}, \mathbf{w}) = w_i \quad \text{and} \quad n_i(\hat{\mathbf{p}}, \mathbf{w}) = 0.$$

Let k be the agent with the minimum report such that for all i with $\hat{p}_i > \hat{p}_k$, $\sigma_i(\hat{\mathbf{p}}, \mathbf{w}) = w_i$ and $n_i(\hat{\mathbf{p}}, \mathbf{w}) = 0$. In particular, either $\sigma_k(\hat{\mathbf{p}}, \mathbf{w}) < w_i$ or $n_k(\hat{\mathbf{p}}, \mathbf{w}) > 0$.

Since we know that there does not exist an agent ℓ satisfying the condition above, there must exist a j such that $\hat{p}_j < \hat{p}_k$ with $\sigma_j(\hat{\mathbf{p}}, \mathbf{w}) < w_i$ or $y_j(\hat{\mathbf{p}}, \mathbf{w}) > 0$ (or else k would be such an agent ℓ).

For some $\delta > 0$, let $\Delta n_k = -\delta$, $\Delta y_k = \delta(1 - \hat{p}_k)/\hat{p}_k$, $\Delta n_j = \delta$, and $\Delta y_j = -\delta(1 - \hat{p}_k)/\hat{p}_k$. We conclude the proof by showing that this trade constitutes a profitable side bet for j and k, by examining the three conditions individually.

1. Clearly holds, since $\Delta n_k = -\Delta n_j$ and $\Delta y_k = -\Delta y_j$.

2. If $n_k > 0$ then

$$\min\{y_k + \Delta y_k, n_k + \Delta n_k\} - (\sigma_k + \Delta\sigma_k) \geqslant -\sigma_k \geqslant -w_i$$

for $\delta < n_k$, and if $\sigma_k < w_i$ then

$$\min\{y_k + \Delta y_k, n_k + \Delta n_k\} - (\sigma_k + \Delta\sigma_k) \geqslant -\delta - \sigma_k \geqslant -w_i$$

for $\delta < w_i - \sigma_k$. Similarly, if $y_j > 0$ then

$$\min\{y_j + \Delta y_j, n_j + \Delta n_j\} - (\sigma_j + \Delta\sigma_j) \geqslant -\sigma_j \geqslant -w_i$$

for $\delta(1 - \hat{p}_k)/\hat{p}_k < y_j$, and if $\sigma_j < w_i$ then

$$\min\{y_k + \Delta y_k, n_k + \Delta n_k\} - (\sigma_k + \Delta\sigma_k) \geqslant -\delta(1 - \hat{p}_k)/\hat{p}_k - \sigma_k \geqslant -w_i$$

for $\delta(1 - \hat{p}_k)/\hat{p}_k + \sigma_k < w_i$.

3. For the final condition,

$$\hat{p}_k \Delta y_k + (1 - \hat{p}_k)\Delta n_k = \hat{p}_k \frac{\delta(1 - \hat{p}_k)}{\hat{p}_k} - \delta(1 - \hat{p}_k) = 0 = \Delta\sigma_k,$$

and

$$\hat{p}_j \Delta y_j + (1 - \hat{p}_j)\Delta n_j = -\hat{p}_j \frac{\delta(1 - \hat{p}_k)}{\hat{p}_k} + \delta(1 - \hat{p}_j) = \delta(1 - \frac{\hat{p}_j}{\hat{p}_k}) > 0 = \Delta\sigma_j.$$

■

Table C.1: The key prices in the execution of the adaptive clinching auction.

p	$D_1(p)$	$D_2(p)$	$D_3(p)$	Notes
$0.3 - \varepsilon$	1	2	6	$D_{-i}(p) \geqslant 3$ for all i; no agent can clinch
0.3	0	2	6	$D_{-3}(p) = 2$; agent 3 can clinch one item
$0.35 + \varepsilon$	0	1	5	$D_{-3}(p) = 1$; agent 3 can clinch one item
0.5	0	1	0	$D_{-2}(p) = 0$; agent 2 can clinch one item

C.2 Example of the Adaptive Clinching Auction

In this section we give an illustrating example of the adaptive clinching auction for indivisible goods.

Consider an adaptive clinching auction for three identical, indivisible items and three agents with (value, budget) pairs $(p_1, w_1) = (0.3, 0.3)$, $(p_2, w_2) = (1, 0.7)$, and $(p_3, w_3) = (0.5, 2)$. When the price p is below 0.3, agent 1 demands at least one item, agent 2 demands at least two items, and agent 3 demands at least six items. Therefore, no agent can clinch any items since the demand of the other two agents is at least three.

When the price reaches 0.3, we have that $D_1(0.3) = 0$ (by the fact that $D_i(p) = 0$ when $p = p_i$) and $D_2(0.3) = 2$, and therefore $D_{-3}(0.3) = 2$, so agent 3 clinches one item for a price of 0.3. We update $q_3(p) = 1$, $c_3(p) = 0.3$, $B_3(p) = 1.7$, and $q(p) = 2$.

Next, note that at any price $p > 0.35$, agent 2 demands only one item, and therefore $D_{-3}(p) = 1$. Since this holds for all prices greater than 0.35, agent 3 clinches an item for a price of 0.35 (we refer the reader to Dobzinski et al. [68] for a complete definition of the auction im terms of limits). The running variables are updated to $q_3(p) = 2$, $c_3(p) = 0.65$, $B_3(p) = 1.35$, and $q(p) = 1$.

Finally, when the price reaches 0.5, agent 3 drops out of the auction ($D_3(p) = 0$), as the price equals her value. Therefore, $D_{-2}(p) = 0$, so agent 2 clinches the final item for a price of 0.5.

Table C.1 shows the evolution of demands over the course of the auction.

C.3   Example of the Double Clinching Auction

Consider Example 9: $N = 4$ agents with reported beliefs $\hat{\mathbf{p}} = (0.1, 0.2, 0.5, 0.7)$ and wagers $\mathbf{w} = (1, 1, 1, 1)$. To compute the outcome of the double clinching auction, we first need to determine m^*.

Note that for $m = 2$, $c_y(m) = \inf\{p : \min_{i \in \mathcal{N}} D^y_{-i}(p) < m\} = 0.5$, since at all prices $p < 0.5$ it is the case that $D^y_3 = D^y_4 = \frac{1}{p} > 2$, and for all prices $p \geqslant 0.5$ it is the case that $D^y_{-4} = 0$. Intuitively, once the price reaches 0.5 all agents except agent 4 have dropped out of the `yes` auction since the price equals or exceeds their value for a security. Similarly, $c_n(m) = \inf\{p : \min_{i \in \mathcal{N}} D^n_{-i}(p) < m\} = 0.5$, because at all prices $p \leqslant 0.5$ it is the case that $D^n_1 = D^n_2 = \frac{1}{p} \geqslant 2$, but for prices $0.5 < p < 0.8$ those demands have dropped to $D^n_1 = D^n_2 = \frac{1}{p} < 2$. Intuitively, once the price exceeds 0.5, neither of agents 1 or 2 demands all the items, due to their budget constraint, so each is able to start clinching.

Performing a similar exercise for $m = 2 - \varepsilon$ for any arbitrarily small ε, it can be shown that $c_y(m) = 0.5$ (it is still the case that agent 3 drops out when the price reaches exactly 0.5, allowing agent 4 to clinch all the `yes` securities at this price), and that $c_n(m) > 0.5$ (because the budget constraints for agents 1 and 2 do not bind until the price slightly exceeds 0.5, and neither can begin clinching until this point). In particular, $c_y(m) + c_n(m) > 1$, so $m \in M = \{m : c_y(m) + c_n(m) > 1\}$.

Since $m \in M$ for all $m < 2$, and $2 \notin M$, it is the case that $m^* = \sup M = 2$ (since c_y and c_n are decreasing functions, we know that 2 is an upper bound on M). Now, running an adaptive clinching auction for 2 `yes` securities yields $(\mathbf{y}, \boldsymbol{\sigma}_y) = ((0, 0, 0, 2), (0, 0, 0, 1))$, and running an adaptive clinching auction for 2 `no` securities yields $(\mathbf{n}, \boldsymbol{\sigma}_n) = ((1.39, 0.61, 0, 0), (1, 0.37, 0, 0))$. Thus, the outcome of the double clinching auction for this instance is $(\mathbf{y}, \mathbf{n}, \boldsymbol{\sigma}) = ((0, 0, 0, 2), (1.39, 0.61, 0, 0), (1, 0.37, 0, 1))$.

C.4 Additional Proofs from Section 5.5

C.4.1 Proof of Lemma 61

We prove the lemma for c_y; the argument for c_n is symmetric.

We first show that c_y is weakly decreasing. Consider any m and m' such that $m' > m > 0$. We have $\{p : \min_{i \in \mathcal{N}} D^y_{-i}(p) < m\} \subseteq \{p : \min_{i \in \mathcal{N}} D^y_{-i}(p) < m'\}$, and therefore

$$c_y(m) = \inf\{p : \min_{i \in \mathcal{N}} D^y_{-i}(p) < m\} \geqslant \inf\{p : \min_{i \in \mathcal{N}} D^y_{-i}(p) < m'\} = c_y(m').$$

Finally, note that for all $m > 0$, $D^y_{-N}(\hat{p}_{N-1}) = 0 < m$, and so $\hat{p}_{N-1} \in \{p : \min_{i \in \mathcal{N}} D^y_{-i}(p) < m\}$. Therefore $c_y(m) = \inf\{p : \min_{i \in \mathcal{N}} D^y_{-i}(p) < m\} \leqslant \hat{p}_{N-1} = c_y(0)$.

We now show that c_y is continuous by showing that c_y is surjective; this, together with the fact that c_y is decreasing, implies continuity. To that end, fix any $x \in (0, \hat{p}_{N-1})$. Set $m = \min_{i \in \mathcal{N}} D^y_{-i}(x) > 0$, where the inequality holds because at least two agents have positive demand at price $x < \hat{p}_{N-1}$. Then $c_y(m) = \inf\{p : \min_{i \in \mathcal{N}} D^y_{-i}(p) < m\} = x$. To see this, note that x is a lower bound on $\{p : \min_{i \in \mathcal{N}} D^y_{-i}(p) < m\}$, because $\min_{i \in \mathcal{N}} D^y_{-i}$ is a decreasing function of p; this follows from the definition of D^y_i. To see that x is indeed the greatest lower bound, note that for every $x' > x$, $\min_{i \in \mathcal{N}} D^y_{-i}(x') < m$, since $\min_{i \in \mathcal{N}} D^y_{-i}$ is strictly decreasing in the neighborhood of x; this follows from the fact that $D^y_i(p)$ is strictly decreasing until it reaches 0, and therefore $\min_{i \in \mathcal{N}} D^y_{-i}$ is also strictly decreasing until it reaches 0. Finally, for $x = \hat{p}_{N-1}$ we have that $c_y(0) = x$. Thus c_y is surjective, and therefore continuous. ∎

C.4.2 Proof of Lemma 62

Clearly, m^* is well-defined when $\hat{p}_2 = \hat{p}_{N-1}$. Suppose that $\hat{p}_2 < \hat{p}_{N-1}$. To show m^* is well-defined, it is sufficient to show that the set M is non-empty and bounded above

since this implies the existence of a unique least upper bound.

M must be non-empty since $c_y(0) + c_n(0) = \hat{p}_{N-1} + 1 - \hat{p}_2 > 1$ and therefore $0 \in M$. To show it is bounded, we show that there exists an m with $c_y(m) + c_n(m) \leqslant 1$. Since c_y and c_n are decreasing, this proves the existence of an upper bound.

Note that both $c_y(3\sum_i w_i) \leqslant \frac{2}{5}$ and $c_n(3\sum_i w_i) \leqslant \frac{2}{5}$, since for all agents i,

$$D^{y/n}_{-i}\left(\frac{2}{5}\right) \leqslant D^{y/n}\left(\frac{2}{5}\right) \leqslant \sum_i \frac{w_i}{\frac{2}{5}} = \frac{5}{2}\sum_i w_i < 3\sum_i w_i.$$

Therefore

$$c_y(3\sum_i w_i) + c_n(3\sum_i w_i) \leqslant \frac{2}{5} + \frac{2}{5} = \frac{4}{5} < 1,$$

which shows that $m = 3\sum_i w_i$ is an upper bound for M.

Finally, we show that $m^* > 0$ when $\hat{p}_{N-1} > \hat{p}_2$. If $\hat{p}_{N-1} > \hat{p}_2$ then $c_y(0) + c_n(0) > 1$, as noted earlier, and therefore there must exist some $m' > 0$ such that $c_y(m') + c_n(m') > 1$; that is, $m' \in M$. Therefore, $m^* = \sup M \geqslant m' > 0$. ∎

C.4.3 Proof of Lemma 66

If $m^* = 0$ then $c_y(m^*) + c_n(m^*) = \hat{p}_{N-1} + (1 - \hat{p}_2) = 1$, where the final equality follows from the second part of Lemma 62. So assume that $m^* > 0$.

Suppose $c_y(m^*) + c_n(m^*) > 1$. Let $\varepsilon = \frac{1}{2}(c_y(m^*) + c_n(m^*) - 1)$. Then by continuity there exists δ_y such that for all m' with $|m' - m^*| < \delta_y$, $|c_y(m^*) - c_y(m')| < \varepsilon$. Similarly there is a δ_n such that for all m' with $|m' - m^*| < \delta_n$, $|c_n(m^*) - c_n(m')| < \varepsilon$. Let $\delta = \min\{\delta_y, \delta_n\}$. Then there exists an $m > m^*$ with $m - m^* < \delta$, such that $c_y(m) + c_n(m) > c_y(m^*) + c_n(m^*) - 2\varepsilon = 1$. This violates the definition of m^* as an upper bound on M, a contradiction.

Next suppose $c_y(m^*) + c_n(m^*) < 1$. Let $\varepsilon = \frac{1}{2}(1 - c_y(m^*) - c_n(m^*))$. Then by continuity there exists δ_y such that for all m' with $|m' - m^*| < \delta_y$, $|c_y(m^*) - c_y(m')| < \varepsilon$. Similarly there is a δ_n such that for all m' with $|m' - m^*| < \delta_n$,

$|c_n(m^*) - c_n(m')| < \varepsilon$. Let $\delta = \min\{\delta_y, \delta_n\}$. Then there exists an $m < m^*$ with $m^* - m < \delta$, such that $c_y(m) + c_n(m) < c_y(m^*) + c_n(m^*) + 2\varepsilon = 1$. Thus m is a lower upper bound on M than m^*, violating the definition of m^* as the least upper bound. ∎

C.4.4 Proof of Lemma 69

We prove that $c_y(m)$ is increasing in agent i's report, $\hat{p}_i$. The statement for $c_n(m)$ can be proved analogously with only small modifications. First, if $m = 0$, then $c_y(m)$ is defined to be the second highest report among the agents. Clearly this can only increase as a result of any agent increasing her report.

Suppose that $m > 0$. Consider two reports $\hat{p}_i$ and $\hat{p}'_i$, with $\hat{p}_i < \hat{p}'_i$. Let c be the value of $c_y(m)$ when i reports $\hat{p}_i$, and c' the value when i reports $\hat{p}'_i$. Suppose that $c' < c$. Then there exists some $p \in (c', c)$ such that $D^y_{-j}(p) < m^*$ for some $j \in \mathcal{N}$ when i reports $\hat{p}'_i$, while $D^y_{-j}(p) \geqslant m^*$ for all j when i reports $\hat{p}_i$. But this is not possible since the demand at any given price can only increase as a result of agent i increasing her report. ∎

C.4.5 Proof of Lemma 70

We show the result for the `yes` securities.

If $m = 0$, then $\min_{i\in\mathcal{N}} D^y_{-i}(c_y(m)) = \min_{i\in\mathcal{N}} D^y_{-i}(\hat{p}_{N-1}) = D^y_{-N}(\hat{p}_{N-1}) = 0 \leqslant m$.

Now consider the case where $m > 0$. Suppose for contradiction that $\min_{i\in\mathcal{N}} D_{-i}(c_y(m)) > m$. Let ε be such that for every agent i, if $c_y(m) < \hat{p}_i$ then $c_y(m) + \varepsilon < \hat{p}_i$. For every agent i with $\hat{p}_i \leqslant c_y(m)$, we have that $D^y_i(c_y(m) + \varepsilon) = D^y_i(c_y(m) + \varepsilon) = 0$. For every agent i with $\hat{p}_i > c_y(m)$, we have that $D^y_i(c_y(m)) = \frac{w_i}{c_y(m)}$ and $D^y_i(c_y(m) + \varepsilon) = \frac{w_i}{c_y(m)+\varepsilon}$. Note that by setting ε small enough, we can force $D^y_i(c_y(m) + \varepsilon)$ to be arbitrarily close to $D^y_i(c_y(m))$. Since, by assumption, $\min_{i\in\mathcal{N}} D_{-i}(c_y(m)) > m$, we can set ε small enough that $\min_{i\in\mathcal{N}} D_{-i}(c_y(m) + \varepsilon) > m$. Thus, $c_y(m) + \varepsilon$ is a greater

lower bound on $\{p : \min_{i \in \mathcal{N}} D^y_{-i}(p) < m\}$ than $c_y(m)$, violating the definition of $c_y(m)$.

Bibliography

[1] (2011), "More Google cluster data," Google research blog, Posted at `http://googleresearch.blogspot.com/2011/11/more-google-cluster-data.html`.

[2] (Accessed April 2018), "Augur," `http://www.augur.net/`.

[3] (Accessed April 2018), "Gnosis," `https://gnosis.pm/`.

[4] (Accessed April 2018), "Hivemind," `hhttp://bitcoinhivemind.com/`.

[5] (Accessed April 2018), "Iowa Electronic Markets," `https://iemweb.biz.uiowa.edu/`.

[6] (Accessed April 2018), "Pnyx," `https://pnyx.dss.in.tum.de/`.

[7] (Accessed April 2018), "PredictIt," `https://www.predictit.org/`.

[8] (Accessed April 2018), "RoboVote," `http://www.robovote.org/`.

[9] (Accessed April 2018), "Spliddit," `http://www.spliddit.org/`.

[10] (Accessed Jan 2018), "Capacity Scheduler," `https://hadoop.apache.org/docs/current/hadoop-yarn/hadoop-yarn-site/CapacityScheduler.html`.

[11] (Accessed Jan 2018), "Fair Scheduler," `https://hadoop.apache.org/docs/current/hadoop-yarn/hadoop-yarn-site/FairScheduler.html`.

[12] (Accessed Jan 2018), "SHARCNET," `https://www.sharcnet.ca`.

[13] (Accessed Jan 2018), "Spark Dynamic Allocation," `https://spark.apache.org/docs/latest/job-scheduling.html`.

[14] Abdulkadiroğlu, A. and Bagwell, K. (2013), "Trust, Reciprocity, and Favors in Cooperative Relationships," *American Economic Journal: Microeconomics*, 5, 213–259.

[15] Agrawal, S. and Devanur, N. R. (2015), "Fast algorithms for online stochastic convex programming," in *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1405–1424.

[16] Aleksandrov, M. and Walsh, T. (2017), "Pure Nash Equilibria in Online Fair Division," in *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 42–48.

[17] Aleksandrov, M., Aziz, H., Gaspers, S., and Walsh, T. (2015), "Online Fair Division: Analysing a Food Bank Problem," in *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 2540–2546.

[18] Ali, M. M. (1977), "Probability and utility estimates for racetrack bettors," *Journal of Political Economy*, 85, 803–816.

[19] Arrow, K. J., Forsythe, R., Gorham, M., Hahn, R., Hanson, R., Ledyard, J. O., Levmore, S., Litan, R., Milgrom, P., Nelson, F. D., Neumann, G. R., Ottaviani, M., Schelling, T. C., Shiller, R. J., Smith, V. L., Snowberg, E., Sunstein, C. R., Tetlock, P. C., Tetlock, P. E., Varian, H. R., Wolfers, J., and Zitzewitz, E. (2008), "The promise of prediction markets," *Science*, 320, 877–878.

[20] Athey, S. and Bagwell, K. (2001), "Optimal Collusion with Private Information," *RAND Journal of Economics*, 32, 428–465.

[21] Ausubel, L. M. (2004), "An Efficient Ascending-Bid Auction for Multiple Objects," *The American Economic Review*, 94, 1452–1475.

[22] Azevedo, E. M. and Budish, E. (2017), "Strategy-Proofness in the Large," Tech. rep., National Bureau of Economic Research.

[23] Aziz, H. and Stursberg, P. (2014), "A Generalization of Probabilistic Serial to Randomized Social Choice." in *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 559–565.

[24] Aziz, H., Brandt, F., and Stursberg, P. (2013), "On popular random assignments," in *Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT)*, pp. 183–194.

[25] Aziz, H., Brill, M., Conitzer, V., Elkind, E., Freeman, R., and Walsh, T. (2015), "Justified Representation in Approval-Based Committee Voting," in *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pp. 784–790.

[26] Aziz, H., Brill, M., Conitzer, V., Elkind, E., Freeman, R., and Walsh, T. (2017), "Justified representation in approval-based committee voting," *Social Choice and Welfare*, 48, 461–485.

[27] Bacon, D. F., Parkes, D. C., Chen, Y., Rao, M., Kash, I., and Sridharan, M. (2012), "Predicting Your Own Effort," in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 695–702.

[28] Bansal, N. and Sviridenko, M. (2006), "The Santa Claus problem," in *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 31–40.

[29] Benade, G., Procaccia, A. D., Nath, S., and Shah, N. (2017), "Preference Elicitation for Participatory Budgeting," in *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pp. 376–382.

[30] Berg, J. E. and Rietz, T. A. (2006), "The Iowa Electronic Markets: Stylized facts and open issues," *Information Markets: A New Way of Making Decisions*, pp. 142–169.

[31] Beygelzimer, A., Langford, J., and Pennock, D. M. (2012), "Learning Performance of Prediction Markets with Kelly Bettors," in *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pp. 1317–1318.

[32] Bhattacharya, S., Conitzer, V., Munagala, K., and Xia, L. (2010), "Incentive Compatible Budget Elicitation in Multi-Unit Auctions," in *Proceedings of the Twenty-First ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

[33] Birkhoff, G. (1946), "Tres observaciones sobre el algebra lineal," *Univ. Nac. Tucumán Rev. Ser. A*, 5, 147–151.

[34] Bogomolnaia, A., Moulin, H., and Stong, R. (2005), "Collective Choice Under Dichotomous Preferences," *Journal of Economic Theory*, 122, 165–184.

[35] Boutilier, C., Caragiannis, I., Haber, S., Lu, T., Procaccia, A. D., and Sheffet, O. (2015), "Optimal Social Choice Functions: A Utilitarian View," *Artificial Intelligence*, 227, 190–213.

[36] Boutin, E., Ekanayake, J., Lin, W., Shi, B., Zhou, J., Qian, Z., Wu, M., and Zhou, L. (2014), "Apollo: Scalable and coordinated scheduling for cloud-scale computing," in *Proceedings of the 11th USENIX conference on Operating Systems Design and Implementation (OSDI)*, pp. 285–300, USENIX Association.

[37] Bouveret, S. and Lang, J. (2008), "Efficiency and Envy-Freeness in Fair Division of Indivisible Goods: Logical Representation and Complexity," *Journal of Artificial Intelligence Research*, 32, 525–564.

[38] Brams, S. J., Kilgour, M., and Klamler, C. (2012), "The Undercut Procedure: An Algorithm for the Envy-Free Division of Indivisible Items," *Social Choice and Welfare*, 39, 615–631.

[39] Brandt, F., Conitzer, V., Endress, U., Lang, J., and Procaccia, A. D. (eds.) (2016), *Handbook of Computational Social Choice*, Cambridge University Press.

[40] Brier, G. W. (1950), "Verification of forecasts expressed in terms of probability," *Monthly Weather Review*, 78, 1–3.

[41] Brill, M., Freeman, R., and Conitzer, V. (2016a), "Computing Possible and Necessary Equilibrium Actions (and Bipartisan Set Winners)," in *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 369–375.

[42] Brill, M., Conitzer, V., Freeman, R., and Shah, N. (2016b), "False-Name-Proof Recommendations in Social Networks," in *Proceedings of the 15th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pp. 332–340.

[43] Brill, M., Freeman, R., Janson, S., and Lackner, M. (2017), "Phragmén's Voting Methods and Justified Representation." in *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, vol. 17, pp. 406–413.

[44] Budish, E. (2011), "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes," *Journal of Political Economy*, 119, 1061–1103.

[45] Cabannes, Y. (2004), "Participatory Budgeting: A Significant Contribution to Participatory Democracy," *Environment and Urbanization*, 16, 27–46.

[46] Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., and Kyropoulou, M. (2012a), "The efficiency of fair division," *Theory of Computing Systems*, 50, 589–610.

[47] Caragiannis, I., Covey, J. A., Feldman, M., Homan, C. M., Kaklamanis, C., Karanikolas, N., Procaccia, A. D., and Rosenschein, J. S. (2012b), "On the Approximability of Dodgson and Young Elections," *Artificial Intelligence*, 187–188, 31–51.

[48] Caragiannis, I., Kaklamanis, C., Karanikolas, N., and Procaccia, A. D. (2014), "Socially desirable approximations for Dodgson's voting rule," *ACM Transactions on Algorithms*, 10, article 6.

[49] Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A. D., Shah, N., and Wang, J. (2016), "The Unreasonable Fairness of Maximum Nash Welfare," in *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pp. 305–322.

[50] Chakraborty, M. and Das, S. (2016), "Trading On A Rigged Game: Outcome Manipulation In Prediction Markets," in *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*.

[51] Chamberlin, J. R. and Courant, P. N. (1983), "Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule," *American Political Science Review*, 77, 718–733.

[52] Chen, Y. and Pennock, D. (2007a), "A utility framework for bounded-loss market makers," in *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence (UAI)*.

[53] Chen, Y. and Pennock, D. (2010), "Designing Markets for Prediction," *AI Magazine*, 31, 42–52.

[54] Chen, Y. and Pennock, D. M. (2007b), "A utility framework for bounded-loss market makers," in *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence (UAI)*.

[55] Chen, Y. and Vaughan, J. W. (2010), "A New Understanding of Prediction Markets via No-regret Learning," in *Proceedings of the 11th ACM Conference on Electronic Commerce*, pp. 189–198, ACM.

[56] Chen, Y., Devanur, N. R., Pennock, D. M., and Vaughan, J. W. (2014), "Removing Arbitrage from Wagering Mechanisms," in *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*.

[57] Clark, J., Bonneau, J., Felten, E. W., Kroll, J. A., Miller, A., and Narayanan, A. (2014), "On Decentralizing Prediction Markets and Order Books," in *Workshop on the Economics of Information Security, State College, Pennsylvania*.

[58] Cole, R. and Gkatzelis, V. (2015), "Approximating the Nash Social Welfare with Indivisible Items," in *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 371–380.

[59] Conitzer, V., Brill, M., and Freeman, R. (2015), "Crowdsourcing Societal Tradeoffs," in *Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pp. 1213–1217.

[60] Conitzer, V., Freeman, R., Brill, M., and Li, Y. (2016), "Rules for Choosing Societal Tradeoffs." in *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 460–467.

[61] Conitzer, V., Freeman, R., and Shah, N. (2017), "Fair Public Decision Making," in *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pp. 629–646.

[62] Cover, T. M. and Thomas, J. A. (2006), *Elements of Information Theory, 2nd Edition*, Wiley, New York.

[63] Cummings, R., Pennock, D. M., and Vaughan, J. W. (2016), "The Possibilities and Limitations of Private Prediction Markets," in *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*.

[64] Darmann, A. and Schauer, J. (2015), "Maximizing Nash Product Social Welfare in Allocating Indivisible Goods," *European Journal of Operational Research*, 247, 548–559.

[65] Dasgupta, A. and Ghosh, A. (2013), "Crowdsourced Judgement Elicitation with Endogenous Proficiency," in *Proceedings of the 22nd International World Wide Web Conference (WWW)*, pp. 319–330.

[66] Dearden, J. A. (1998), "Serial cost sharing of excludable public goods: general cost functions," *Economic Theory*, 12, 189–198.

[67] Demers, A., Keshav, S., and Shenker, S. (1989), "Analysis and Simulation of a Fair Queueing Algorithm," in *ACM SIGCOMM Computer Communication Review*, vol. 19, pp. 1–12.

[68] Dobzinski, S., Lavi, R., and Nisan, N. (2012), "Multi-Unit Auctions with Budget Limits," *Games and Economic Behavior*, 74, 486–503.

[69] Dubey, P., Geanakoplos, J., and Shubik, M. (1987), "The Revelation of Information in Strategic Market Games: A Critique of Rational Expectations Equilibrium," *Journal of Mathematical Economics*, 16, 3105–3137.

[70] Duke University (2017), "The Duke Computer Cluster," `http://rc.duke.edu/the-duke-compute-cluster.edu`.

[71] Ehlers, L., Peters, H., and Storcken, T. (2004), "Threshold Strategy-Proofness: On Manipulability in Large Voting Problems," *Games and Economic Behavior*, 49, 103–116.

[72] Eisenberg, E. and Gale, D. (1959), "Consensus of Subjective Probabilities: The Pari-Mutuel Method," *The Annals of Mathematical Statistics*, 30, 165–168.

[73] Elkind, E., Faliszewski, P., Skowron, P., and Slinko, A. (2014), "Properties of Multiwinner Voting Rules," in *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pp. 53–60.

[74] Fain, B., Goel, A., and Munagala, K. (2016), "The Core of the Participatory Budgeting Problem," in *Proceedings of the 12th Conference on Web and Internet Economics (WINE)*, pp. 384–399.

[75] Fain, B., Munagala, K., and Shah, N. (2018), "Fair Allocation of Indivisible Public Goods," *arXiv preprint arXiv:1805.03164*.

[76] Foley, D. (1967), "Resource allocation and the public sector," *Yale Economics Essays*, 7, 45–98.

[77] Freeman, R. and Pennock, D. M. (2018), "An Axiomatic View of the Parimutuel Consensus Mechanism," in *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, Forthcoming.

[78] Freeman, R., Brill, M., and Conitzer, V. (2014), "On the Axiomatic Characterization of Runoff Voting Rules," in *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 675–681.

[79] Freeman, R., Brill, M., and Conitzer, V. (2015), "General Tiebreaking Schemes for Computational Social Choice," in *Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.

[80] Freeman, R., Haney, S., and Panigrahi, D. (2016), "On the Price of Stability of Undirected Multicast Games," in *Proceedings of the 12th Conference on Web and Internet Economics (WINE)*, pp. 354–368.

[81] Freeman, R., Lahaie, S., and Pennock, D. M. (2017a), "Crowdsourced Outcome Determination in Prediction Markets." in *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pp. 523–529.

[82] Freeman, R., Pennock, D. M., and Vaughan, J. W. (2017b), "The Double Clinching Auction for Wagering," in *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pp. 43–60.

[83] Freeman, R., Zahedi, S., and Conitzer, V. (2017c), "Fair and Efficient Social Choice in Dynamic Settings," in *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 4580–4587.

[84] Freeman, R., Zahedi, S. M., Conitzer, V., and Lee, B. C. (2018), "Dynamic Proportional Sharing: A Game-Theoretic Approach," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 2, 3:1–3:36.

[85] Friedman, E., Psomas, C.-A., and Vardi, S. (2017), "Controlled Dynamic Fair Division," in *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pp. 461–478.

[86] Frongillo, R. and Witkowski, J. (2016), "A Geometric Method to Construct Minimal Peer Prediction Mechanisms," in *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*.

[87] Galebach, B. (2004), "ProbabilitySports," http://probabilitysports.com/.

[88] Genest, C. and Zidek, J. V. (1986), "Combining Probability Distributions: A Critique and an Annotated Bibliography," *Statistical Science*, 1, 114–148.

[89] Ghodsi, A., Zaharia, M., Hindman, B., Konwinski, A., Shenker, S., and Stoica, I. (2011), "Dominant Resource Fairness: Fair Allocation of Multiple Resource Types," in *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation (NSDI)*, pp. 24–37.

[90] Ghodsi, A., Sekar, V., Zaharia, M., and Stoica, I. (2012), "Multi-Resource Fair Queueing for Packet Processing," *ACM SIGCOMM Computer Communication Review*, 42, 1–12.

[91] Ghodsi, A., Zaharia, M., Shenker, S., and Stoica, I. (2013), "Choosy: Max-min Fair Sharing for Datacenter Jobs with Constraints," in *Proceedings of the 8th ACM European Conference on Computer Systems*, pp. 365–378, ACM.

[92] Gneiting, T. and Raftery, A. E. (2007), "Strictly Proper Scoring Rules, Prediction, and Estimation," *Journal of the American Statistical Association*, 102, 359–378.

[93] Goel, A., Krishnaswamy, A. K., Sakshuwong, S., and Aitamurto, T. (2016), "Knapsack Voting: Voting mechanisms for Participatory Budgeting," Manuscript.

[94] Goel, G., Mirrokni, V., and Leme, R. P. (2015), "Clinching Auctions with Online Supply," *Games and Economic Behavior*.

[95] Goldman, J. and Procaccia, A. D. (2015), "Spliddit: Unleashing Fair Division Algorithms," *ACM SIGecom Exchanges*, 13, 41–46.

[96] Gorokh, A., Banerjee, S., and Iyer, K. (2016), "Near-Efficient Allocation Using Artificial Currency in Repeated Settings," .

[97] Gorokh, A., Banerjee, S., and Iyer, K. (2017), "From Monetary to Non-Monetary Mechanism Design Via Artificial Currencies," .

[98] Gulati, A., Shanmuganathan, G., Zhang, X., and Varman, P. J. (2012), "Demand Based Hierarchical QoS Using Storage Resource Pools." in *USENIX Annual Technical Conference*, pp. 1–13.

[99] Guo, M., Conitzer, V., and Reeves, D. M. (2009), "Competitive Repeated Allocation Without Payments," in *Proceedings of the 5th Workshop on Internet and Network Economics (WINE)*, pp. 244–255.

[100] Haake, C.-J., Raith, M. G., and Su, F. E. (2002), "Bidding for envy-freeness: A procedural approach to n-player fair-division problems," *Social Choice and Welfare*, 19, 723–749.

[101] Hanson, R. (2003), "Combinatorial Information Market Design," *Information Systems Frontiers*, 5, 107–119.

[102] Jacobs, R. A. (1995), "Methods for Combining Experts' Probability Assessments," *Neural Computation*, 7, 867–888.

[103] Jalali, A., Padovani, R., and Pankaj, R. (2000), "Data Throughput of CDMA-HDR a High Efficiency-High Data Rate Personal Communication Wireless System," in *Proceedings of the IEEE Vehicular Technology Conference*, pp. 1854–1858.

[104] Johnstone, D. J. (2007), "The Parimutuel Kelly Probability Scoring Rule," *Decision Analysis*, 4, 66–75.

[105] Jordan, J. and Radner, R. (1979), "The nonexistence of Rational expectations equilibrium: A robust example," Working paper.

[106] Jurca, R. and Faltings, B. (2008), "Incentives for Expressing Opinions in Online Polls," in *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*, pp. 119–128.

[107] Jurca, R. and Faltings, B. (2011), "Incentives for Answering Hypothetical Questions," in *Workshop on Social Computing and User Generated Content, EC-11*.

[108] Kash, I., Procaccia, A. D., and Shah, N. (2014), "No Agent Left Behind: Dynamic Fair Division of Multiple Resources," *Journal of Artificial Intelligence Research*, pp. 579–603.

[109] Kilgour, D. M. and Gerchak, Y. (2004), "Elicitation of Probabilities Using Competitive Scoring Rules," *Decision Analysis*, 1, 108–113.

[110] Klijn, F. (2000), "An algorithm for envy-free allocations in an economy with indivisible objects and money," *Social Choice and Welfare*, 17, 201–215.

[111] Kurokawa, D., Procaccia, A. D., and Shah, N. (2015), "Leximin Allocations in the Real World," in *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*, pp. 345–362.

[112] Lambert, N. S., Langford, J., Wortman, J., Chen, Y., Reeves, D., Shoham, Y., and Pennock, D. M. (2008), "Self-Financed Wagering Mechanisms for Forecasting," in *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*.

[113] Lambert, N. S., Langford, J., Vaughan, J. W., Chen, Y., Reeves, D., Shoham, Y., and Pennock, D. M. (2015), "An axiomatic characterization of wagering mechanisms," *Journal of Economic Theory*, 156, 389–416.

[114] Lang, J. and Xia, L. (2016), "Voting in Combinatorial Domains," in *Handbook of Computational Social Choice*, eds. F. Brandt, V. Conitzer, U. Endress, J. Lang, and A. D. Procaccia, chap. 8, Cambridge University Press.

[115] Lawrence Berkeley National Laboratory (2017), "National Energy Research Scientific Computing Center," https://www.nersc.gov.

[116] Lee, E. (2017), "APX-hardness of maximizing Nash social welfare with indivisible items," *Information Processing Letters*, 122, 17–20.

[117] Lipton, R. J., Markakis, E., Mossel, E., and Saberi, A. (2004), "On Approximately Fair Allocations of Indivisible Goods," in *Proceedings of the 6th ACM Conference on Electronic Commerce (EC)*, pp. 125–131.

[118] Manski, C. F. (2006), "Interpreting the Predictions of Prediction Markets," *Economics Letters*, 91, 425–429.

[119] Miller, N., Resnick, P., and Zeckhauser, R. (2005), "Eliciting Informative Feedback: The Peer-Prediction Method," *Management Science*, 51, 1359–1373.

[120] Monroe, B. L. (1995), "Fully Proportional Representation," *American Political Science Review*, 89, 925–940.

[121] Morris, P. A. (1977), "Combining Expert Judgments: A Bayesian Approach," *Management Science*, 23, 679–693.

[122] Moulin, H. (1994), "Serial Cost-Sharing of Excludable Public Goods," *Review of Economic Studies*, 61, 305–325.

[123] Moulin, H. and Shenker, S. (2001), "Strategyproof sharing of submodular costs: budget balance versus efficiency," *Economic Theory*, 18, 511–533.

[124] Nash, J. (1950), "The Bargaining Problem," *Econometrica*, 18, 155–162.

[125] Nemhauser, G. L. and Wolsey, L. A. (1988), "Integer and Combinatorial Optimization. Interscience Series in Discrete Mathematics and Optimization," *ed: John Wiley & Sons*.

[126] Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V. (2007), *Algorithmic Game Theory*, Cambridge University Press.

[127] Othman, A., Pennock, D. M., Reeves, D. M., and Sandholm, T. (2013), "A Practical Liquidity-Sensitive Automated Market Maker," *ACM Transactions on Economics and Computation*, 1, 14.

[128] Parekh, A. K. and Gallager, R. G. (1993), "A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Single-Node Case," *IEEE/ACM Transactions on Networking*, 1, 344–357.

[129] Parkes, D. C. and Procaccia, A. D. (2013), "Dynamic Social Choice with Evolving Preferences," in *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 767–773.

[130] Parkes, D. C., Cavallo, R., Constantin, F., and Singh, S. (2010), "Dynamic Incentive Mechanisms," *AI Magazine*, 31, 79–94.

[131] Pennock, D. M., Syrgkanis, V., and Vaughan, J. W. (2016), "Bounded Rationality in Wagering Mechanisms," in *Proceedings of the 32nd Annual Conference on Uncertainty in Artificial Intelligence (UAI)*.

[132] Peterson, J. and Krug, J. (2015), "Augur: A Decentralized, Open-Source Platform for Prediction Markets," *arXiv preprint arXiv:1501.01042*.

[133] Plott, C. R., Wit, J., and Yang, W. C. (2003), "Parimutuel Betting Markets as Information Aggregation Devices: Experimental Results," *Economic Theory*, 22, 311–351.

[134] Posner, E. A. and Weyl, E. G. (2014), "Quadratic Voting as Efficient Corporate Governance," *The University of Chicago Law Review*, 81, 251–272.

[135] Posner, E. A. and Weyl, E. G. (2015), "Voting Squared: Quadratic Voting in Democratic Politics," *Vanderbilt Law Review*, 68, 441.

[136] Prelec, D. (2004), "A Bayesian Truth Serum for Subjective Data," *Science*, 306, 462–466.

[137] Procaccia, A. D. and Rosenschein, J. S. (2006), "The Distortion of Cardinal Preferences in Voting," in *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*, pp. 317–331.

[138] Procaccia, A. D. and Wang, J. (2014), "Fair Enough: Guaranteeing Approximate Maximin Shares," in *Proceedings of the 14th ACM Conference on Economics and Computation (EC)*, pp. 675–692.

[139] Pu, Q., Li, H., Zaharia, M., Ghodsi, A., and Stoica, I. (2016), "FairRide: Near-Optimal, Fair Cache Sharing." in *Proceedings of the 13th USENIX Conference on Networked Systems Design and Implementation (NSDI)*, pp. 393–406.

[140] Ramezani, S. and Endriss, U. (2010), "Nash Social Welfare in Multiagent Resource Allocation," in *Proceedings of the 12th International Workshop on Agent-Mediated Electronic Commerce (AMEC)*, pp. 117–131.

[141] Reijnierse, J. and Potters, J. A. M. (1998), "On finding an envy-free Pareto-optimal division," *Mathematical Programming*, 83, 291–311.

[142] Reiss, C., Wilkes, J., and Hellerstein, J. L. (2011), "Google cluster-usage traces: format + schema," Technical report, Google Inc., Mountain View, CA, USA, Revised 2014-11-17 for version 2.1. Posted at `https://github.com/google/cluster-data`.

[143] Roberts, D. J. and Postlewaite, A. (1976), "The Incentives for Price-Taking Behavior in Large Exchange Economies," *Econometrica: Journal of the Econometric Society*, pp. 115–127.

[144] Roth, A. E., Sönmez, T., and Ünver, M. U. (2004), "Kidney Exchange," *The Quarterly Journal of Economics*, 119, 457–488.

[145] Roth, A. E., Sönmez, T., and Ünver, M. U. (2005), "Pairwise kidney exchange," *Journal of Economic Theory*, 125, 151–188.

[146] Sandholm, T. and Lai, K. (2010), "Dynamic Proportional Share Scheduling in Hadoop," in *Job scheduling strategies for parallel processing*, pp. 110–131.

[147] Savage, L. J. (1971), "Elicitation of Personal Probabilities and Expectations," *Journal of the American Statistical Association*, 66, 783–801.

[148] Shi, P., Conitzer, V., and Guo, M. (2009), "Prediction Mechanisms That Do Not Incentivize Undesirable Actions," in *Proceedings of the 5th International Workshop on Internet and Network Economics (WINE)*, pp. 89–100.

[149] Shieh, A., Kandula, S., Greenberg, A. G., Kim, C., and Saha, B. (2011), "Sharing the Data Center Network." in *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation (NSDI)*, vol. 11, pp. 23–23.

[150] Shnayder, V., Agarwal, A., Frongillo, R., and Parkes, D. C. (2016), "Informed Truthfulness in Multi-Task Peer Prediction," in *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pp. 179–196.

[151] Snyder, W. W. (1978), "Horse Racing: Testing the Efficient Markets Model," *Journal of Finance*, 33, 1109–1118.

[152] Spann, M. and Skiera, B. (2003), "Internet-Based Virtual Stock Markets for Business Forecasting," *Management Science*, 49, 1310–1326.

[153] Steinhaus, H. (1948), "The problem of fair division," *Econometrica*, 16, 101–104.

[154] Tang, S., Lee, B.-S., He, B., and Liu, H. (2014), "Long-Term Resource Fairness: Towards Economic Fairness on Pay-as-you-use Computing Systems," in *Proceedings of the 28th ACM International Conference on Supercomputing*, pp. 251–260.

[155] Thaler, R. H. and Ziemba, W. T. (1988), "Anomalies: Parimutuel Betting Markets: Racetracks and lotteries," *Journal of Economic Perspectives*, 2, 161–174.

[156] Varian, H. (1974), "Equity, envy and efficiency," *Journal of Economic Theory*, 9, 63–91.

[157] Verma, A., Pedrosa, L., Korupolu, M., Oppenheimer, D., Tune, E., and Wilkes, J. (2015), "Large-scale Cluster Management at Google with Borg," in *Proceedings of the 10th European Conference on Computer Systems (EuroSys)*, pp. 18:1–18:17, ACM.

[158] Viswanath, P., Tse, D. N., and Laroia, R. (2002), "Opportunistic beamforming using dumb antennas," *IEEE Transactions on Information Theory*, 48, 1277–1294.

[159] Walsh, T. (2011), "Online Cake Cutting," in *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*, pp. 292–305.

[160] Wang, X. and Martínez, J. F. (2015), "XChange: Scalable Dynamic Multi-Resource Allocation in Multicore Architectures," in *Proceedings of the 21st IEEE International Symposium on High Performance Computer Architecture (HPCA)*.

[161] Weyl, E. G. (2017), "The robustness of quadratic voting," *Public choice*, 172, 75–107.

[162] Witkowski, J. (2014), "Robust Peer Prediction Mechanisms," Ph.D. thesis, Department of Computer Science, Albert-Ludwigs-Universität Freiburg.

[163] Witkowski, J. and Parkes, D. C. (2012), "A Robust Bayesian Truth Serum for Small Populations." in *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*.

[164] Witkowski, J., Atanasov, P., Ungar, L. H., and Krause, A. (2017), "Proper Proxy Scoring Rules," in *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*.

[165] Witkowski, J., Freeman, R., Vaughan, J. W., Pennock, D., and Krause, A. (2018), "Incentive-Compatible Forecasting Competitions," in *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*.

[166] Wolfers, J. and Zitzewitz, E. (2006), "Interpreting Prediction Market Prices as Probabilities," Tech. Rep. Working Paper 12200, National Bureau of Economic Research.

[167] Xu, H., Freeman, R., Conitzer, V., Dughmi, S., and Tambe, M. (2016), "Signaling in Bayesian Stackelberg Games," in *Proceedings of the 15th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pp. 150–158.

[168] Zaharia, M., Chowdhury, M., Franklin, M. J., Shenker, S., and Stoica, I. (2010), "Spark: Cluster Computing with Working Sets," in *Proceedings of the 2nd USENIX Conference on Hot Topics in Cloud Computing*, vol. 10, p. 10.

[169] Zahedi, S. M. and Lee, B. C. (2014), "REF: Resource Elasticity Fairness with Sharing Incentives for Multiprocessors," in *Proceedings of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pp. 145–160, ACM.

Biography

Rupert Freeman was born on October 16, 1989, in Auckland, New Zealand. He is a Ph.D. candidate in computer science at Duke University, advised by Vincent Conitzer. He holds a B.Sc.(Hons) degree from the University of Auckland. He is the recipient of a Facebook Ph.D. Fellowship, and Duke Computer Science outstanding Teaching Assistant and Outstanding Preliminary Exam awards. In July 2018, Rupert will commence a postdoc at Microsoft Research New York City.