

An Equivalence Between Wagering and Fair-Division Mechanisms*

Rupert Freeman, David M. Pennock, Jennifer Wortman Vaughan

Abstract

We draw a surprising and direct mathematical equivalence between the class of allocation mechanisms for divisible goods studied in the context of fair division and the class of weakly budget-balanced wagering mechanisms designed for eliciting probabilities. The equivalence rests on the intuition that wagering is an allocation of financial securities among bettors, with a bettor’s value for each security proportional to her belief about the likelihood of a future event. The equivalence leads to theoretical advances and new practical approaches for both fair division and wagering. Known wagering mechanisms based on proper scoring rules yield fair allocation mechanisms with desirable properties, including the first strictly incentive compatible fair division mechanism. At the same time, allocation mechanisms make for novel wagering rules, including one that requires only ordinal uncertainty judgments and one that outperforms existing rules in a range of simulations.

1 Introduction

Consider the following two scenarios. In the first, an information-seeking principal would like to elicit credible probabilistic forecasts from a group of agents. She does so by collecting wagers from the agents along with their predictions, and redistributing these wagers in such a way that agents with more accurate predictions are more highly rewarded, potentially keeping a cut for herself. In the second, a neutral mediator needs to allocate a set of divisible goods or resources, such as plots of land or compute cycles, to a group of agents who all have different preferences over the goods and potentially different entitlements. His top priority is ensuring that each agent walks away with her fair share.

On the surface, these scenarios appear quite distinct. Central to wagering is the idea of money changing hands, with each agent’s rewards contingent on an uncertain future state of the world. In contrast, the fair division of goods or resources typically involves no exchange of money and no inherent uncertainty.

Despite these apparent differences, we show that these scenarios are two sides of the same coin. In particular, we show that there is a one-to-one correspondence between the class of weakly budget-balanced wagering mechanisms—those for which the principal is guaranteed no loss—and the class of allocation mechanisms for divisible goods under additive valuations. Furthermore, we show that commonly studied properties of wagering mechanisms correspond precisely to commonly studied properties of fair division. *Individual rationality*, which ensures that risk-neutral agents benefit in expectation from participation in a wagering mechanism, corresponds to the notion of *proportionality* in fair division, which, loosely speaking, guarantees each agent at least a minimal share of allocated goods. And properties common to the literature in both settings, including *incentive compatibility* and *Pareto optimality*, carry over immediately, with minor caveats.

This correspondence has implications for both fair division and wagering. For example, weighted score wagering mechanisms (Lambert et al., 2008, 2015), which build on the machinery of proper scoring rules (Gneiting and Raftery, 2007), can be imported into the fair division setting to yield allocation mechanisms that are proportional, envy-free, and strictly incentive compatible—to the best of our knowledge, the first strictly incentive compatible allocation mechanisms. On the flip side, the strong demand matching allocation mechanism (Cole et al., 2013) can be imported into the wagering setting to produce a mechanism that, in our simulations, outperforms standard wagering mechanisms in terms of amount of trade facilitated and agent utility. As another example, since the constrained serial dictatorship (Aziz and Ye, 2014) allocation mechanism requires only ordinal agent preferences, the correspondence yields the first wagering mechanism that allows agents to report ordinal beliefs.

*A preliminary version of this paper appears in AAAI’19.

The correspondence we show builds on an earlier observation that the output of a wagering mechanism can be interpreted as a sale of Arrow-Debreu securities to bettors (Freeman et al., 2017). Once all wagers have been collected, a wagering mechanism can be viewed as allocating a fixed number of securities of different types among the bettors (possibly with some leftover for the principal), with a bettor’s value for each security proportional to her belief about the likelihood of a future event. Under this interpretation, it is natural to consider the use of an allocation mechanism to allocate the securities. Perhaps more surprisingly, our correspondence shows that the class of weakly budget-balanced wagering mechanisms is general enough to capture the entire space of allocation mechanisms for divisible goods under additive valuations, so any wagering mechanism can be used for allocation of general goods with no exchange of money.

2 Preliminaries and Background

To set the stage for our main results, we begin with a review of the formal models of both fair division and wagering, and define the properties most commonly sought out in allocation and wagering mechanisms.

Throughout the paper, we use the notation $[n]$ to denote the set $\{1, \dots, n\}$ and the notation \mathbf{x}_{-i} to denote all components of the vector \mathbf{x} except the i th.

2.1 Fair Division

The study of fair division dates back to the 1940s (Steinhaus, 1948) and many variants of the problem have been examined. In this paper, we focus on the problem of allocating divisible goods to agents with additive values, a special case of the cake-cutting problem (Robertson and Webb, 1998; Procaccia, 2013).

Formally, consider a set of n agents indexed $\{1, \dots, n\}$ and a set of m divisible goods indexed $\{1, \dots, m\}$. Each agent i has a private set of values $v_{i,j} \geq 0$ for each good j . As is standard in the cake-cutting literature, we assume that values are normalized so that $\sum_{j=1}^m v_{i,j} = 1$. Let \mathbf{z} denote a bundle of goods, with each component $z_j \in [0, 1]$ specifying a fractional piece of good j . Agent i ’s value for \mathbf{z} is then $v_i(\mathbf{z}) = \sum_{j=1}^m v_{i,j} z_j$. Each agent also has a fixed weight $e_i > 0$, capturing the agent’s entitlement or priority. Let $E = \sum_{i=1}^n e_i$.

An *allocation* consists of a bundle of goods \mathbf{a}_i assigned to each agent i . It is permitted to allocate goods incompletely, resulting in waste. A *feasible allocation* is therefore any allocation for which $\sum_{i=1}^n a_{i,j} \leq 1$ for all j , and $a_{i,j} \geq 0$ for all i and j . An *allocation mechanism* \mathcal{A} takes as input a vector of reported values $\hat{\mathbf{v}}$ and weights \mathbf{e} for each agent i and outputs a feasible allocation. We use $\mathcal{A}_{i,j}(\hat{\mathbf{v}}, \mathbf{e})$ to denote the fraction of good j allocated to agent i by the allocation mechanism \mathcal{A} with input $\hat{\mathbf{v}}$ and \mathbf{e} , and $\mathcal{A}_i(\hat{\mathbf{v}}, \mathbf{e})$ for i ’s full allocation.

We note that some of the allocation mechanisms we discuss are either ill-defined or lose some of their fairness properties when weights are unequal. As long as weights are rational, these mechanisms can be extended to the weighted setting by creating multiple copies of each agent i such that the number of copies of i is proportional to e_i without sacrificing any of the properties described below (Brams and Taylor 1996, page 152). However, this method can result in an exponential increase in time and space requirements, making some unweighted mechanisms infeasible to run in the weighted setting.

Several notions of fairness have been proposed in the fair division literature. Proportionality (Steinhaus, 1948) requires that each agent get an e_i/E share of her total (reported) value for the set of all goods.

Definition 1. An allocation mechanism \mathcal{A} is *proportional* if, for all $\hat{\mathbf{v}}$ and \mathbf{e} , and for all $i \in [n]$, $\hat{v}_i(\mathcal{A}_i(\hat{\mathbf{v}}, \mathbf{e})) \geq e_i/E$.

Envy-freeness (Foley, 1967) requires that no agent can ever prefer another agent’s allocation to her own, after accounting for differences in weights.

Definition 2. An allocation mechanism \mathcal{A} is *envy-free* if, for all $\hat{\mathbf{v}}$ and \mathbf{e} , and for all $i, i' \in [n]$, $(1/e_i)\hat{v}_i(\mathcal{A}_i(\hat{\mathbf{v}}, \mathbf{e})) \geq (1/e_{i'})\hat{v}_{i'}(\mathcal{A}_{i'}(\hat{\mathbf{v}}, \mathbf{e}))$.

Note that a trivial envy-free and proportional allocation always exists; simply allocate each agent i an e_i/E fraction of each good.

Pareto optimality is a basic notion of economic efficiency. It says that a mechanism should not output allocations for which it is possible to make some agent better off without making any other agent worse off.

Definition 3. An allocation mechanism \mathcal{A} is *Pareto optimal* if for all $\hat{\mathbf{v}}$ and \mathbf{e} , and all alternative allocations $\{\mathbf{a}'_1, \dots, \mathbf{a}'_n\}$, if $\hat{v}_i(\mathbf{a}'_i) > \hat{v}_i(\mathcal{A}_i(\hat{\mathbf{v}}, \mathbf{e}))$ for some i , there exists a j with $\hat{v}_j(\mathbf{a}'_j) < \hat{v}_j(\mathcal{A}_j(\hat{\mathbf{v}}, \mathbf{e}))$.

Finally, we might desire mechanisms that do not incentivize agents to misreport their (private) values.

Definition 4. An allocation mechanism \mathcal{A} is *weakly incentive compatible* if, for every agent i with values \mathbf{v}_i , and all reports $\hat{\mathbf{v}}$ and weights \mathbf{e} , $v_i(\mathcal{A}_i((\mathbf{v}_i, \hat{\mathbf{v}}_{-i}), \mathbf{e})) \geq v_i(\mathcal{A}_i(\hat{\mathbf{v}}, \mathbf{e}))$. \mathcal{A} is *strictly incentive compatible* if this inequality is strict whenever $\mathbf{v}_i \neq \hat{\mathbf{v}}_i$.

In fair division, fairness is typically emphasized over incentive compatibility, and weak incentive compatibility (or equivalently, dominant strategy truthfulness) is often considered sufficient (e.g., Cole et al. 2013).

2.2 Wagering

While the study of wagering mechanisms also has a long history (Eisenberg and Gale, 1959), the model we study here was formalized by Lambert et al. (2008), building on ideas from Kilgour and Gerchak (2004).

Consider a random variable X that takes a value in $\{1, \dots, m\}$. There is a set of n agents (or bettors) indexed $\{1, \dots, n\}$. Each agent i has private, subjective, immutable beliefs \mathbf{p}_i , where $p_{i,j}$ is the probability i assigns to the event $X = j$, and a budget or *wager* $w_i > 0$, the maximum amount of money i is prepared to lose.

A *wagering mechanism* Π is used to elicit beliefs from the agents. Each agent reports beliefs $\hat{\mathbf{p}}_i$ along with her wager w_i . The wagering mechanism defines a net payoff $\Pi_i(\hat{\mathbf{p}}, \mathbf{w}, x)$ to each agent that depends on the full set of agent reports $\hat{\mathbf{p}}$, the vector \mathbf{w} of agent wagers, and x , the observed value of X . To be a valid wagering mechanism, it must be the case that no agent loses more than her wager (i.e., $\Pi_i(\hat{\mathbf{p}}, \mathbf{w}, x) \geq -w_i$ for all $\hat{\mathbf{p}}, \mathbf{w}$, and x). Let $W = \sum_{i=1}^n w_i$.

In prior work, we observed that the output of a wagering mechanism can be interpreted as an allocation of securities with payoffs that depend on the realization of X (Freeman et al., 2017). For any $j \in [m]$, we define a *type- j security* to be a contract that pays off \$1 in the event that $X = j$ and \$0 otherwise. If i is risk neutral, she should be willing to buy a type- j security at any price up to $p_{i,j}$.

To define the output of a wagering mechanism in terms of securities, we can think of each agent i paying a price of w_i up front in exchange for $b_{i,j}$ securities of each type j , where $b_{i,j} = w_i + \Pi_i(\hat{\mathbf{p}}, \mathbf{w}, j) \geq 0$. In the event that $X = j$, agent i receives net payoff $b_{i,j} - w_i = \Pi_i(\hat{\mathbf{p}}, \mathbf{w}, j)$. Therefore, the output of a wagering mechanism can be completely specified by an allocation of securities to agents. We take this interpretation of wagering mechanisms throughout the remainder of the paper. To emphasize the distinction, we talk about wagering mechanisms \mathcal{B} , letting $\mathcal{B}_{i,j}(\hat{\mathbf{p}}, \mathbf{w})$ denote the number of type- j securities allocated to agent i by the mechanism \mathcal{B} on input $\hat{\mathbf{p}}$ and \mathbf{w} .

We note that this interpretation differs slightly from the one used in Freeman et al. (2017), in which complete sets of securities were transformed into cash and bettors' payments were explicitly calculated. The two interpretations are mathematically equivalent, but the one used here makes the correspondence with fair division more direct by removing any need to reason about payments.

Lambert et al. (2008) introduced a set of properties desirable for wagering mechanisms, which we translate into our security-based notation here. Individual rationality says that, under the assumption that agents are risk neutral and have immutable beliefs, agents don't lose money in expectation and therefore would willingly participate in the mechanism.

Definition 5. A wagering mechanism \mathcal{B} is *individually rational* if, for any agent i with beliefs \mathbf{p}_i , there exists a report $\hat{\mathbf{p}}_i$ such that for all $\hat{\mathbf{p}}_{-i}$ and \mathbf{w} , $\sum_{j=1}^m p_{i,j} \mathcal{B}_{i,j}(\hat{\mathbf{p}}, \mathbf{w}) \geq w_i$.

Strict budget balance guarantees that the principal never makes or loses money. Weak budget balance allows the principal to profit, but never record a loss.

Definition 6. A wagering mechanism \mathcal{B} is *weakly budget balanced* if, for all $\hat{\mathbf{p}}$ and \mathbf{w} , and for all $j \in [m]$, $\sum_{i=1}^n \mathcal{B}_{i,j}(\hat{\mathbf{p}}, \mathbf{w}) \leq W$. A wagering mechanism is *strictly budget balanced* if the inequality holds with equality for every $j \in [m]$.

Incentive compatibility says that agents should have incentive to truthfully report their beliefs.

Definition 7. A wagering mechanism \mathcal{B} is *weakly incentive compatible* if, for every agent i with beliefs \mathbf{p}_i , and all reports $\hat{\mathbf{p}}$ and wagers \mathbf{w} , $\sum_{j=1}^m p_{i,j} \mathcal{B}_{i,j}((\mathbf{p}_i, \hat{\mathbf{p}}_{-i}), \mathbf{w}) \geq \sum_{j=1}^m p_{i,j} \mathcal{B}_{i,j}(\hat{\mathbf{p}}, \mathbf{w})$. A wagering mechanism is *strictly incentive compatible* if the inequality is strict whenever $\mathbf{p}_i \neq \hat{\mathbf{p}}_i$.

Efficiency is also a concern for wagering mechanisms. A specific notion of Pareto optimality, which we term *side-bet Pareto optimality* to distinguish it from the corresponding fair division notion, was defined in Freeman et al. (2017). It says that the wagering mechanism should facilitate all available trade. To formally define this, we first need the notion of a profitable side bet.

Definition 8. Given reports $\hat{\mathbf{p}}$ and allocations $b_{i,j}$ of type- j securities to each agent i , we say that $\Delta \mathbf{b}$ is a *profitable side bet* if the following conditions hold:

1. For all $j \in [m]$, $\sum_{i=1}^n \Delta b_{i,j} = 0$.
2. For all $i \in [n]$, $\min_{j \in [m]} (b_{i,j} + \Delta b_{i,j}) \geq 0$.
3. For all $i \in [n]$, $\sum_{j=1}^m \hat{p}_{i,j} \Delta b_{i,j} \geq 0$, with strict inequality for at least one i .

Definition 9. A wagering mechanism \mathcal{B} is *side-bet Pareto optimal* if, for all reports $\hat{\mathbf{p}}$ and wagers \mathbf{w} , the allocation $\mathcal{B}(\hat{\mathbf{p}}, \mathbf{w})$ does not permit a profitable side bet.

Finally, we note that envy-freeness has been defined in the context of wagering (Freeman and Pennock, 2018), but in such a way that an agent i by definition cannot envy another agent j if j 's maximum loss is greater than i 's wager. While this definition makes sense for wagering, it does not lend itself to a natural interpretation for allocation.

3 The Correspondence

In this section, we show that there is a one-to-one correspondence between weakly budget-balanced wagering mechanisms and allocation mechanisms for divisible goods. This correspondence builds on the intuition that the output of a wagering mechanism is a division of m “goods” among the agents, where each good represents a security that pays off $W = \sum_{i=1}^n w_i$ if a particular outcome occurs. A bettor's value for the type- j security is proportional to her estimate of the probability that event j will occur, and wagers can be viewed as analogs of the priority weights used for allocations.

Definition 10 (Corresponding mechanisms). We say that an allocation mechanism \mathcal{A} and wagering mechanism \mathcal{B} are *corresponding mechanisms* if for all $\hat{\mathbf{p}}, \mathbf{w}$, $i \in [n]$, and $j \in [m]$,

$$\mathcal{B}_{i,j}(\hat{\mathbf{p}}, \mathbf{w}) = W \mathcal{A}_{i,j}(\hat{\mathbf{p}}, \mathbf{w}),$$

or equivalently, if for all $\hat{\mathbf{v}}, \mathbf{e}$, $i \in [n]$, and $j \in [m]$,

$$\mathcal{A}_{i,j}(\hat{\mathbf{v}}, \mathbf{e}) = \frac{\mathcal{B}_{i,j}(\hat{\mathbf{v}}, \mathbf{e})}{E}.$$

If \mathcal{A} is a valid allocation mechanism, then its corresponding wagering mechanism satisfies $\mathcal{B}_{i,j}(\hat{\mathbf{p}}, \mathbf{w}) \geq 0$ and for all $\hat{\mathbf{p}}, \mathbf{w}$, and $j \in [m]$,

$$\sum_{i=1}^n \mathcal{B}_{i,j}(\hat{\mathbf{p}}, \mathbf{w}) = \sum_{i=1}^n W \mathcal{A}_{i,j}(\hat{\mathbf{p}}, \mathbf{w}) \leq W.$$

It is therefore a well-defined, weakly budget-balanced wagering mechanism.

Similarly, if \mathcal{B} is a valid weakly budget-balanced wagering mechanism, then the corresponding allocation mechanism is valid since it satisfies $\mathcal{A}_{i,j}(\hat{\mathbf{v}}, \mathbf{e}) \geq 0$ and for all $\hat{\mathbf{v}}, \mathbf{e}$, and $j \in [m]$,

$$\sum_{i=1}^n \mathcal{A}_{i,j}(\hat{\mathbf{v}}, \mathbf{e}) = \sum_{i=1}^n \frac{\mathcal{B}_{i,j}(\hat{\mathbf{v}}, \mathbf{e})}{E} \leq 1.$$

This correspondence highlights the parallels between properties traditionally studied in the context of wagering and properties more often studied in the context of fair division, as we explore next.

	Allocation		Both				Wagering	
	Envy-Free	Pareto Optimal	Proportional/ Ind. Rational	Incentive Compatible	General m	General e / w	Budget Balanced	Side-bet Pareto Opt.
ME	Yes	Yes	Yes	No	Yes	Yes	Strict	Yes
WSWM	Yes	No	Yes	Strict	Yes	Yes	Strict	No
NAWM	No	No	Yes	Strict	Yes	Yes	Weak	No
DCA	Eq. weights	No	Yes	Weak	No	Yes	Weak	No
CSD	No	No	Yes	Weak	Yes	No	Strict	No
PA	Eq. weights	No	No	Weak	Yes	Yes	Weak	Yes
SDM	Yes	No	No	Weak	Yes	No	Weak	Yes

Table 1: Comparison of properties satisfied by market equilibrium (ME), weighted score wagering mechanisms (WSWM), no-arbitrage wagering mechanisms (NAWM), the double clinching auction (DCA), constrained serial dictatorship (CSD), partial allocation (PA), and strong demand matching (SDM).

3.1 Corresponding Properties

We first show an equivalence between incentive compatibility in both settings. A misreport that produces a more preferred allocation of goods also produces a more preferred allocation of securities. All omitted proofs are included in the appendix.

Theorem 1. *A random assignment mechanism \mathcal{A} is (weakly, strictly) incentive compatible if and only if the corresponding wagering mechanism \mathcal{B} is (weakly, strictly) incentive compatible.*

We next show an equivalence between individual rationality and proportionality. Though it is not typically described this way, the equivalence highlights that individual rationality can be viewed as guaranteeing an agent her “fair share” of the collected wagers. The restriction to incentive compatible mechanisms stems from a minor difference between the most common variants of the definitions in the two settings: proportionality is defined with respect to reported values, whereas individual rationality is defined with respect to true (not reported) beliefs.

Theorem 2. *A (weakly) incentive compatible random assignment mechanism \mathcal{A} is proportional if and only if the corresponding (weakly) incentive compatible wagering mechanism \mathcal{B} is individually rational.*

Pareto optimality in the wagering setting—what we term side-bet Pareto optimality—does not always imply Pareto optimality in the fair division setting. This is because side-bet Pareto optimality does not require the principal to allocate the maximum number of securities possible. The principal may instead keep a profit for himself. In the fair division setting, this maps to a wasteful allocation that leaves some (fractional) goods unallocated, which is not permitted under Pareto optimality. To obtain the same guarantee in the wagering setting, we must require the mechanism to be strictly budget-balanced, precluding any profit for the principal.

Theorem 3. *A random assignment mechanism \mathcal{A} is Pareto optimal if and only if the corresponding wagering mechanism \mathcal{B} is both side-bet Pareto optimal and strictly budget balanced.*

4 A Comparison of Mechanisms

The correspondence described in the previous section implies that any wagering mechanism can be viewed as an allocation mechanism and vice versa. In this section, we explore several immediate implications of this equivalence. In particular, we describe a variety of mechanisms that have been proposed in each setting and examine their potential in the other. Table 1 summarizes the properties of these mechanisms.

We start by pointing out a pair of equivalent mechanisms that have been explored in both settings.

Market Equilibrium (ME).

In the context of allocation, the market equilibrium solution endows each agent i with e_i units of currency and simulates a competitive equilibrium in which each good has a price and all agents spend their entire budget on goods that maximize their utility:price ratio (Walras, 2013). In the wagering context, this is equivalent to the parimutuel consensus mechanism of Eisenberg and Gale (1959).

In the allocation setting, the market equilibrium solution is envy-free, proportional, and Pareto optimal, but not incentive compatible. In the wagering setting, the parimutuel consensus mechanism is individually rational, strictly budget balanced, and side-bet Pareto optimal, but of course still fails to be incentive compatible.

4.1 Wagering Mechanisms

We next explore the implications of exporting existing wagering mechanisms into the fair division setting.

Weighted Score Wagering Mechanisms (WSWM).

The class of weighted score wagering mechanisms (Lambert et al., 2008, 2015), are built on proper scoring rules, payment functions that elicit truthful predictions from individual agents (Savage, 1971; Gneiting and Raftery, 2007). A scoring rule s maps a prediction $p \in [0, 1]^m$ and an outcome $j \in [m]$ to a score in $\mathbb{R} \cup \{-\infty\}$. We say s is *proper* if for all $\mathbf{p}, \mathbf{q} \in [0, 1]^m$, $\sum_{j=1}^m p_j s(\mathbf{p}, j) \geq \sum_{j=1}^m p_j s(\mathbf{q}, j)$, and *strictly proper* if this inequality is strict for $\mathbf{p} \neq \mathbf{q}$.

Fixing any strictly proper scoring rule s bounded in $[0, 1]$, the net payoff of a weighted score wagering mechanism is defined as

$$\Pi_i(\hat{\mathbf{p}}, \mathbf{w}, j) = w_i \left(s(\hat{\mathbf{p}}_i, j) - \frac{\sum_{i'=1}^n s(\hat{\mathbf{p}}_{i'}, j) w_{i'}}{W} \right).$$

Weighted score wagering mechanisms are individually rational, strictly incentive compatible, and strictly budget balanced, but not side-bet Pareto optimal.

Theorem 4. *An allocation mechanism corresponding to a weighted score wagering mechanism is proportional, strictly incentive compatible, and envy-free, but not Pareto optimal.*

Proportionality, strict incentive compatibility, and the absence of Pareto optimality follow from the correspondence. Envy-freeness holds because each agent’s expected utility is an affine transformation of her expected score according to s (and each agent’s score undergoes the same transformation). By properness of s , every agent believes her own score to be highest in expectation, so does not envy any other agent’s allocation.

To our knowledge, allocation mechanisms in this class are the first strictly incentive compatible allocation mechanisms, and the first non-trivial allocation mechanisms that are proportional, incentive compatible, and envy-free.

No-Arbitrage Wagering Mechanisms (NAWM).

No-arbitrage wagering mechanisms (Chen et al., 2014) are defined as follows. Let s be a strictly proper scoring rule bounded in $[0, 1]$, and let $\bar{p} : [0, 1]^{m \times (n-1)} \times \mathbb{R}^{n-1} \rightarrow [0, 1]^m$ be a function that maps reports and wagers of $n - 1$ agents to an “average” prediction. A no-arbitrage wagering mechanism defines net payoffs

$$\Pi_i(\hat{\mathbf{p}}, \mathbf{w}, j) = \frac{w_i W_{-i}}{W} [s(\hat{\mathbf{p}}_i, j) - s(\bar{p}(\hat{\mathbf{p}}_{-i}, \mathbf{w}_{-i}), j)].$$

For certain choices of \bar{p} , such mechanisms are weakly budget balanced, individually rational, and strictly incentive compatible, but not side-bet Pareto optimal.

Theorem 5. *An allocation mechanism corresponding to a no-arbitrage wagering mechanism is proportional and strictly incentive compatible, but not envy-free or Pareto optimal.*

Again, proportionality, incentive compatibility, and lack of Pareto optimality follow immediately from the properties of the corresponding wagering mechanism. The following example illustrates a violation of envy-freeness.

Example 1 (NAWM violates envy-freeness). Consider the NAWM with quadratic scoring rule ($s(p, x) = 1 - (p - x)^2$), and $\bar{p}(\mathbf{p}_{-i}, \mathbf{w}_{-i}) = \sum_{i' \neq i} \frac{w_{i'} \hat{p}_{i'}}{W_{-i}}$ be the weighted mean of other agents' reports.

Let $n = 3$ and $m = 2$, and suppose that reports are given by $p_1 = (0, 1)$, $p_2 = (0.9, 0.1)$, and $p_3 = (1, 0)$, and wagers $\mathbf{w} = (1, 1, 1)$. If $X = 1$, the net payoff of agent 2 is $\frac{2}{3}(0.99 - 0.75) = 0.16$, and if $X = 0$ the net payoff of agent 1 is $\frac{2}{3}(0.19 - 0.75) = -0.37$. For agent 3, the payoffs are $\frac{2}{3}(1 - (1 - 0.55^2)) = 0.2017$ and $\frac{2}{3}(0 - (1 - 0.45^2)) = -0.5317$ respectively.

The expected payoff that agent 2 achieves is therefore $0.9 \cdot 0.16 - 0.1 \cdot 0.37 = 0.1067$, while she would achieve expected payoff $0.9 \cdot 0.2017 - 0.1 \cdot 0.5317 = 0.1284$. Therefore, agent 2 envies agent 3.

The Double Clinching Auction (DCA).

The double clinching auction (Freeman et al., 2017) carefully selects a number of securities to allocate and then allocates them via an adaptive clinching auction (Dobzinski et al., 2008). We refer the reader to Freeman et al. (2017) for additional details, but note that the double clinching auction is only defined for the case of $m = 2$.

The double clinching auction is individually rational, weakly budget balanced, and weakly incentive compatible, but not side-bet Pareto optimal.

Theorem 6. *The allocation mechanism corresponding to the double clinching auction is proportional and weakly incentive compatible, but not envy-free or Pareto optimal.*

The double clinching auction violates envy-freeness on the following example.

Example 2 (DCA violates envy-freeness). Let $n = 5$ and $m = 2$, and suppose $p_1 = p_2 = (0.2, 0.8)$, $p_3 = (0.8, 0.2)$, $p_4 = p_5 = (0.81, 0.19)$, and $\mathbf{w} = (2, 2, 2, 2, 1)$. Then the double clinching auction allocates 1.25 securities of type 1 to agent 5, and 2.66 securities of type 1 to agent 4 (and no securities of type 2 to either). Therefore agent 5 envies agent 4.

4.2 Allocation Mechanisms

We next consider exporting existing allocation mechanisms into the wagering setting.

Constrained Serial Dictatorship (CSD).

Constrained serial dictatorship (Aziz and Ye, 2014) is defined only for equal weights ($e_i = 1$ for all i). Imagine fixing a particular ordering of the agents, allocating to each agent in order her most preferred m/n of the remaining goods (allowing partial goods to be chosen). The output of constrained serial dictatorship is the expected allocation that would arise from this process from an ordering chosen uniformly at random.

Constrained serial dictatorship is weakly incentive compatible and proportional, but not envy-free or Pareto optimal. There is no known natural extension for unequal weights, so agent duplication must be employed.

Theorem 7. *The wagering mechanism corresponding to constrained serial dictatorship is strictly budget balanced, individually rational, and weakly incentive compatible. It is not side-bet Pareto optimal.*

Strict budget balance holds because the allocation mechanism always allocates goods completely.

Interestingly, the constrained serial dictatorship allocation depends only on ordinal, not cardinal, preferences. Thus it yields a wagering mechanism that can be used for agents who report only ordinal beliefs (outcome j is more likely than outcome j').

Computing the constrained serial dictatorship allocation is #P-complete (Aziz et al., 2013, 2015), but a standard Chernoff and union bound argument shows that an arbitrary additive approximation can be found via sampling. We adopt this approach in our simulations.

	Value Ratio	Fraction of Securities Allocated	Fraction of Agents with an IR Violation	Value Ratio for Agents with an IR Violation
WSWM	1.092	1.000	0	n/a
NAWM	1.046	0.954	0	n/a
DCA	1.326	0.997	0	n/a
CSD	1.168	1.000	0.004	0.997
PA	0.490	0.369	0.995	0.487
SDM	1.329	0.999	0.042	0.999

Table 2: Comparison of mechanisms on the full ProbabilitySports dataset with uniform wagers.

Partial Allocation (PA).

The partial allocation mechanism (Cole et al., 2013) is designed to approximate market equilibrium, providing each agent at least a $1/e$ fraction of the utility she would receive in the market equilibrium allocation. Each agent i is allocated some $f_i \leq 1$ fraction of her market equilibrium allocation, where f_i is determined according to how costly agent i 's presence is to the other agents. See Cole et al. (2013) for further details.

Partial allocation is known to be weakly incentive compatible and envy-free when the weights of all agents are equal. However, it is not envy-free when weights are unequal, and it is not proportional or Pareto optimal.

Theorem 8. *The wagering mechanism corresponding to the partial allocation mechanism is weakly incentive compatible, weakly budget balanced, and side-bet Pareto optimal. It is not individually rational.*

Strong Demand Matching (SDM).

Strong demand matching (Cole et al., 2013) is also designed to approximate market equilibrium. Its approximation factor is particularly good when all items are highly demanded, for instance, when there are many more agents than goods and no good is uniformly disliked. It works by computing minimal prices at which each agent's complete demand (when she has a single unit of currency to spend) can be met by a single good. Again, we refer the reader to Cole et al. (2013) for further details.

Strong demand matching is envy-free and weakly incentive-compatible, but not proportional or Pareto optimal. There is no known, natural way to extend the mechanism to unequal weights that retains the desirable properties, other than the agent-duplication method.

Theorem 9. *The wagering mechanism corresponding to strong demand matching is weakly budget balanced, weakly incentive compatible, and side-bet Pareto optimal. It is not individually rational.*

5 Simulations

We next compare the performance of these mechanisms in practice. Our test data comes from a binary outcome ($m = 2$) wagering setting, but the equivalence allows us to make observations about allocation-specific properties like envy-freeness too. In the appendix, we additionally report results from simulations with larger outcome spaces using synthetic data. We restrict attention to incentive compatible mechanisms, excluding market equilibrium.

We use reports gathered from an online prediction contest called ProbabilitySports (Galebach, 2017). The dataset consists of probabilistic predictions about the outcomes of 1643 U.S. National Football League matches between 2000 and 2004. For each match, between 64 and 1574 people reported their subjective probability of the home team winning the game and earned points according to a Brier scoring rule.

ProbabilitySports participants did not submit wagers. For our simulations, we follow previous work (Freeman et al., 2017) and generate wagers in two ways. First, we consider uniform wagers, modeling the scenario in which agents are required to risk the same amount. Second, for smaller instances on which it is tractable to employ the agent duplication method for unequal weights, we generate wagers according to a Pareto distribution with shape parameter

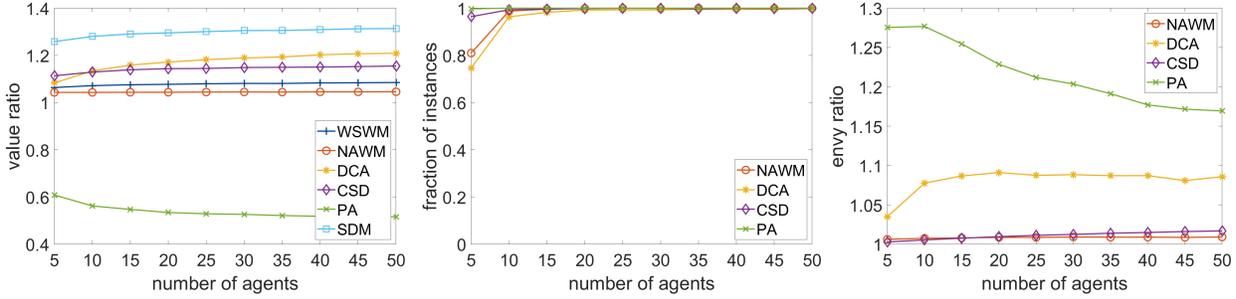


Figure 1: Value ratio (left), fraction of instances with envy (center), and average maximum envy across all instances with envy (right) on a subsample of the ProbabilitySports data with Pareto distribution wagers.

1.16 and scale parameter 1. Since agent duplication requires rational wagers, we scale the generated wagers to lie in $[0, 50]$, and then take the ceiling of each, yielding integral wagers between 1 and 50.

Table 2 summarizes the results of our first simulation, in which the mechanisms were run on the complete ProbabilitySports dataset with uniform wagers. Weighted score and no-arbitrage wagering mechanisms utilized the Brier score (Brier, 1950). We report four statistics, all averaged over the 1643 matches considered. *Value ratio* is the sum over all agents of the expected value of their allocated securities (with respect to their own subjective beliefs) normalized by the total amount wagered, or $\frac{1}{W} \sum_{i=1}^n \sum_{j=1}^m \hat{p}_{i,j} \mathcal{B}_{i,j}(\hat{\mathbf{p}}, \mathbf{w})$. In the allocation setting, the value ratio is simply the (normalized) social welfare, a natural measure of efficiency. In the wagering setting, a high value ratio can be seen as encouraging participation, since forecasters will be more likely to participate when they expect to make a higher profit. High participation is crucial to harnessing the wisdom-of-crowds principle and constructing a good aggregate forecast. *Fraction of securities allocated* is the quantity $\frac{1}{mW} \sum_{i=1}^n \sum_{j=1}^m \mathcal{B}_{i,j}(\hat{\mathbf{p}}, \mathbf{w})$. In the allocation setting, this is a measure of waste, with lower values implying higher waste, while in the wagering setting any “wasted” securities become mechanism profit when the corresponding outcome occurs. *Fraction of agents with an IR violation* is the fraction of agents for whom the expected value of their security allocation is less than the value of their wager. The final column is similar to value ratio, but calculated only over agents with an individual rationality violation, giving a sense of the magnitude of the violations.

Judging by these metrics, the double clinching auction (DCA) and strong demand matching (SDM) show the strongest performance, with comparably high value ratios, negligible waste, and negligible violations of individual rationality (none in the case of DCA). Despite their strong theoretical guarantees but in line with previous observations (Freeman et al., 2017), weighted score (WSWM) and no-arbitrage (NAWM) wagering do not facilitate much trade, creating little value for the agents. At the other end of the spectrum, partial allocation (PA) suffers poor performance across the board, violating individual rationality for almost all agents, and often substantially so. We note that the small number of individual rationality violations observed for constrained serial dictatorship (CSD) stem from the fact that CSD allocations can only be computed approximately.

To understand how performance scales with the size of the instance, we next subsampled n agent reports and generated wagers for each $n \in \{5, 10, \dots, 50\}$. Figure 1 shows a selection of the results for agents with Pareto distribution wagers. More are included in the appendix. Results for uniform wagers were similar, though the double clinching auction and partial allocation exhibit no envy in those cases.

We see similar trends in value ratios (left plot), but with strong demand matching exhibiting a clear advantage over the double clinching auction. The center plot shows the fraction of instances for which at least one agent envies another, limited to those mechanisms that are not guaranteed to be envy-free. In all cases, envy almost always exists. For a pair of agents (i, j) , define the envy ratio as the expected utility that i would receive if she were allocated j ’s securities, divided by the expected utility she receives from her own allocation, and scaled to adjust for wagers. The right plot shows the maximum envy ratio on instances for which at least one agent envies another. We see that partial allocation often leads to significant amounts of envy.

6 Discussion

While we have focused on the positive implications of the equivalence between wagering and fair division mechanisms, we note that it immediately implies new impossibility results as well. For instance, a classical result of Schummer (1996) in the allocation setting says that the only incentive compatible and Pareto optimal mechanisms are dictatorial. This result carries over to the wagering setting, and complements the impossibility result of Freeman et al. (2017) that no weakly budget-balanced wagering mechanism can simultaneously satisfy incentive compatibility, side-bet Pareto optimality, and individual rationality. As another example, Han et al. (2011) showed that, for large enough n , no incentive compatible allocation mechanism can achieve better than a $1/m$ approximation to the optimal social welfare. This result also carries over to wagering.

Despite the equivalence, we expect there to be continued value in studying the two problems in isolation. Different design criteria have traditionally been emphasized for the two—most notably, strict incentive compatibility (drawing heavily on the scoring rule literature) for wagering and envy-freeness for fair division—and different mechanisms may be more appropriate in one setting than the other. However, the connection we revealed opens up the possibility of applying any new results developed in one setting to the other, an approach we expect to be fruitful as both fields develop.

References

- Haris Aziz and Chun Ye. Cake cutting algorithms for piecewise constant and piecewise uniform valuations. In *Proceedings of the 10th Conference on Web and Internet Economics (WINE)*, pages 1–14, 2014.
- Haris Aziz, Felix Brandt, and Markus Brill. The computational complexity of random serial dictatorship. *Economics Letters*, 121(3):341–345, 2013.
- Haris Aziz, Felix Brandt, Markus Brill, and Julián Mestre. Computational aspects of random serial dictatorship. *ACM SIGecom Exchanges*, 13(2):26–30, 2015.
- Steven J. Brams and Alan D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.
- Glenn W. Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1):1–3, 1950.
- Yiling Chen, Nikhil R. Devanur, David M. Pennock, and Jennifer Wortman Vaughan. Removing arbitrage from wagering mechanisms. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, 2014.
- Richard Cole, Vasilis Gkatzelis, and Gagan Goel. Mechanism design for fair division: allocating divisible items without payments. In *Proceedings of the 14th ACM Conference on Economics and Computation (EC)*, pages 251–268, 2013.
- Shahar Dobzinski, Ron Lavi, and Noam Nisan. Multi-unit auctions with budget limits. In *FOCS*, pages 260–269, 2008.
- Miroslav Dudík, Sébastien Lahaie, Ryan M Rogers, and Jennifer Wortman Vaughan. A decomposition of forecast error in prediction markets. In *Advances in Neural Information Processing Systems*, pages 4371–4380, 2017.
- Edmund Eisenberg and David Gale. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959.
- Duncan Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:45–98, 1967.
- Rupert Freeman and David M Pennock. An axiomatic view of the parimutuel consensus mechanism. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 254–260, 2018.
- Rupert Freeman, David M Pennock, and Jennifer Wortman Vaughan. The double clinching auction for wagering. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 43–60, 2017.
- Brian Galebach. ProbabilitySports. <http://probabilitysports.com/>, 2017.
- Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- Li Han, Chunzhi Su, Linpeng Tang, and Hongyang Zhang. On strategy-proof allocation without payments or priors. In *Proceedings of the 7th Conference on Web and Internet Economics (WINE)*, pages 182–193, 2011.
- D. Marc Kilgour and Yigal Gerchak. Elicitation of probabilities using competitive scoring rules. *Decision Analysis*, 1(2):108–113, 2004.
- Nicolas S. Lambert, John Langford, Jennifer Wortman, Yiling Chen, Daniel M. Reeves, Yoav Shoham, and David M. Pennock. Self-financed wagering mechanisms for forecasting. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*, pages 170–179, 2008.
- Nicolas S. Lambert, John Langford, Jennifer Wortman Vaughan, Yiling Chen, Daniel Reeves, Yoav Shoham, and David M. Pennock. An axiomatic characterization of wagering mechanisms. *Journal of Economic Theory*, 156:389–416, 2015.

- Ariel D. Procaccia. Cake cutting: Not just child's play. *Communications of the ACM*, 56(7):78–87, 2013.
- Jack M. Robertson and William A. Webb. *Cake Cutting Algorithms: Be Fair If You Can*. A. K. Peters, 1998.
- Leonard J. Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.
- James Schummer. Strategy-proofness versus efficiency on restricted domains of exchange economies. *Social Choice and Welfare*, 14(1):47–56, 1996.
- Hugo Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.
- Leon Walras. *Elements of pure economics*. Routledge, 2013.

A Omitted Proofs from Section 3

A.1 Proof of Theorem 1

First, suppose that \mathcal{B} is weakly incentive compatible. Let $i \in [n]$ and suppose that $\hat{\mathbf{v}}_i \neq \mathbf{v}_i$. Then

$$\begin{aligned}
v_i(\mathcal{A}_i((\mathbf{v}_i, \hat{\mathbf{v}}_{-i}), \mathbf{e})) &= \sum_{j=1}^m v_{i,j} \mathcal{A}_{i,j}((\mathbf{v}_i, \hat{\mathbf{v}}_{-i}), \mathbf{e}) \\
&= \sum_{j=1}^m v_{i,j} \frac{\mathcal{B}_{i,j}((\mathbf{v}_i, \hat{\mathbf{v}}_{-i}), \mathbf{e})}{E} \\
&\geq \sum_{j=1}^m v_{i,j} \frac{\mathcal{B}_{i,j}(\hat{\mathbf{v}}, \mathbf{e})}{E} \\
&= \sum_{j=1}^m v_{i,j} \mathcal{A}_{i,j}(\hat{\mathbf{v}}, \mathbf{e}) \\
&= v_i(\mathcal{A}_i(\hat{\mathbf{v}}, \mathbf{e})).
\end{aligned}$$

The second and fourth equalities follows from the definition of the correspondence, and the inequality from weak incentive compatibility of \mathcal{B} . Note that if \mathcal{B} is strictly incentive compatible then the inequality is strict, implying that \mathcal{A} is strictly incentive compatible.

Next, suppose that \mathcal{A} is weakly incentive compatible. Let $i \in [n]$ and suppose that $\hat{\mathbf{p}}_i \neq \mathbf{p}_i$. We have:

$$\begin{aligned}
&\sum_{j=1}^m p_{i,j} \mathcal{B}_{i,j}((\mathbf{p}_i, \hat{\mathbf{p}}_{-i}), \mathbf{w}) \\
&= W \sum_{j=1}^m p_{i,j} \mathcal{A}_{i,j}((\mathbf{p}_i, \hat{\mathbf{p}}_{-i}), \mathbf{w}) \\
&\geq W \sum_{j=1}^m p_{i,j} \mathcal{A}_{i,j}(\hat{\mathbf{p}}, \mathbf{w}) \\
&= \sum_{j=1}^m p_{i,j} \mathcal{B}_{i,j}(\hat{\mathbf{p}}, \mathbf{w}).
\end{aligned}$$

The equalities follows from the definition of the correspondence, and the inequality from weak incentive compatibility of \mathcal{A} . Once again if \mathcal{A} is strictly incentive compatible then the inequality is strict, implying that \mathcal{B} is strictly incentive compatible. □

A.2 Proof of Theorem 2

First, suppose that \mathcal{B} is individually rational and weakly incentive compatible. Let $i \in [n]$. We have:

$$\begin{aligned}\hat{v}_i(\mathcal{A}_i(\hat{\mathbf{v}}, \mathbf{e})) &= \sum_{j=1}^m \hat{v}_{i,j} \cdot \mathcal{A}_{ij}(\hat{\mathbf{v}}, \mathbf{e}) \\ &= \sum_{j=1}^m \hat{v}_{i,j} \left(\frac{\mathcal{B}_{j,i}(\hat{\mathbf{v}}, \mathbf{e})}{E} \right) \geq \frac{e_i}{E}.\end{aligned}$$

The second equality follows from the definition of the correspondence, and the inequality from individual rationality (there exists a report for which the inequality holds) and weak incentive compatibility (no report yields higher utility than the truthful report).

Next, suppose that \mathcal{A} is proportional. Let $i \in [n]$, and let $\hat{\mathbf{p}}_i = \mathbf{p}_i$. We have:

$$\begin{aligned}\sum_{j=1}^m p_{i,j} \mathcal{B}_{i,j}(\hat{\mathbf{p}}, \mathbf{w}) &= W \sum_{j=1}^m \hat{p}_{i,j} \mathcal{A}_{i,j}(\hat{\mathbf{p}}, \mathbf{w}) \\ &\geq W \frac{w_i}{W} = w_i.\end{aligned}$$

The first equality follows from the definition of the correspondence and the inequality from proportionality.

A.3 Proof of Theorem 3

We prove the “only if” direction first. Suppose that \mathcal{B} is not side-bet Pareto optimal. That is, there exist reports $\hat{\mathbf{p}}$ and wagers \mathbf{w} , producing allocation $\mathcal{B}(\hat{\mathbf{p}}, \mathbf{w})$, for which there exists a profitable side bet $\Delta \mathbf{b}$.

Consider the allocation $\mathcal{A}(\hat{\mathbf{p}}, \mathbf{w})$. For every $i \in [n]$ and $j \in [m]$, we have

$$\mathcal{A}_{ij}(\hat{\mathbf{p}}, \mathbf{w}) = \frac{\mathcal{B}_{i,j}(\hat{\mathbf{p}}, \mathbf{w})}{W}.$$

We will show that $\mathcal{A}(\hat{\mathbf{p}}, \mathbf{w})$ is not Pareto optimal by considering an alternative allocation defined by

$$a_{i,j} = \frac{\mathcal{B}_{i,j}(\hat{\mathbf{p}}, \mathbf{w}) + \Delta b_{i,j}}{W}$$

for every $i \in [n]$ and $j \in [m]$. For every $i \in [n]$, we have

$$\sum_{j=1}^m \hat{p}_{i,j} \mathcal{A}_{i,j}(\hat{\mathbf{p}}, \mathbf{w}) = \sum_{j=1}^m \hat{p}_{i,j} \frac{\mathcal{B}_{i,j}(\hat{\mathbf{p}}, \mathbf{w})}{W}$$

and

$$\sum_{j=1}^m \hat{p}_{i,j} a_{i,j}(\hat{\mathbf{p}}, \mathbf{w}) = \sum_{j=1}^m \hat{p}_{i,j} \frac{\mathcal{B}_{i,j}(\hat{\mathbf{p}}, \mathbf{w})}{W} + \sum_{j=1}^m \hat{p}_{i,j} \frac{\Delta b_{i,j}}{W}.$$

Taking the difference,

$$\begin{aligned}\sum_{j=1}^m \hat{p}_{i,j} a_{i,j}(\hat{\mathbf{p}}, \mathbf{w}) - \sum_{j=1}^m \hat{p}_{i,j} \mathcal{A}_{i,j}(\hat{\mathbf{p}}, \mathbf{w}) \\ = \sum_{j=1}^m \hat{p}_{i,j} \frac{\Delta b_{i,j}}{W} \geq 0,\end{aligned}$$

where the inequality follows from Condition 3 in Definition 8, and holds strictly for some i . Therefore, \mathcal{A} is not Pareto optimal.

Next, suppose that \mathcal{B} is weakly budget balanced but not strictly budget balanced. Then there must exist a $j \in [m]$ for which

$$\sum_{i=1}^n \mathcal{A}_{ij}(\hat{\mathbf{v}}, \mathbf{e}) = \sum_{i=1}^n \frac{\mathcal{B}_{i,j}(\hat{\mathbf{v}}, \mathbf{e})}{E} < 1.$$

We may therefore define a Pareto-improving allocation in which the unallocated portion of good j is allocated to some agent who values it positively, while holding the rest of the allocation constant.

To prove the “if” direction, suppose that \mathcal{B} is side-bet Pareto optimal and strictly budget balanced, but that \mathcal{A} is not Pareto optimal. By strict budget balance, it must be the case that $\sum_{i=1}^n \mathcal{A}_{ij}(\hat{\mathbf{v}}, \mathbf{e}) = 1$ for all $j \in [m]$. Therefore, a Pareto-improving allocation consists of agents swapping (fractions of) goods among themselves; let $\Delta a_{i,j}$ denote the amount of good $j \in [m]$ that agent $i \in [n]$ gains or loses from this swap. It is easy to check that $\Delta \mathbf{a}_j$ corresponds to a profitable side bet, violating side-bet Pareto optimality of \mathcal{B} . \square

B Omitted Proofs from Section 4

B.1 Proof of Theorem 4

Proportionality and strict incentive compatibility follow from the properties of the corresponding wagering mechanism. To prove envy-freeness, let $i, i' \in [n]$. For all $\hat{\mathbf{v}}, \mathbf{e}$, we have

$$\begin{aligned} & \frac{\hat{v}_i(\mathcal{A}_i(\hat{\mathbf{v}}, \mathbf{e}))}{e_i} \\ &= \frac{\sum_{j=1}^m \hat{v}_{i,j} \mathcal{A}_{i,j}(\hat{\mathbf{v}}, \mathbf{e})}{e_i} \\ &= \frac{\sum_{j=1}^m \hat{v}_{i,j} \mathcal{B}_{i,j}(\hat{\mathbf{v}}, \mathbf{e})}{E e_i} \\ &= \frac{1}{E} \sum_{j=1}^m \hat{v}_{i,j} \left(1 + s(\hat{\mathbf{v}}_i, j) - \frac{\sum_{k=1}^n s(\hat{\mathbf{v}}_k, j) e_k}{E} \right) \\ &= \frac{1}{E} \left(1 + \sum_{j=1}^m \hat{v}_{i,j} s(\hat{\mathbf{v}}_i, j) - \sum_{j=1}^m \hat{v}_{i,j} \frac{\sum_{k=1}^n e_k s(\hat{\mathbf{v}}_k, j)}{E} \right) \\ &\geq \frac{1}{E} \left(1 + \sum_{j=1}^m \hat{v}_{i,j} s(\hat{\mathbf{v}}_{i'}, j) - \sum_{j=1}^m \hat{v}_{i,j} \frac{\sum_{k=1}^n e_k s(\hat{\mathbf{v}}_k, j)}{E} \right) \\ &= \frac{1}{E} \sum_{j=1}^m \hat{v}_{i,j} \left(1 + s(\hat{\mathbf{v}}_{i'}, j) - \frac{\sum_{k=1}^n s(\hat{\mathbf{v}}_k, j) e_k}{E} \right) \\ &= \frac{\sum_{j=1}^m \hat{v}_{i,j} \mathcal{B}_{i',j}(\hat{\mathbf{v}}, \mathbf{e})}{E e_{i'}} \\ &= \frac{\sum_{j=1}^m \hat{v}_{i,j} \mathcal{A}_{i',j}(\hat{\mathbf{v}}, \mathbf{e})}{e_{i'}} \\ &= \frac{\hat{v}_i(\mathcal{A}_{i'}(\hat{\mathbf{v}}, \mathbf{e}))}{e_{i'}}, \end{aligned}$$

where the inequality follows from properness of the scoring rule s ; agent i maximizes her score by reporting v_i . \square

B.2 Proof of Theorem 5

Proportionality, incentive compatibility, and (lack of) Pareto optimality follow immediately from the properties of the corresponding wagering mechanism. We present an example in which an allocation mechanism corresponding to a no arbitrage wagering mechanism violates envy-freeness.

Consider the NAWM that uses the quadratic scoring rule $s(\mathbf{p}, j) = \frac{1}{2} + p_j - \frac{1}{2} \sum_{j \in [m]} p_j^2$, and

$$\bar{p}(\hat{\mathbf{p}}_{-i}, \mathbf{w}_{-i}) = \left(\sum_{i' \neq i} \frac{w_{i'} \hat{p}_{i', j}}{W_{-i}} \right)_{j \in [m]}$$

be the weighted mean of other agents' reports.

Let $n = 3$ and $m = 2$, and suppose that reports are given by $\mathbf{p}_1 = (0, 1)$, $\mathbf{p}_2 = (0.9, 0.1)$, and $\mathbf{p}_3 = (1, 0)$, and wagers $\mathbf{w} = (1, 1, 1)$. If $X = 1$, the net payoff of agent 2 is $\frac{2}{3}(0.99 - 0.75) = 0.16$, and if $X = 0$ the net payoff of agent 2 is $\frac{2}{3}(0.19 - 0.75) = -0.37$. For agent 3, the payoffs are $\frac{2}{3}(1 - (1 - 0.55^2)) = 0.2017$ and $\frac{2}{3}(0 - (1 - 0.45^2)) = -0.5317$ respectively.

The expected payoff that agent 2 achieves is therefore $0.9 \cdot 0.16 - 0.1 \cdot 0.37 = 0.1067$, while she would achieve expected payoff $0.9 \cdot 0.2017 - 0.1 \cdot 0.5317 = 0.1284$ if she got 3's bundle of securities. Therefore, agent 2 envies agent 3. □

B.3 Proof of Theorem 6

Proportionality, incentive compatibility, and (lack of) Pareto optimality follow immediately from the properties of the corresponding wagering mechanism. We present an example in which the allocation mechanism corresponding to the double clinching auction violates envy-freeness.

Let $n = 5$ and $m = 2$, and suppose $\mathbf{p}_1 = \mathbf{p}_2 = (0.2, 0.8)$, $\mathbf{p}_3 = (0.8, 0.2)$, $\mathbf{p}_4 = \mathbf{p}_5 = (0.81, 0.19)$, and $\mathbf{w} = (2, 2, 2, 2, 1)$. Then the double clinching auction allocates 1.25 securities of type 1 to agent 5, and 2.66 securities of type 1 to agent 4 (and no securities of type 2 to either). Therefore agent 5 envies agent 4. □

B.4 Proof Sketch of Theorem 8

Weak incentive compatibility, weak budget balance, and lack of individual rationality follow immediately from the properties of the allocation mechanism. Nothing about side-bet Pareto optimality is immediately implied, since it is a weaker requirement than Pareto optimality. To see that it is satisfied, note that the partial allocation mechanism is equivalent to setting prices for the goods equal to the market equilibrium prices, but giving each agent only $f_i e_i$ units of currency to spend. (See Cole et al. (2013) for full details.) Agents still buy items that maximize their utility:price ratio. This fact is sufficient to conclude that there is no mutually beneficial trade that the agents can make using only goods that are allocated to them (that is, a side-bet). □

B.5 Proof Sketch of Theorem 9

Weak budget balance, weak incentive compatibility, and lack of individual rationality follow immediately from the properties of the allocation mechanism. Side-bet Pareto optimality holds for the same reason as for the partial allocation mechanism. The strong demand matching mechanism computes prices at which agents buy goods that maximize their utility:price ratio (and some goods are left unallocated). Therefore, there is no profitable side bet available. □

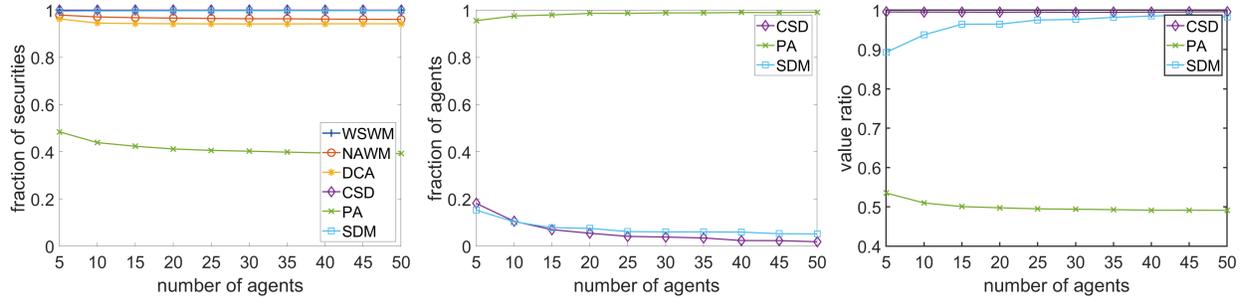


Figure 2: Fraction of securities allocated (left), fraction of agents with an individual rationality violation (center), and value ratio across all agents with an individual rationality violation (right) on a subsample of the ProbabilitySports data with Pareto distribution wagers.

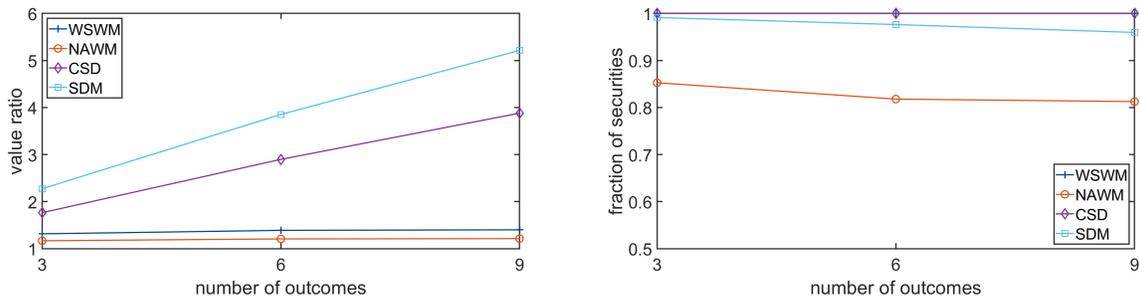


Figure 3: Value ratio (left) and fraction of securities allocated (right) for randomly generated data with $n = 30$.

C Additional Simulation Results

In Figure 2 we display some additional results for agents with Pareto wagers and beliefs subsampled from the ProbabilitySports dataset. The left plot shows the total fraction of securities allocated to the agents. We see that all mechanisms except partial allocation always allocate a high ($> 90\%$) fraction of securities. We note that for Pareto wagers, the double clinching auction allocates relatively fewer securities than on the full dataset with uniform wagers.

In the center plot, we consider the fraction of agents for whom the expected utility from the securities allocated to them by the mechanism is less than the utility they would receive from not participating. As on the full instances, close to 100% of agents experience an individual rationality violation under the partial allocation mechanism. For both the strong demand matching and constrained serial dictatorship mechanisms, there are a fairly high number of individual rationality violations for small n , with decreasing numbers for larger n . The right figure plots the value ratio calculated only over agents with an individual rationality violation. As on the full instances, the violations tend to be small for strong demand matching and constrained serial dictatorship, but large for partial allocation.

Next, we consider instances with larger outcome spaces. Since the ProbabilitySports dataset consists entirely of binary events, we generate synthetic data according to Dudík et al. (2017). To generate beliefs, we fix parameters $\theta^{\text{true}} = (\theta_j)_{j \in [m]} = \log(1 - 0.05(m - 1), 0.05, \dots, 0.05)$ and ground truth probabilities e^{θ_j} for each outcome j . For each agent i and each outcome j , we sample $x_{i,j}$ independently from $\text{Normal}(\theta_j, \sigma^2)$ with $\sigma = 3$. For every agent i , we set $\hat{p}_{i,j}$ proportional to $e^{x_{i,j}}$ so that each agent’s beliefs sum to 1.

For non-binary outcomes, the double clinching auction is not defined. Further, the partial allocation mechanism is significantly more computationally intensive for more than two outcomes. Since it does not perform well in the binary setting, we exclude it from these simulations.

In Figure 3 we consider the effect of increasing the number of outcomes on the value ratio and the number of securities allocated. On the left, we see that the value ratio for strong demand matching and constrained serial dictatorship increases quite dramatically with the number of outcomes. Weighted score and no arbitrage wagering, on the other

hand, exhibit no change. In the right plot, we see that for the weakly budget balanced mechanisms, the total fraction of securities allocated decreases slightly as the number of outcomes increases. Of course, the strictly budget balanced mechanisms continue to allocate all securities.