# An Equivalence Between Fair Division and Wagering Mechanisms

Rupert Freeman[†]
University of Virginia, Charlottesville, VA, USA, FreemanR@darden.virginia.edu
Jens Witkowski[†]
Frankfurt School of Finance & Management, Frankfurt, Germany, j.witkowski@fs.de
Jennifer Wortman Vaughan
Microsoft Research, New York, NY, USA, jenn@microsoft.com
David M. Pennock
Rutgers University, Piscataway, NJ, USA, dpennock@dimacs.rutgers.edu

We draw a surprising and direct mathematical equivalence between the class of fair division mechanisms, designed to allocate divisible goods without money, and the class of weakly budget-balanced wagering mechanisms, designed to elicit probabilities. While this correspondence between fair division and wagering has applications in both settings, we focus on its implications for the design of incentive-compatible fair division mechanisms. In particular, we show that applying the correspondence to Competitive Scoring Rules, a prominent class of wagering mechanisms based on proper scoring rules, yields the first incentive-compatible fair division mechanism that is both fair (proportional and envy-free) and responsive to agent preferences. Moreover, for two agents, we show that Competitive Scoring Rules characterize the whole class of non-wasteful and incentive-compatible fair division mechanisms, subject to mild technical conditions. As one of several consequences, this allows us to resolve an open question about the best possible approximation to optimal utilitarian welfare that can be achieved by any incentive-compatible mechanism. Finally, since the equivalence greatly expands the set of known incentive-compatible fair division mechanisms, we conclude with an evaluation of this entire set, comparing the mechanisms' axiomatic properties and examining their welfare performance in simulation.

## 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. The principal does so by collecting a wager from each agent along with the agent's prediction, and redistributing the wagers in such a way that agents with more accurate predictions are more highly rewarded. In the second scenario, a neutral mediator needs to allocate a set of divisible resources, such as produce donated to a food bank or compute cycles on a shared server, to a group of agents with differing preferences over the

[†] Rupert Freeman and Jens Witkowski contributed equally to the manuscript.

resources. The mediator's primary objective is ensuring that every agent walks away with her fair share.[1]

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 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 fair division 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 guarantees that every agent is at least as satisfied as if each good were divided equally among the agents. Similarly, *incentive compatibility*, a property common to the literature in both settings, carries over immediately. Finally, wagering mechanisms that satisfy *normality*, which, loosely speaking, means that an agent's payment can be interpreted as her performance relative to other agents, yield fair division mechanisms that are *envy-free*, meaning that no agent prefers any other agent's allocation to her own. (For this last pair of properties, the reverse direction does not hold.)

While this correspondence has implications for both fair division and wagering, we focus on the implications it has on the design of incentive-compatible fair division mechanisms. In particular, Competitive Scoring Rules (CSRs, Kilgour and Gerchak 2004, Lambert et al. 2008, 2015), a family of wagering mechanisms built on the machinery of proper scoring rules (e.g., Gneiting and Raftery 2007), can be imported into the fair division setting. To the best of our knowledge, this newly established family of fair division mechanisms is the first to jointly satisfy proportionality, envy-freeness, and incentive compatibility. Moreover, these mechanisms are also the first *strictly* incentive-compatible fair division mechanisms, i.e., agents do worse by misreporting than by truthfully reporting their preferences.

Generalizing a characterization result by Han et al. (2011) from two goods to any number of goods, we show that CSRs encompass the entire space of non-wasteful, incentive-compatible fair division mechanisms for two agents, subject to mild technical conditions. As a corollary, we strengthen an impossibility result pertaining to social welfare approximation. Exploring the characterization result and knowing that all non-wasteful, incentive-compatible mechanisms for two

[1] For clarity, we adopt the convention of using feminine pronouns for the agents.

agents must be CSRs, we cast several mechanisms from the literature in that framework and identify their corresponding proper scoring rules. Interestingly, the designers of some of these mechanisms turn out to have rediscovered classic scoring rules from the literature, such as the spherical scoring rule (Roby 1965, Jose 2009), whereas others invented new ones. Concluding the analysis of CSRs, we provide closed-form expressions for utilitarian and Nash welfare as functions of the proper scoring rule associated with a given CSR. As it turns out, there is a tight connection between these welfare measures and the scoring rule's expected loss function.

Due to the correspondence, the set of known incentive-compatible mechanisms expands greatly, so that we can now compare the existing mechanisms from the fair division literature with the new mechanisms gained through the correspondence. In doing so, we first provide a complete overview of the axiomatic properties of this updated set of known incentive-compatible mechanisms. Moreover, we also compare the mechanisms via simulation, evaluating their empirical performance according to utilitarian and Nash welfare. For the base case of two agents and two goods, we find that CSRs perform better than any of the wasteful mechanisms, including those with constant social welfare approximation guarantees. Furthermore, we observe that the best choice of CSR depends on context, confirming our theoretical analysis regarding the connection between loss functions and welfare. For the case where the number of agents is large relative to the number of goods, Strong Demand Matching (SDM, Cole et al. 2013b) performs best overall. However, for the special case of two goods, the Double Clinching Auction (DCA, Freeman et al. 2017), which was originally developed as a wagering mechanism, achieves similarly high welfare and, in contrast to SDM, also satisfies proportionality. In the setting where the number of goods is large relative to the number of agents, the Constrained Serial Dictatorship (CSD, Aziz and Ye 2014) proves surprisingly robust.

**Related Work**

The study of fair division dates back to the 1940s (Steinhaus 1948) and many variants of the problem have been examined (see Chapters 11–13 in the survey by Brandt et al. 2016). In its general form, fair division is concerned with allocating a set of resources among competing agents without monetary transfers. The goal is to achieve allocations that are fair to the agents while also allocating goods as efficiently as possible (Bertsimas et al. 2011, 2012). Arguably the most prominent fair division setting is that of cake cutting, in which the resource to be allocated is a single heterogenous good (Robertson and Webb 1998, Deng et al. 2012, Procaccia 2013).

We consider agents who report their valuations strategically, joining a growing line of work in fair division. Several papers derive impossibility results for designing deterministic, incentive-compatible, and fair mechanisms for cake cutting (Brânzei and Miltersen 2015, Menon and Larson 2017, Tao 2022). Others circumvent these results by designing randomized algorithms (Mossel and

Tamuz 2010, Chen et al. 2013) or considering special cases such as agents holding piecewise uniform valuations (Maya and Nisan 2012, Chen et al. 2013, Bei et al. 2020), a setting incomparable to ours.

In this paper, we focus on the problem of allocating divisible goods to agents with linear-additive valuation functions, a special case of cake cutting. For the special case of two agents, several incentive-compatible mechanisms focusing on approximating the optimal utilitarian welfare have been proposed (Guo and Conitzer 2010, Han et al. 2011, Cheung 2016, Cole et al. 2013a). We discuss these mechanisms in more detail in Section 5. In a similar spirit but for any number of agents, Cole et al. (2013b, 2022) design mechanisms that guarantee every agent a constant fraction of the utility they would achieve in the allocation that maximizes Nash welfare. Aziz and Ye (2014) introduce the Constrained Serial Dictatorship (CSD) mechanism for cake cutting, which, in the setting we study, corresponds to an incentive-compatible and proportional mechanism.

Beyond cake cutting, fair division is also concerned with the allocation of indivisible goods (e.g., Chevaleyre et al. 2006, Brânzei et al. 2022). Analogous to the literature on divisible goods, a line of work has focused on designing mechanisms for allocating indivisible goods under additive valuations. In particular, Caragiannis et al. (2009) provide a characterization of all incentive-compatible mechanisms for two agents and two goods in this space, which was later generalized to arbitrary numbers of goods by Amanatidis et al. (2017). This can be seen as analogous to our result in Section 5 characterizing all incentive-compatible mechanisms for two agents and *divisible* goods. For additional discussion of work focusing on allocating indivisible goods with strategic agents, see Section 6.2 of the survey by Amanatidis et al. (2022).

On the other side of the equivalence is the wagering setting, which dates back to the work of Eisenberg and Gale (1959), who studied the equilibrium behavior of bettors with probabilistic beliefs about horse race outcomes. More recently, Lambert et al. (2008, 2015) initiated the study of incentive-compatible wagering mechanisms, building on Competitive Scoring Rules (Kilgour and Gerchak 2004) to introduce Weighted Score Wagering Mechanisms (WSWMs). Chen et al. (2014) identify arbitrage opportunities in WSWMs and design a class of No-Arbitrage Wagering Mechanisms (NAWMs), which do not permit these opportunities. Freeman et al. (2017) point out that, in existing incentive-compatible wagering mechanisms, bettors typically only stand to lose a small fraction of their wager, and design the Double Clinching Auction (DCA) as a more efficient alternative. Chen et al. (2019) address the efficiency problem by allowing mechanisms to randomize.

Finally, while we are unaware of any other work that draws a direct connection between wagering (or even forecasting more generally) and fair division, the connection between uncertainty and item allocation dates back at least as far as the work by Arrow and Debreu (1954). More recently, Hanson (2003) utilizes scoring rules in the design of prediction markets and Frongillo and Kash (2021) present a convex analysis framework, unifying mechanism design and scoring rule design.

## 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 fair division and wagering mechanisms.

### 2.1. Fair Division Mechanisms

Consider a set of $n \geq 2$ agents $[n] = \{1, \ldots, n\}$ and a set of $m \geq 2$ divisible items $[m] = \{1, \ldots, m\}$. Each agent $i$ has cardinal valuations $\boldsymbol{v}_i = (v_{i,1}, \ldots, v_{i,m})$ such that $v_{i,k} \geq 0$ for each item $k \in [m]$, capturing the agents' marginal rates of substitution across individual items. For example, $v_{i,1} = 0.75, v_{i,2} = 0.25$ means that agent $i$ values one part of item 1 as much as three parts of item 2. As is standard in the cake cutting literature, we assume these valuations are normalized so that $\sum_{k=1}^{m} v_{i,k} = 1$ for all agents $i$. Let $\boldsymbol{a}_i = (a_{i,1}, \ldots, a_{i,m})$ denote a bundle of items allocated to agent $i$, with each component $a_{i,k} \in [0,1]$ specifying the fraction of item $k$ that $i$ receives. Agent $i$'s utility[2] for $\boldsymbol{a}_i$ with valuations $\boldsymbol{v}_i$ is then $u_i(\boldsymbol{a}_i, \boldsymbol{v}_i) = \sum_{k=1}^{m} v_{i,k} a_{i,k} = \boldsymbol{v}_i \cdot \boldsymbol{a}_i$.

An *allocation* $\boldsymbol{a} = (\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n)$ consists of bundles of items $\boldsymbol{a}_i$, one for each agent $i$, such that no item is overallocated, i.e, $\sum_{i=1}^{n} a_{i,k} \leq 1$ for all $k \in [m]$. Observe that we do allow items to be allocated incompletely, resulting in waste. We say that an allocation is *non-wasteful* if $\sum_{i=1}^{n} a_{i,k} = 1$ for all $k \in [m]$. A *fair division mechanism* $\mathcal{A}$ (for "allocation") takes as input a profile of all agents' reports $\boldsymbol{y} = (\boldsymbol{y}_1, \ldots, \boldsymbol{y}_n)$, where $\boldsymbol{y}_i = (y_{i,1}, \ldots, y_{i,m})$ is a vector of reported (normalized) valuations for each agent $i$, and outputs an allocation $\mathcal{A}(\boldsymbol{y})$. We use $\mathcal{A}_i(\boldsymbol{y})$ for the bundle of items allocated to agent $i$ by $\mathcal{A}$ and $\mathcal{A}_{i,k}(\boldsymbol{y})$ for the implied fraction of item $k \in [m]$. We say that a fair division mechanism is non-wasteful if it is guaranteed to output non-wasteful allocations.

One way to evaluate the quality of an allocation $\boldsymbol{a}$ from a mechanism is according to a *social welfare* function that takes into account the utilities of all agents. We will make use of two prominent social welfare functions: *utilitarian welfare* $\mathbf{U}(\boldsymbol{a}, \boldsymbol{y}) = \frac{1}{n} \sum_{i=1}^{n} u_i(\boldsymbol{a}_i, \boldsymbol{y}_i) = \frac{1}{n} \sum_{i=1}^{n} \sum_{k=1}^{m} y_{i,k} a_{i,k}$, defined as the arithmetic mean of the agents' utilities (according to their reported valuations), and *Nash welfare* $\mathbf{N}(\boldsymbol{a}, \boldsymbol{y}) = \left(\prod_{i=1}^{n} u_i(\boldsymbol{a}_i, \boldsymbol{y}_i)\right)^{\frac{1}{n}} = \left(\prod_{i=1}^{n} \sum_{k=1}^{m} y_{i,k} a_{i,k}\right)^{\frac{1}{n}}$, defined as the geometric mean of the agents' utilities (Nash 1950).[3] Observe that Nash welfare can be seen as a tradeoff between utilitarian welfare maximization (since the product rewards high utilities) and egalitarian fairness (in the extreme case, a single agent with zero utility results in zero Nash welfare).

[2] Technically, $u_i$ does not require a subscript since it is fully specified by the bundle and valuation vector that it takes as input, but we add the subscript for clarity.

[3] Note that, although common in the literature, the assumption of normalized valuations is not without loss of generality in the context of utilitarian welfare. More precisely, the relative welfare of two allocations can be reversed by a change in the scale of the agents' valuations. In contrast, this is not the case under Nash welfare. For example, if the sum of an agent's valuations doubled, the Nash welfare of any allocation would also double.

Another way to evaluate mechanisms is via axiomatic properties. Several of these have been proposed in the fair division literature. Anonymity requires that the mechanism's allocation does not depend on the identities of the agents.

DEFINITION 1. A fair division mechanism $\mathcal{A}$ is *anonymous* if for all profiles of reports $\boldsymbol{y}$, all permutations $\sigma$ of $[n]$, all agents $i$, and all items $k \in [m]$, it holds that $\mathcal{A}_{i,k}((\boldsymbol{y}_1, \ldots, \boldsymbol{y}_n)) = \mathcal{A}_{\sigma(i),k}((\boldsymbol{y}_{\sigma^{-1}(1)}, \ldots, \boldsymbol{y}_{\sigma^{-1}(n)}))$.

Proportionality (Steinhaus 1948) requires that each agent receives a $1/n$ share of her total (reported) valuations for the set of all items.

DEFINITION 2. A fair division mechanism $\mathcal{A}$ is *proportional* if, for all profiles of reports $\boldsymbol{y}$ and all agents $i$, it holds that $u_i(\mathcal{A}_i(\boldsymbol{y}), \boldsymbol{y}_i) \geq 1/n$.

Envy-freeness (Foley 1967) requires that no agent ever prefer another agent's allocation to her own.

DEFINITION 3. A fair division mechanism $\mathcal{A}$ is *envy-free* if, for all profiles of reports $\boldsymbol{y}$, and for all agents $i, j$, it holds that $u_i(\mathcal{A}_i(\boldsymbol{y}), \boldsymbol{y}_i) \geq u_i(\mathcal{A}_j(\boldsymbol{y}), \boldsymbol{y}_i)$.

Note that a trivial envy-free and proportional allocation always exists; simply allocate each agent a $1/n$ fraction of each item. Furthermore, note that envy-freeness implies proportionality for non-wasteful mechanisms.

Pareto optimality is a prominent notion of economic efficiency. It requires that a mechanism outputs allocations for which it is impossible to make any agent better off without making another agent worse off.

DEFINITION 4. A fair division mechanism $\mathcal{A}$ is *Pareto optimal* if for all profiles of reports $\boldsymbol{y}$, there does not exist an allocation $\boldsymbol{a}$ such that $u_i(\boldsymbol{a}_i, \boldsymbol{y}_i) \geq u_i(\mathcal{A}_i(\boldsymbol{y}), \boldsymbol{y}_i)$ for all agents $i$, with the inequality being strict for at least one agent.

Finally, we desire mechanisms that do not incentivize agents to misreport their (private) valuations. The definition we use is that of dominant-strategy incentive compatibility, which requires that it is in each agent's best interest to report truthfully, regardless of the reports of others. This is in contrast to the weaker concept of Bayes-Nash incentive compatibility, which only requires truthful reporting to be an equilibrium.

DEFINITION 5. A fair division mechanism $\mathcal{A}$ is *(weakly) incentive compatible* if, for all agents $i$, all valuations $\boldsymbol{v}_i$, and all profiles of reports $\boldsymbol{y}$, it holds that $u_i(\mathcal{A}_i((\boldsymbol{y}_1, \ldots, \boldsymbol{v}_i, \ldots, \boldsymbol{y}_n)), \boldsymbol{v}_i) \geq u_i(\mathcal{A}_i((\boldsymbol{y}_1, \ldots, \boldsymbol{y}_i, \ldots, \boldsymbol{y}_n)), \boldsymbol{v}_i)$. $\mathcal{A}$ is *strictly incentive compatible* if this inequality is strict whenever $\boldsymbol{y}_i \neq \boldsymbol{v}_i$.

Note that, for $n = 2$, it is known that any fair division mechanism satisfying incentive compatibility and Pareto optimality must be a dictatorship (Schummer 1996), in which one agent receives all items that she has positive value for, with the other agent receiving the remaining items. In particular, this implies that no incentive-compatible and Pareto-optimal fair division mechanism can also satisfy anonymity, proportionality, or envy-freeness.

### 2.2. Wagering Mechanisms

The model of wagering mechanisms we study here was first formalized by Lambert et al. (2008), who build on ideas of Kilgour and Gerchak (2004). For intuition, consider a group of friends betting (probabilistically) on the outcome of a sports game, where everyone pays a certain amount of money into a (common) pot and later—depending on the outcome that materializes—receives some money back.

Let $X$ be a random variable that takes values in $[m] = \{1, \ldots, m\}$ with $m \geq 2$ and a set of $n \geq 2$ agents (or bettors) $[n] = \{1, \ldots, n\}$. Let $\Delta_m$ denote the space of probability distributions over the $m$ possible outcomes. Each agent $i$ has a private and subjective belief $\boldsymbol{p}_i = (p_{i,1}, \ldots, p_{i,m}) \in \Delta_m$ with $p_{i,k}$ denoting the probability mass that agent $i$ assigns to event outcome $k \in [m]$. These beliefs are assumed to be immutable, i.e., agent $i$ does not update her belief upon learning any other agent's belief. A *wagering mechanism* $\mathcal{B}$ (for "betting") elicits beliefs from the agents. Formally, each agent $i$ pays a "wager"[4] $w \geq 0$ and reports a belief $\boldsymbol{y}_i = (y_{i,1}, \ldots, y_{i,m}) \in \Delta_m$, and, once the event has materialized, gets paid an amount $\mathcal{B}_i(\boldsymbol{y}, k)$, which depends on the full profile of agent reports $\boldsymbol{y} = (\boldsymbol{y}_1, \ldots, \boldsymbol{y}_n) \in \Delta_m^n$ and the eventually observed outcome $X = k$. Wagering mechanism $\mathcal{B}$ is defined as the set $\{\mathcal{B}_i\}_{i \in [n]}$. For notational simplicity, we normalize the total amount paid by the agents to $\sum_{i=1}^{n} w := 1$, resulting in a uniform wager $w = 1/n$. To be a well-defined wagering mechanism, it must be the case that the payments later (partly) returned to the agents are non-negative, i.e., $\mathcal{B}_i(\boldsymbol{y}, k) \geq 0$ for all $i, \boldsymbol{y}$ and $k$.

Lambert et al. (2008), extending the model of Kilgour and Gerchak (2004), introduce a set of properties desirable for wagering mechanisms: anonymity, individual rationality, budget balance, incentive compatibility, and normality. We present each of them in turn. Anonymity requires that the payments do not depend on the identities of the agents.

Definition 6. A wagering mechanism $\mathcal{B}$ is *anonymous* if, for all profiles of reports $\boldsymbol{y}$, all permutations $\sigma$ of $[n]$, all agents $i$, and all outcomes $k \in [m]$, it holds that $\mathcal{B}_i((\boldsymbol{y}_1, \ldots, \boldsymbol{y}_n), k) = \mathcal{B}_{\sigma(i)}((\boldsymbol{y}_{\sigma^{-1}(1)}, \ldots, \boldsymbol{y}_{\sigma^{-1}(n)}), k)$.

Individual rationality says that, under the assumption that agents are risk neutral, agents do not lose money in expectation and therefore would willingly participate in the mechanism.

Definition 7. A wagering mechanism $\mathcal{B}$ is *individually rational* if, for all agents $i$ and all beliefs $\boldsymbol{p}_i$, there exists a report $\boldsymbol{y}_i$ such that for all reports of the other agents $(\boldsymbol{y}_1, \ldots, \boldsymbol{y}_{i-1}, \boldsymbol{y}_{i+1}, \ldots, \boldsymbol{y}_n)$, it holds that $\mathbf{E}_{X \sim \boldsymbol{p}_i} \mathcal{B}_i(\boldsymbol{y}, X) \geq w = 1/n$.

[4] Lambert et al. (2008) extend the model of Kilgour and Gerchak (2004) to incorporate individual wagers but since the main ideas of our work are already exhibited without this extra degree of freedom, we restrict to the uniform-wager case presented here.

Strict budget balance guarantees that the amount paid to the mechanism is fully returned to the participants. Weak budget balance allows for some money to remain with the mechanism.

DEFINITION 8. A wagering mechanism $\mathcal{B}$ is *weakly budget balanced* if, for all profiles of reports $\boldsymbol{y}$ and all outcomes $k \in [m]$, it holds that $\sum_{i=1}^{n} \mathcal{B}_i(\boldsymbol{y}, k) \leq 1$. Wagering mechanism $\mathcal{B}$ is *strictly budget balanced* if the inequality holds with equality for all $\boldsymbol{y}$ and all $k \in [m]$.

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

DEFINITION 9. A wagering mechanism $\mathcal{B}$ is *(weakly) incentive compatible* if, for all agents $i$, all beliefs $\boldsymbol{p}_i$, and all profiles of reports $\boldsymbol{y}$, $\mathbf{E}_{X\sim \boldsymbol{p}_i} \mathcal{B}_i((\boldsymbol{y}_1, \ldots, \boldsymbol{p}_i, \ldots, \boldsymbol{y}_n), X) \geq \mathbf{E}_{X\sim \boldsymbol{p}_i} \mathcal{B}_i((\boldsymbol{y}_1, \ldots, \boldsymbol{y}_i, \ldots, \boldsymbol{y}_n), X)$. $\mathcal{B}$ is *strictly incentive compatible* if the inequality is strict whenever $\boldsymbol{y}_i \neq \boldsymbol{p}_i$.

Normality gives meaning to the idea that an agent's payment can be interpreted as a measure of her prediction's accuracy relative to the accuracies of the other agents. It says that if the expected payment of an agent $i$ increases (decreases) as a result of $i$ changing her report, then the expected payment of every other agent should not increase (decrease). Similarly, if agent $i$'s expected payment does not change when she changes her report, then the expected payment of all other agents should also stay unchanged.

DEFINITION 10. A wagering mechanism $\mathcal{B}$ is *normal* if, for every probability distribution $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_m) \in \Delta_m$, every agent $i$, all reports $\boldsymbol{y}_i, \boldsymbol{y}'_i$ with $\boldsymbol{y}'_i \neq \boldsymbol{y}_i$, and all profiles of reports of the other agents $(\boldsymbol{y}_1, \ldots, \boldsymbol{y}_{i-1}, \boldsymbol{y}_{i+1}, \ldots, \boldsymbol{y}_n)$, the following two conditions hold:

1. If $\mathbf{E}_{X\sim \boldsymbol{\theta}} \mathcal{B}_i((\boldsymbol{y}_1, \ldots, \boldsymbol{y}'_i, \ldots, \boldsymbol{y}_n), X) > \mathbf{E}_{X\sim \boldsymbol{\theta}} \mathcal{B}_i((\boldsymbol{y}_1, \ldots, \boldsymbol{y}_i, \ldots, \boldsymbol{y}_n), X)$ then $\mathbf{E}_{X\sim \boldsymbol{\theta}} \mathcal{B}_j((\boldsymbol{y}_1, \ldots, \boldsymbol{y}'_i, \ldots, \boldsymbol{y}_n), X) \leq \mathbf{E}_{X\sim \boldsymbol{\theta}} \mathcal{B}_j((\boldsymbol{y}_1, \ldots, \boldsymbol{y}_i, \ldots, \boldsymbol{y}_n), X)$ for all $j \neq i$.
2. If $\mathbf{E}_{X\sim \boldsymbol{\theta}} \mathcal{B}_i((\boldsymbol{y}_1, \ldots, \boldsymbol{y}'_i, \ldots, \boldsymbol{y}_n), X) = \mathbf{E}_{X\sim \boldsymbol{\theta}} \mathcal{B}_i((\boldsymbol{y}_1, \ldots, \boldsymbol{y}_i, \ldots, \boldsymbol{y}_n), X)$ then $\mathbf{E}_{X\sim \boldsymbol{\theta}} \mathcal{B}_j((\boldsymbol{y}_1, \ldots, \boldsymbol{y}'_i, \ldots, \boldsymbol{y}_n), X) = \mathbf{E}_{X\sim \boldsymbol{\theta}} \mathcal{B}_j((\boldsymbol{y}_1, \ldots, \boldsymbol{y}_i, \ldots, \boldsymbol{y}_n), X)$ for all $j \neq i$.

The first case covers the situations of agent $i$ increasing her expected score by changing her report from $\boldsymbol{y}_i$ to $\boldsymbol{y}'_i$, and, symmetrically, of agent $i$ decreasing her expected score by changing from $\boldsymbol{y}'_i$ to $\boldsymbol{y}_i$.

## 3. The Correspondence

In this section, we show that there is a one-to-one correspondence between fair division mechanisms and weakly budget-balanced wagering mechanisms. To build intuition, imagine yourself in a wagering setting and consider the situation before the outcome materializes but with all reports already reported. Even by then, the wagering mechanism already describes, for each possible outcome, who shall receive how much of the full \$1. That is, the wagering mechanism defines conditional payments for each outcome. Now imagine that the $m$ possible outcomes in the wagering setting correspond to $m$ items in a fair division setting, and consider a single outcome/item $k$. What is an agent's valuation for this item? In the wagering setting, the agent believes outcome $k$ to occur with probability $p_{i,k}$ and hence that outcome's contribution to her expected payment is $p_{i,k}$ times the fraction of \$1 she receives in that outcome, i.e., $p_{i,k} \cdot \mathcal{B}_i(\boldsymbol{y}, k)$. In the fair division setting, the agent values the entire item at $v_{i,k}$ and, given that she only receives fraction $\mathcal{A}_{i,k}(\boldsymbol{y})$, her valuation for her fraction of the item is $v_{i,k} \cdot \mathcal{A}_{i,k}(\boldsymbol{y})$.

By comparing the agent's utility in both settings, it is natural to define a fair division mechanism $\mathcal{A}$, corresponding to the given wagering mechanism, which allocates to each agent $i$ a $\mathcal{B}_i(\boldsymbol{y}, k)$ fraction of item $k$ given reports $\boldsymbol{y}$, i.e., $\mathcal{A}_{i,k}(\boldsymbol{y}) = \mathcal{B}_i(\boldsymbol{y}, k)$. The total (expected) utility that the agent receives in the wagering setting is $\mathbf{E}_{X \sim \boldsymbol{p}_i} \mathcal{B}_i(\boldsymbol{y}, X) = \sum_{k=1}^m p_{i,k} \cdot \mathcal{B}_i(\boldsymbol{y}, k)$, while the utility that she receives in the fair division setting is $u_i(\mathcal{A}_i(\boldsymbol{y}), \boldsymbol{v}_i) = \boldsymbol{v}_i \cdot \mathcal{A}_i(\boldsymbol{y}) = \sum_{k=1}^m v_{i,k} \cdot \mathcal{A}_{i,k}(\boldsymbol{y}) = \sum_{k=1}^m v_{i,k} \cdot \mathcal{B}_i(\boldsymbol{y}, k)$. Since we have assumed linear additive valuations in the fair division setting, this correspondence between mechanisms preserves (expected) utility when beliefs $\boldsymbol{p}_i$ are mapped to valuations $\boldsymbol{v}_i$.

**Corresponding Mechanisms**

The following definition formalizes this correspondence.

Definition 11 (Corresponding mechanisms). Fair division mechanism $\mathcal{A}$ and wagering mechanism $\mathcal{B}$ are *corresponding mechanisms* if, for all profiles of reports $\boldsymbol{y}$, all $i \in [n]$, and all $k \in [m]$,

$$\mathcal{A}_{i,k}(\boldsymbol{y}) = \mathcal{B}_i(\boldsymbol{y}, k).$$

Definition 11 gives a one-to-one correspondence between fair division mechanisms and weakly budget-balanced wagering mechanisms. We illustrate both directions separately.

If $\mathcal{A}$ is a fair division mechanism, then, by definition of a fair division mechanism, its corresponding wagering mechanism always returns non-negative payments. That is, it satisfies $\mathcal{B}_i(\boldsymbol{y}, k) \geq 0$ for all profiles of reports $\boldsymbol{y}$, all $i \in [n]$, and all $k \in [m]$. Furthermore, because a fair division mechanism never overallocates items, the corresponding wagering mechanism is weakly budget balanced. That is,

$$\sum_{i=1}^n \mathcal{B}_i(\boldsymbol{y}, k) = \sum_{i=1}^n \mathcal{A}_{i,k}(\boldsymbol{y}) \leq 1.$$

Similarly, if $\mathcal{B}$ is a weakly budget-balanced wagering mechanism, then the corresponding fair division mechanism satisfies $\mathcal{A}_{i,k}(\boldsymbol{y}) \geq 0$ for all profiles of reports $\boldsymbol{y}$, all $i \in [n]$, and all $k \in [m]$. Moreover, for all profiles of reports $\boldsymbol{y}$ and all $k \in [m]$,

$$\sum_{i=1}^n \mathcal{A}_{i,k}(\boldsymbol{y}) = \sum_{i=1}^n \mathcal{B}_i(\boldsymbol{y}, k) \leq 1.$$

### Corresponding Properties

In the remainder of the section, we show that the correspondence from Definition 11 maps several desirable properties from the fair division to the wagering setting, and vice versa. In particular, incentive compatibility is preserved, and, subject to incentive compatibility, proportionality maps to individual rationality and envy-freeness is implied by normality.

We first show an equivalence between incentive compatibility in both settings. That is, a misreport that produces a more preferred allocation of items also produces a higher expected payment in the wagering setting. All proofs are in the appendix.

THEOREM 1. *A fair division 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 proportionality and individual rationality. Though typically not described this way, the equivalence highlights that proportionality can be viewed as providing incentives for agents to participate over an outside option of being allocated $1/n$ of each item.

THEOREM 2. *An incentive-compatible fair division mechanism $\mathcal{A}$ is proportional if and only if the corresponding incentive-compatible wagering mechanism $\mathcal{B}$ is individually rational.*

The restriction to incentive-compatible mechanisms is needed because the definition of individual rationality allows the belief/valuation to be different from the report, whereas proportionality guarantees that the inequality in the definition holds when the report equals the belief/valuation.[5] Incentive compatibility nullifies this difference as it guarantees that the report equals the belief/valuation.

Finally, we show that, subject to incentive compatibility and anonymity,[6] normality implies envy-freeness.

THEOREM 3. *If an anonymous, incentive-compatible wagering mechanism is normal then the corresponding anonymous, incentive-compatible fair division mechanism is envy-free.*

Note that, as part of the proof of Theorem 3, we show that anonymity is preserved by the correspondence. Further note that the converse to Theorem 3 does not hold; in particular, we show in the e-companion (Section EC.1) that Strong Demand Matching is anonymous, (weakly) incentive compatible, and envy-free, but not normal.

[5] Without restricting to incentive compatibility, proportionality implies individual rationality but the reverse implication does not hold. For example, the following mechanism is individually rational but not proportional: each agent receives 1/n of every item if she reports $(1, 0, ..., 0)$ and otherwise receives nothing.

[6] In fact, the theorem continues to hold when we exchange anonymity for the strictly weaker property of "equal treatment of equals," which states that two agents with the same report should receive the same payment (wagering) / bundle (fair division).

## 4. Competitive Scoring Rules: A Family of Envy-Free, Proportional, and Incentive-Compatible Fair Division Mechanisms

A primary focus of the fair division literature is to design mechanisms that are "fair," with proportionality and envy-freeness typically understood to be the gold standard. To our knowledge, the only known incentive-compatible mechanism to achieve allocations that are fair in this sense is the uniform allocation mechanism, which allocates each agent a $1/n$ fraction of every item. Of course, this uniform allocation is hardly satisfactory as it ignores the agents' preferences and hence potential benefits from trade. Instead, one would want a mechanism to facilitate beneficial trade whenever agent preferences differ, and thus to Pareto dominate the uniform allocation with respect to the agents' utilities.

In this section, we introduce *Competitive Scoring Rules* (CSRs, Kilgour and Gerchak 2004), a class of wagering mechanisms that—when imported into the fair division setting—turn out to be the first mechanisms to jointly satisfy the aforementioned properties.[7] We begin by introducing proper scoring rules, which form the basis of Competitive Scoring Rules.

**Proper Scoring Rules.** Proper scoring rules (Brier 1950, Good 1952, Gneiting and Raftery 2007) are scoring functions, which are used in two ways: to evaluate the accuracy of probabilistic forecasts and to incentivize the truthful reporting of privately held, probabilistic beliefs. Given the focus of this work, we introduce them taking the incentive perspective. Consider a future event, such as a presidential election, that will take one of $m$ possible outcomes (i.e., one of the $m$ presidential candidates will eventually be elected). Furthermore, consider an agent with a private, subjective belief $\boldsymbol{p}_i \in \Delta_m$ about the probability of the $m$ outcomes, and a principal seeking to truthfully elicit $\boldsymbol{p}_i$ from the agent. A proper scoring rule is a scoring function that ensures the agent maximizes her expected score by reporting her private belief truthfully to the principal. The temporal order is as follows: first, the agent reports a forecast $\boldsymbol{y}_i \in \Delta_m$, which may or may not coincide with her private belief $\boldsymbol{p}_i$. Second, one of the $m$ outcomes materializes, and, third, the proper scoring rule assigns the agent a score that depends on the agent's reported forecast and the materialized outcome.

Definition 12. A *scoring rule* $R$ is a function that maps a report $\boldsymbol{y}_i \in \Delta_m$ and an outcome $k \in [m]$ to a score $R(\boldsymbol{y}_i, k) \in \mathbb{R} \cup \{-\infty\}$. Scoring rule $R$ is *(weakly) proper* if, for all $\boldsymbol{y}_i, \boldsymbol{p}_i \in \Delta_m$,

$$\mathop{\mathbf{E}}_{X \sim \boldsymbol{p}_i} R(\boldsymbol{p}_i, X) \geq \mathop{\mathbf{E}}_{X \sim \boldsymbol{p}_i} R(\boldsymbol{y}_i, X).$$

Scoring rule $R$ is *strictly proper* if the inequality is strict for all $\boldsymbol{y}_i \neq \boldsymbol{p}_i$. Scoring rule $R$ is *bounded* if there exist $\underline{R}, \overline{R} \in \mathbb{R}$ such that $R(\boldsymbol{y_i}, k) \in [\underline{R}, \overline{R}]$ for all $\boldsymbol{y}_i \in \Delta_m, k \in [m]$.

[7] As we will see in Section 5, four mechanisms from the literature can be cast as CSRs and hence also satisfy these properties, but they are only defined for the restricted case of $n = 2$ agents. Furthermore, in Section 6, we will see that two other mechanisms satisfy these properties but again only for restricted cases of $n = 2$ ($\text{PA}_{\max}$) and $m = 2$ (DCA). In addition to these restrictions on $n$ or $m$, both of these latter mechanisms are wasteful. Moreover, and analogous to CSRs, DCA is only known to be a fair division mechanism due to the correspondence in the first place.

It will be helpful to have notation $G^R(\boldsymbol{p}_i) := \mathbf{E}_{X\sim \boldsymbol{p}_i} R(\boldsymbol{p}_i, X)$ for the expected score function given truthful reporting and $L^R(\boldsymbol{p}_i, \boldsymbol{y}_i) := \mathbf{E}_{X\sim \boldsymbol{p}_i}[R(\boldsymbol{p}_i, X) - R(\boldsymbol{y}_i, X)]$ for the expected loss of reporting $\boldsymbol{y}_i \in \Delta_m$ instead of true belief $\boldsymbol{p}_i \in \Delta_m$ under proper scoring $R$.

Positive-affine transformations of proper scoring rules preserve (strict) properness and there exist infinitely many proper scoring rules since any (strictly) convex function $f$ yields a (strictly) proper scoring rule $R$ with $G^R = f$ (Gneiting and Raftery, 2007; Theorem 1). A widely used bounded scoring rule is the *quadratic scoring rule* (Brier 1950), which we will regularly refer to throughout the paper and give here in its standard, normalized form, yielding scores between 0 and 1.

Proposition 1. *(Brier 1950) The* quadratic scoring rule $R_q(\boldsymbol{y}_i, k) = \boldsymbol{y}_i(k) - 0.5\sum_{\ell=1}^{m} \boldsymbol{y}_i(\ell)^2 + 0.5$ *is strictly proper.*

**Competitive Scoring Rules.** Competitive Scoring Rules, first introduced by Kilgour and Gerchak (2004) and further analyzed by Lambert et al. (2008), can be thought of as multi-agent versions of proper scoring rules. Indeed, they are parameterized by the choice of a proper scoring rule, yielding an infinite family of mechanisms. In this section, we analyze this family of mechanisms in the context of fair division, exploiting the correspondence from Section 3.

Given a proper scoring rule bounded by the interval $[0,1]$, a Competitive Scoring Rule can be thought of as paying each agent their wager of $1/n$, plus an adjustment equal to $1/n$ times the difference between their own score and the average score of all other agents.

Definition 13. A *Competitive Scoring Rule*[8] pays every agent $i$

$$\mathcal{C}^R_{i,k}(\boldsymbol{y}) = \mathcal{C}^R_i(\boldsymbol{y}, k) = \frac{1}{n} + \frac{1}{n}\left(R(\boldsymbol{y}_i, k) - \frac{1}{n-1}\sum_{j\neq i} R(\boldsymbol{y}_j, k)\right), \tag{1}$$

where $R$ is a proper scoring rule bounded by $[0,1]$.

Every Competitive Scoring Rule has a corresponding fair division mechanism according to Definition 11. (In Definition 13, we first use the fair division and then the wagering notation.) To see this, first observe that the boundedness condition on $R$ immediately implies that the Competitive Scoring Rule makes non-negative payments. Second, summing over all agents, it is easy to see that every Competitive Scoring Rule is strictly budget balanced.

Competitive Scoring Rules are known to satisfy a number of desirable properties in the wagering setting.[9]

[8] Kilgour and Gerchak (2004) originally defined Competitive Scoring Rules to pay each agent the difference between their own score and the average score of the other agents, that is, only the part between the large brackets. The formulation that we use is a modification by Lambert et al. (2008), who adapted Competitive Scoring Rules to guarantee non-negative payments.

[9] In their Theorem 1, Lambert et al. (2008) only consider CSRs that are defined using strictly proper scoring rules. The extension to weakly proper scoring rules is an easy adaptation of their proof.

Theorem 4 (**Kilgour and Gerchak, 2004; Lambert et al., 2008, Theorem 1**). *Competitive Scoring Rules are anonymous, incentive compatible, individually rational, and normal. Additionally, if a Competitive Scoring Rule is defined using a strictly proper scoring rule, then it is strictly incentive compatible.*

As a corollary of Theorem 4 combined with Theorems 1, 2, and 3 from the correspondence, we obtain the following.

Corollary 1. *Fair division mechanisms corresponding to Competitive Scoring Rules are anonymous, incentive compatible, proportional, and envy-free. Additionally, if a Competitive Scoring Rule is defined using a strictly proper scoring rule, then the corresponding fair division mechanism is strictly incentive compatible.*

As mentioned earlier, the only fair division mechanism known to be incentive compatible, proportional, and envy-free is the uniform allocation mechanism.[10] The following proposition shows that CSRs with strictly proper scoring rules provide strictly higher utility to every agent compared to the uniform allocation as long as some agents disagree about the relative valuation of the items.[11] In the context of cake cutting, Mossel and Tamuz (2010) describe mechanisms with this property as "super fair."

Proposition 2. *Let $R$ be a strictly proper scoring rule and let $\boldsymbol{y}$ be a profile of reports for which there exist $j, \ell$ with $\boldsymbol{y}_j \neq \boldsymbol{y}_\ell$. Then the utility of every agent $i \in [n]$ under CSR $\mathcal{C}^R$ is*

$$u_i(\mathcal{C}_i^R(\boldsymbol{y}), \boldsymbol{y}_i) > 1/n.$$

This implies that CSRs are the first incentive-compatible, proportional, and envy-free fair division mechanisms that Pareto dominate the uniform allocation mechanism. Moreover, when defined using a strictly proper scoring rule, Competitive Scoring Rules are, to our knowledge, the first *strictly* incentive-compatible fair division mechanisms.

## 5. Characterizing Incentive-Compatible Fair Division Mechanisms for $n = 2$

Building on the correspondence, we have identified CSRs as a family of incentive-compatible fair division mechanisms with a unique set of desirable properties among known mechanisms. Further capitalizing on the correspondence, we now exploit a known result in wagering to show that, in fact,

[10] In fact, the uniform allocation mechanism is a member of the Competitive Scoring Rule family, namely the CSR with the constant scoring rule $R(\boldsymbol{y}_i, k) = 1$ for all $\boldsymbol{y}_i \in \Delta_m, k \in [m]$.

[11] In the proof of Proposition 2, we show a more general version that applies whenever there exists an agent $j$ who obtains a strictly lower expected score than agent $i$, according to $i$'s report $\boldsymbol{y}_i$. When $R$ is strictly proper and at least two agents make different reports, such a $j$ is guaranteed to exist.

CSRs comprise the complete set of non-wasteful, incentive-compatible fair division mechanisms for $n = 2$, subject to mild technical conditions (Section 5.1). In Section 5.2, as an immediate consequence of this characterization, we close an open question regarding the best possible approximation to the optimal utilitarian welfare that can be achieved by a non-wasteful, incentive-compatible fair division mechanism for $n = 2$ agents. Moreover, this characterization also reduces the design space of non-wasteful, incentive-compatible fair division mechanisms for $n = 2$ to that of proper scoring rules. We first explore this in Section 5.3, where we show that CSRs encompass several known mechanisms from the literature as special cases by identifying their corresponding scoring rules. We then discuss this further in Section 5.4, where we exploit known structural properties of proper scoring rules to give novel closed-form expressions for utilitarian and Nash welfare in terms of the associated loss functions, demonstrating that insights from the scoring rule literature can be used to inform the particular fair division mechanism that should be used in practice.

### 5.1. The Characterization

Han et al. (2011) characterized the class of incentive-compatible fair division mechanisms for the setting with $n = 2$ agents and $m = 2$ items. They showed that any incentive-compatible mechanism that is anonymous and item symmetric (i.e., if the valuations of two items are swapped by every agent, then the allocations of these two items are also swapped) must be swap-dictatorial. A *swap-dictatorial* mechanism is defined by a (potentially infinite) set of bundles[12] $D \subseteq [0,1]^m$. The mechanism randomly selects one agent as the dictator, who is then allocated their most preferred bundle in $D$ (ties broken arbitrarily), with the other agent receiving the remaining items. The output of the mechanism is then the expectation over both possible choices of dictator. As the first step towards a characterization that applies to any $m$, we show that the class of swap-dictatorial mechanisms is exactly the class of CSRs.

Theorem 5. *For $n = 2$ agents, a fair division mechanism $\mathcal{A}$ is swap-dictatorial if and only if $\mathcal{A}$ is a Competitive Scoring Rule $\mathcal{C}^R$ for some proper scoring rule $R$, i.e., if and only if $\mathcal{A}_{i,k}(\boldsymbol{y}) = \frac{1}{2} + \frac{1}{2}\big(R(\boldsymbol{y}_i, k) - R(\boldsymbol{y}_{3-i}, k)\big)$ for all $i \in \{1,2\}$ and all $k \in [m]$.*

Knowing that swap-dictatorial mechanisms and CSRs are equivalent, we can now leverage a characterization by Lambert et al. (2008), which was originally developed in the wagering setting. We show that, subject to mild technical conditions, the class of CSRs in fact coincides with the class of non-wasteful, incentive-compatible mechanisms. In order to state the result, we need to define a technical condition. We present it in the fair division setting for concreteness but, by the

[12] Guo and Conitzer (2010) and Han et al. (2011) define swap-dictatorial mechanisms using individual sets of bundles $D_i$. For ease of exposition, since we refer to swap-dictatorial mechanisms only in the context of anonymous mechanisms, we have only one set of bundles for all agents.

correspondence (Section 3), it can equally be presented in the wagering setting. Smoothness is largely a mathematical convenience but it does imply that the allocations are continuous in every agent's report, which means that allocations do not change drastically for small changes in reports.

DEFINITION 14. A fair division mechanism $\mathcal{A}$ is *smooth* if, for all $i \in [n]$, $\mathcal{A}_i$ is twice continuously differentiable with respect to every report $\boldsymbol{y}_j$, $j \in [n]$. A proper scoring rule $R$ is smooth if it is twice continuously differentiable with respect to the report $\boldsymbol{y}_i$ for all $k \in [m]$.

LEMMA 1. *For $n = 2$ agents, if a fair division mechanism $\mathcal{A}$ is incentive compatible, anonymous, smooth, and non-wasteful, then it is a Competitive Scoring Rule $\mathcal{C}^R$ with some smooth proper scoring rule $R$.*

Lemma 1 immediately follows from applying the correspondence to Lemma 4 of Lambert et al. (2008) since normality is implied by $n = 2$ for strictly budget-balanced wagering mechanisms. Note that Lambert et al. (2008) consider *strictly* incentive-compatible mechanisms and show that they are defined by *strictly* proper scoring rules, but a simple adaptation of their proof allows the incentive compatibility condition to be relaxed to weak incentive compatibility by allowing the more general class of weakly proper scoring rules instead. Theorem 6 combines Lemma 1 with Theorem 5 and strengthens Han et al. (2011)'s characterization result along two dimensions: first, their characterization applies only to $m = 2$ items and, second, they require item symmetry.

THEOREM 6. *For $n = 2$ agents, if a mechanism $\mathcal{A}$ is anonymous, smooth, and non-wasteful, then $\mathcal{A}$ is incentive compatible if and only if $\mathcal{A}$ is swap-dictatorial.*

### 5.2. Optimal Welfare Approximation

Theorem 6 has immediate implications for a line of work that has focused on using incentive-compatible mechanisms to approximate the optimal welfare for fair division settings with $n = 2$ agents. The goal is to maximize the worst case ratio between the utilitarian welfare implemented by the mechanism and the optimal utilitarian welfare (i.e., the average of the agents' utilities in the allocation in which every item is allocated to the agent with highest valuation for it). Guo and Conitzer (2010) introduced the problem and considered the class of *increasing-price mechanisms*. They showed that no increasing-price mechanism can obtain an approximation factor better than 0.5 as the number of items goes to infinity. However, they left open the problem of whether better mechanisms can be found by focusing on more general families of incentive-compatible mechanisms. Han et al. (2011, Theorem 4) generalized this result to all swap-dictatorial mechanisms and posed as an open question whether the impossibility generalizes beyond swap-dictatorial mechanisms. In later work, Cole et al. (2013a) and Cheung (2016) designed incentive-compatible fair division mechanisms that achieve approximation factors better than 0.5, but these mechanisms allow some items (or parts of items) to remain unallocated, whereas both Guo and Conitzer and Han et al. restricted to non-wasteful mechanisms. It has remained an open question whether any non-wasteful fair division mechanism exists that achieves an approximation factor better than 0.5.

By combining Theorem 6 with the result of Han et al. (2011, Theorem 4) that no swap-dictatorial mechanism can achieve an approximation to optimal utilitarian welfare better than 0.5 for all $m$, we arrive at the following corollary, closing the open question posed by Han et al. (2011) subject to smoothness. Note that anonymity does not appear as a condition in Corollary 2 because it is known that if a non-anonymous mechanism achieves an approximation factor of $\alpha$, there exists an anonymous mechanism that achieves an approximation factor of at least $\alpha$ (Guo and Conitzer 2010).

Corollary 2. *For $n=2$ agents, if a mechanism $\mathcal{A}$ is smooth, weakly incentive compatible and non-wasteful, then $\mathcal{A}$ does not achieve an approximation factor better than 0.5 for all $m$.*

Observe that Corollary 2 implies that any smooth and incentive-compatible mechanism achieving an approximation factor better than 0.5—as those of Cole et al. (2013a) and Cheung (2016) do—must be wasteful.

### 5.3. Revisiting Incentive-Compatible Fair Division Mechanisms from the Literature

Given Lemma 1, we now know that all incentive-compatible, anonymous, smooth, and non-wasteful mechanisms for $n=2$ agents do take the form of a CSR with a particular choice of proper scoring rule $R$. In this section, we revisit four mechanisms that were suggested in the literature and cast them into the CSR framework, spelling out the concrete proper scoring rules that they correspond to. Interpreting these known mechanisms in this way often allows for a more compact representation. Going forward, we believe this will also facilitate the design and simplify the analysis of new mechanisms in this space.

**Revisiting Guo and Conitzer.** Guo and Conitzer (2010) initiated the study of incentive-compatible utilitarian welfare approximation. In Sections 4 and 5, they introduce a family of linear increasing-price (LIP) mechanisms, which are parameterized by a parameter $a$. Interestingly, this family of mechanisms is tightly connected to the well-known logarithmic scoring rule (Good 1952) defined by $R_l(\boldsymbol{y}_i, k) = \ln(y_{i,k})$. In particular, each member of the LIP family can be expressed as a CSR. For $m=2$, the one-dimensional report space allows for a compact expression[13] of the corresponding scoring rule, which, surprisingly, turns out to be the (weakly) proper truncated-log scoring rule $R_{t-l}$:

[13] For $m>2$, the corresponding CSR, as for $m=2$, also makes a case distinction depending on $\boldsymbol{y}_i$. If $\frac{\min_k\{y_{i,k}\}}{\max_k\{y_{i,k}\}} \geq \frac{1}{e^a}$, i.e., if no outcome is assigned too little probability mass, then $R_{t-l}(\boldsymbol{y}_i, k) = \frac{1}{a}\ln(y_{i,k}) + \frac{1}{a}(\ln(e^a + m - 1))$. If this condition does not hold, $R_{t-l}$ does not have a concise description, but is defined in such a way that it remains proper.

$$R_{t-l}(\boldsymbol{y}_i, k) = \begin{cases} 0 \text{ if } y_{i,k} \in \left[0, \frac{1}{1+e^a}\right] \\ \frac{1}{a}\ln(y_{i,k}) + \frac{\ln(1+e^a)}{a} \text{ if } y_{i,k} \in \left(\frac{1}{1+e^a}, \frac{e^a}{1+e^a}\right) \\ 1 \text{ if } y_{i,k} \in \left[\frac{e^a}{1+e^a}, 1\right] \end{cases}$$

Observe that, if the probability report $y_{i,k}$ is near 1 or 0—with the concrete range around 1 and 0 determined by $a$—then the report is scored with the highest or lowest possible score (depending on the outcome). For reports in the "middle range" the score is computed using the logarithmic scoring rule, transformed to yield scores in between 0 and 1. In practice, this procedure of truncation followed by transformation to yield scores in between 0 and 1 is often applied to the logarithmic scoring rule in an effort to preserve the rule's beneficial information-theoretic properties for some range of reports while also guaranteeing bounded, non-negative scores.

Proposition 3. *For $m = 2$, the Linear Increasing Price Mechanism (LIP) with parameter $a$ of Guo and Conitzer (2010) is $\mathcal{C}^{R_{t-l}}$.*

**Revisiting Han et al.** Han et al. (2011, Definition 5) suggest the Sphere Mechanism, which, for $n = 2$, achieves an approximation factor that is strictly higher than 0.5 if the valuation vectors are bounded in the sense that $y_{i,k} \leq C/m$ for some constant $C$ and all $i \in \{1,2\}$, $k \in [m]$. Recall from Corollary 2 in Section 5.2 that any approximation better than 0.5 is impossible for smooth mechanisms and general valuations. As it turns out, Han et al. implicitly reinvent the well-known spherical scoring rule (Roby 1965, Jose 2009), scaled by $\frac{\sqrt{m}}{C}$:

$$R_s(\boldsymbol{y_i}, k) = \frac{\sqrt{m}}{C} \cdot \frac{y_{i,k}}{||\boldsymbol{y}_i||}$$

Proposition 4. *The Sphere Mechanism of Han et al. (2011, Definition 5) is $\mathcal{C}^{R_s}$.*

Note that for the remainder of the paper, we set $\frac{\sqrt{m}}{C} := 1$.

**Revisiting Cheung.** Cheung (2016, Section 4) designs a mechanism for $n = 2$ and $m = 2$, achieving a $\frac{5}{6}$ utilitarian welfare approximation, currently the best-known for this setting.[14] His mechanism is the CSR with the following proper scoring rule, defined for $k \in \{1,2\}$:[15]

[14] Cheung provides a computer proof that this bound cannot be improved to $\frac{5}{6} + \epsilon$ with $\epsilon < 10^9$, so that $\frac{5}{6}$ is likely optimal. Note that this bound applies only to the case of $m = 2$, so it does not contradict the bound of 0.5 for general $m$.

[15] As a side note, applying the correspondence to this fair division mechanism has an interesting implication for the corresponding *wagering* mechanism $\mathcal{C}^{R^{\text{Cheung}}}$. In particular, Theorem 4 of Cheung (2016) and Proposition 5 immediately yield a formal guarantee on the sum of subjective expected utilities that the mechanism enables.

$$R_{\text{Cheung}}(\boldsymbol{y}_i, k) = \begin{cases} 0 \text{ if } & y_{i,k} \in [0, \frac{1}{5}] \\ \frac{1}{3}\left(-\frac{1}{y_{i,k}} - \ln(y_{i,k}) + 5 - \ln(5)\right) \text{ if } & y_{i,k} \in (\frac{1}{5}, \frac{1}{2}] \\ 1 - \frac{1}{3}\ln\left(5y_{i,3-k}\right) \text{ if } & y_{i,k} \in (\frac{1}{2}, \frac{4}{5}] \\ 1 \text{ if } & y_{i,k} \in (\frac{4}{5}, 1] \end{cases}$$

PROPOSITION 5. *The mechanism of Cheung (2016, Section 4) is $\mathcal{C}^{R_{\text{Cheung}}}$.*

**Revisiting Aziz and Ye.** Aziz and Ye (2014) present the Constrained Serial Dictatorship (CSD) mechanism for cake cutting. Since cake cutting is a generalization of the fair division setting we consider in this paper, CSD can be applied to our setting in a straightforward manner. To define CSD, consider fixing an order of the agents. Going through this order, each agent receives the $m/n$ items (if $m/n$ is not an integer, agents receive only a fraction of some items) that they value the most out of the remaining items, with ties broken in favor of lower-indexed items.[16] For example, if $m = 3$ and $n = 2$, then the first agent receives their most-preferred item and half of their second-most-preferred item while the second agent receives the remaining items. The output of CSD is the expected allocation over all agent orderings.

For $n = 2$, CSD corresponds to the CSR with the proper scoring rule that rewards an agent $i$ with a score of 1 if the realized outcome is among the $m/2$ most likely outcomes according to $\boldsymbol{y}_i$. Let $\overline{S}$ denote the set of $\lfloor \frac{m}{2} \rfloor$ outcomes that have the most probability mass according to $\boldsymbol{y}_i$, with ties broken in favor of lower-indexed outcomes. Similarly, let $\underline{S}$ denote the set of $\lfloor \frac{m}{2} \rfloor$ outcomes that have the least probability mass according to $\boldsymbol{y}_i$, with ties broken in favor of higher-indexed outcomes. In case $m$ is odd, let $S = [m] \setminus (\underline{S} \cup \overline{S})$ denote the remaining outcome.

$$R_{\text{CSD}}(\boldsymbol{y}_i, k) = \begin{cases} 1 \text{ if } & k \in \overline{S} \\ \frac{1}{2} \text{ if } & k \in S \\ 0 \text{ if } & k \in \underline{S} \end{cases}$$

PROPOSITION 6. *For $n = 2$, the Constrained Serial Dictatorship (CSD) mechanism of Aziz and Ye (2014) is $\mathcal{C}^{R_{\text{CSD}}}$.*

### 5.4. Closed Form Expression for Social Welfare

In Section 5.1, Lemma 1 we have seen that for $n = 2$, all anonymous, incentive-compatible, and non-wasteful mechanisms take the form of a Competitive Scoring Rule (subject to smoothness). We have also seen that they are parameterized by a choice of scoring rule $R$. But which $R$ should we choose? In Section 5.3 we saw four possible choices, but in principle we can choose from infinitely many (because every convex function gives rise to a proper scoring rule). In this section we take the perspective of utilitarian and Nash welfare and provide closed-form expressions for them as a

[16] The tiebreaking order is arbitrary but we align it with the index for simplicity.

function of $R$. In doing so, we gain a deeper understanding of which rules perform well for which reports and why. As it turns out, there is a tight connection between these welfare measures and $R$'s expected score function. Interestingly, visually inspecting the plot of the expected score functions already allows a judgment as to which rule will perform best for a given pair of reports.

THEOREM 7. *Let $n=2$ and consider Competitive Scoring Rule $\mathcal{C}^R$, defined using scoring rule $R$ with associated expected loss function $L^R$. Then, the utility of agent $i$ is*

$$u_i(\mathcal{C}_i^R(\boldsymbol{y}), \boldsymbol{y}_i) = \frac{1}{2} + \frac{1}{2} L^R(\boldsymbol{y}_i, \boldsymbol{y}_{3-i}). \tag{2}$$

*Hence, Competitive Scoring Rule $\mathcal{C}^R$ achieves utilitarian welfare*

$$\mathbf{U}(\mathcal{C}^R(\boldsymbol{y}), \boldsymbol{y}) = \frac{1}{2} + \frac{1}{4}\left(L^R(\boldsymbol{y}_1, \boldsymbol{y}_2) + L^R(\boldsymbol{y}_2, \boldsymbol{y}_1)\right) \tag{3}$$

*and Nash welfare*

$$\mathbf{N}(\mathcal{C}^R(\boldsymbol{y}), \boldsymbol{y}) = \frac{1}{2} \cdot \sqrt{(1 + L^R(\boldsymbol{y}_1, \boldsymbol{y}_2))\,(1 + L^R(\boldsymbol{y}_2, \boldsymbol{y}_1))}. \tag{4}$$

Observe that utilitarian welfare is simply 0.5 plus half the average of the two losses. The expected loss functions of standard scoring rules are known (Gneiting and Raftery 2007, p. 363), so plugging them into Equations 3 and 4 of Theorem 7 immediately gives rise to the corresponding closed-form solutions for utilitarian and Nash welfare. For example, the expected loss of the quadratic scoring rule is the Euclidean distance $L^{R_q}(\boldsymbol{p}_i, \boldsymbol{y}_i) = \sum_{k=1}^m (\boldsymbol{p}_i(k) - \boldsymbol{y}_i(k))^2$ and so the CSR with the quadratic scoring rule yields utilitarian welfare

$$\begin{aligned}
\mathbf{U}(\mathcal{C}^{R_q}(\boldsymbol{y}), \boldsymbol{y}) =& \frac{1}{2} + \frac{1}{4}\left(L^{R_q}(\boldsymbol{y}_1, \boldsymbol{y}_2) + L^{R_q}(\boldsymbol{y}_2, \boldsymbol{y}_1)\right) \\
=& \frac{1}{2} + \frac{1}{4}\left(\sum_{k=1}^m \big(\boldsymbol{y}_1(k) - \boldsymbol{y}_2(k)\big)^2 + \sum_{k=1}^m \big(\boldsymbol{y}_2(k) - \boldsymbol{y}_1(k)\big)^2\right) \\
=& \frac{1}{2} + \frac{1}{2}\sum_{k=1}^m \big(\boldsymbol{y}_1(k) - \boldsymbol{y}_2(k)\big)^2 \\
=& \frac{1}{2} + \frac{1}{2} L^{R_q}(\boldsymbol{y}_1, \boldsymbol{y}_2).
\end{aligned}$$

The CSR with the quadratic scoring rule is actually an interesting rule in its own right because it is the unique incentive-compatible fair division mechanism that is *equitable* (Brams and Taylor 1996), subject to the technical conditions of Lemma 1. That is, it guarantees equal utility to both agents. To see this, note that $L^{R_q}$ is symmetric in its two arguments and that $R_q$ is the only proper scoring rule with this property (Savage 1971).

In order to provide a more visual interpretation of Theorem 7, we restrict attention to the case of $m=2$, where a belief $\boldsymbol{p}_i = (p_{i,1}, p_{i,2})$ can be represented by a single number. In particular, we

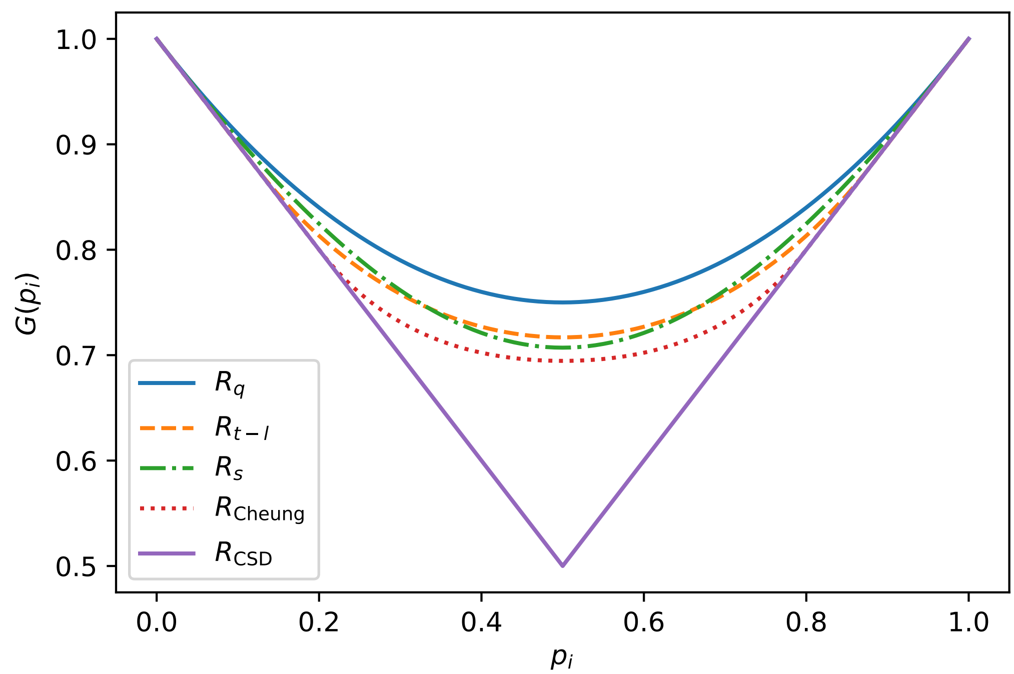


**Figure 1** **The expected score functions $G^R(p_i)$ (see p. 12) for $m = 2$ and different choices of $R$. Note that the order in which the functions are presented in the legend matches their order at $p_i = 0.5$.**

write $p_i$ to represent the probability that agent $i$ places on outcome 2, i.e., $p_i := p_{i,2}$. (Analogously, we write $y_i$ for the corresponding report about outcome 2.) In Figure 1, we plot the expected score functions $G^R(p_i)$ associated with each of the four scoring rules described in Section 5.3 as well as that of the quadratic scoring rule.

Using Theorem 7, we can write the utilitarian welfare in terms of the expected score function $G^R(p_i)$. Consider Competitive Scoring Rule $\mathcal{C}^R$, defined using scoring rule $R$ with associated expected score function $G^R$. The utilitarian welfare that $\mathcal{C}^R$ achieves on reports $y_1, y_2$ is

$$\mathbf{U}(\mathcal{C}^R(\boldsymbol{y}), \boldsymbol{y}) = \frac{1}{2} + \frac{1}{4}(y_1 - y_2)\big(dG^R(y_1) - dG^R(y_2)\big), \tag{5}$$

where $dG^R(y_i)$ is the derivative[17] of $G^R$ at $y_i$. To see this, note that $L^R(y_i, y_{3-i})$ can be expressed as $L^R(y_i, y_{3-i}) = G^R(y_i) - G^R(y_{3-i}) - (y_i - y_{3-i})(dG^R(y_i))$, as is well known (Savage 1971, Gneiting and Raftery 2007), so that Equation 5 follows from Equation 3 using simple algebra.[18]

Let us conclude by reflecting on what Theorem 7 and its implied Equation 5 tell us about designing incentive-compatible mechanisms for two agents and two items. Suppose that $y_1$ and $y_2$

[17] If $G^R(y_i)$ is not differentiable, $dG^R(y_i)$ is a subgradient.

[18] One can perform a similar exercise to write the Nash welfare in terms of the expected score function $G^R$, but the resulting expression is not as informative.

are fixed. Seeking to maximize utilitarian welfare, the mechanism designer should then choose a scoring rule with corresponding expected score function $G^R$ such that the difference $|dG^R(y_1) - dG^R(y_2)|$ is maximized. If the mechanism designer has distributional knowledge about reports $y_1$ and $y_2$, then the optimal scoring rule is the one that maximizes the value of Equation 5 in expectation. Informally speaking, the utilitarian welfare is higher when using a scoring rule whose $G^R$ has high curvature (high rate of change in gradient) in the region in which $y_1$ and $y_2$ are expected to lie. For example, if $y_1$ and $y_2$ are guaranteed to lie on opposite sides of 0.5, then CSD is the best choice since the corresponding expected score function maximizes $|dG^R(y_1) - dG^R(y_2)|$. However, if $y_1$ and $y_2$ are expected to both be greater (or both be less) than 0.5, then CSD is a bad choice since its expected score function is linear on both sides of 0.5. From this it follows that the utilitarian welfare will be exactly 1 because the expected losses of reporting $y_1$ instead of $y_2$, and vice versa, are both 0 (the agents are indifferent between reporting $y_1$ and $y_2$). In contrast, if $y_1$ and $y_2$ are expected to both be close to 0 or both be close to 1, then setting $R$ to be the quadratic score yields the highest utilitarian welfare since it has the highest curvature around the endpoints. We explore this further in the following section using simulations.

## 6. Comparison of Mechanisms

The correspondence from Section 3 implies that any wagering mechanism can be viewed as a fair division mechanism and vice versa. Table 1 provides, to our knowledge, the complete set of known incentive-compatible mechanisms from both literatures, including a summary of their respective properties. We provide more details of these mechanisms in the e-companion (Section EC.1). For clarity, these properties are presented in the fair division setting, and they are either (1) known results in fair division, (2) following from known results in wagering and the correspondence developed in Section 3, or (3) proven separately (in Section EC.1). The properties proven separately are marked with a cross in the table. In the remainder of this section, we explore, via simulation, the performance of the mechanisms from Table 1, evaluating them according to both utilitarian and Nash welfare.

We begin, in Section 6.1, by considering the base setting with $n = 2$ agents and $m = 2$ items, which is a commonly considered special case (Guo and Conitzer 2010, Han et al. 2011, Cheung 2016). As seen in Section 5.1, CSRs comprise the complete set of non-wasteful mechanisms in this setting, subject to mild technical conditions. We include five representative members of the CSR family parameterized by their respective scoring rules: the four identified in Section 5.3 and the commonly used quadratic scoring rule (Proposition 1). Then, in Section 6.2, we consider larger numbers of agents and items.

| | Incentive Compatible | Pareto Optimal | Proportional | Envy-Free | Non-Wasteful |
|---|---|---|---|---|---|
| CEEI | No | Yes | Yes | Yes | Yes |
| CSR | Strict* | No | Yes | Yes | Yes |
| NAWM | Strict* | No | Yes | No† | No |
| CSD** | Weak | No | Yes | No | Yes |
| DCA*** | Weak | No | Yes | Yes† | No |
| PA | Weak | No | No | Yes | No |
| SDM | Weak | No | No | Yes | No |
| $PA_{max}$**** | Weak | No | Yes† | Yes† | No† |

**Table 1** **Comparison of properties satisfied by Competitive Equilibrium from Equal Incomes (CEEI, Foley 1967, Varian 1974), Competitive Scoring Rules (CSR, Kilgour and Gerchak 2004, Lambert et al. 2008), No-Arbitrage Wagering Mechanisms (NAWM, Chen et al. 2014), Constrained Serial Dictatorship (CSD, Aziz and Ye 2014), the Double Clinching Auction (DCA, Freeman et al. 2017), Partial Allocation (PA, Cole et al. 2013b), Strong Demand Matching (SDM, Cole et al. 2013b), and $PA_{max}$ (Cole et al. 2013a). *CSRs and NAWMs are strictly incentive-compatible when used with a strictly proper scoring rule. **Computing the CSD allocation is known to be #P-complete (Aziz et al. 2013), which makes it infeasible to compute for larger $n$; an approximate, sampled version is not guaranteed to output proportional solutions. ***DCA is only defined for $m=2$ items and $n \geq 4$ agents. It relies on a continuous ascending auction, of which we implement an approximate version that is not guaranteed to inherit all of the original mechanism's properties. ****$PA_{max}$ is only defined for $n=2$ agents. The same holds for a slightly extended version due to Cheung (2016), which loses proportionality for a small improvement of the approximation factor (it is unknown if this extension preserves envy-freeness). †These properties are proven in Section EC.1 of the e-companion.**

### 6.1. Two Agents, Two Items

To generate reports in this base setting, we draw from three sets of beta distributions. The first contains beta distributions with $\alpha = \beta$, generating reports centered at $(0.5, 0.5)$, with the uniform distribution corresponding to the special case of $\alpha = \beta = 1$. Increasing $\alpha$ and $\beta$ leads to higher concentration around the mean $(0.5, 0.5)$ and hence more similar reports. We report the value of $\alpha$ as the concentration, which we increase in increments of 0.1. The beta distributions in the second set are no longer centered at $(0.5, 0.5)$. Instead, we generate "biased" reports with mean $(0.25, 0.75)$ by using $\alpha = 3\beta$. Distributions in the third set are centered at $(0.05, 0.95)$ by setting $\alpha = 19\beta$. Analogous to the first set, we again vary the concentration of the reports around the mean by increasing $\alpha$, now subject to $\alpha = 3\beta$ and $\alpha = 19\beta$, respectively.

Figures 2 and 3 display the results, averaged over 1000 samples per parameter choice. For both social welfare measures, the values are normalized so that (1) the allocation that maximizes the respective welfare measure achieves a welfare of 100 and (2) the uniform allocation achieves a welfare of zero. Thus, the welfare we report is the average fraction of the maximum possible improvement over the uniform allocation that the mechanisms achieve. Note that the uniform allocation mechanism's robust incentive and fairness guarantees make it preferable to any mechanism that

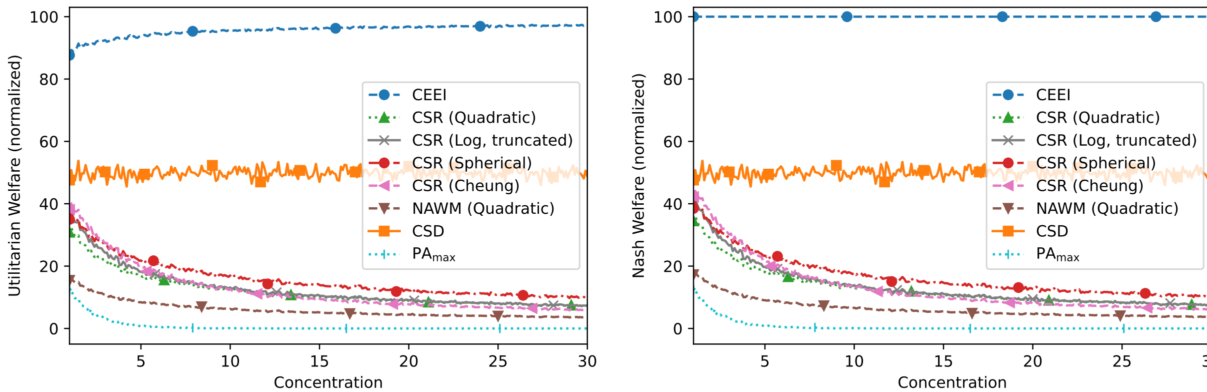


**Figure 2** **Utilitarian welfare (left side) and Nash welfare (right side) of different mechanisms with $n = 2$ agents and $m = 2$ items. Reports are centered around $(0.5, 0.5)$ and the x axis varies the level of concentration around that center.**

does not outperform it in welfare, so that the mechanisms of practical interest obtain normalized welfare between 0 and 100. The allocation maximizing utilitarian welfare allocates each item to the agent who values it the highest. For Nash welfare, Competitive Equilibrium from Equal Incomes (CEEI, Foley 1967, Varian 1974) always achieves a value of 100 because it maximizes that welfare measure (Arrow and Intriligator 1982, Volume 2, Chapter 14). The incentive-compatible mechanisms that we consider achieve welfare values less than 100 since they additionally incorporate incentive compatibility constraints.[19]

The mechanisms included in Figures 2 and 3 are those that are defined for the setting with $n = 2$ agents and $m = 2$ items while also achieving normalized welfare values of at least 0 for some parameter choice. This excludes the Double Clinching Auction (DCA, Freeman et al. 2017), because it is only defined for $n \geq 4$, as well as Partial Allocation (PA, Cole et al. 2013b), Strong Demand Matching (SDM, Cole et al. 2013b), and the extended version of $\text{PA}_{\text{max}}$ due to Cheung (2016), all of which had welfare values below 0 for all parameter choices. For the family of No-Arbitrage Wagering Mechanisms (NAWM, Chen et al. 2014), we include the representative with the quadratic scoring rule. The right side of Figure 3 excludes CEEI and $\text{PA}_{\text{max}}$ (Cole et al. 2013a) for presentational reasons.

We now discuss some interesting behavior that can be seen in the plots. First, observe in the uniform setting presented in Figure 2 that the relative performance of the mechanisms is essentially identical across the two welfare measures. This is also the case for the two biased settings, and the discussion thus focuses on utilitarian welfare. Figure 3 shows the results for biased reports

[19] As an example, consider $n = 2$ agents and $m = 2$ items with $\boldsymbol{v}_1 = (1, 0)$ and $\boldsymbol{v}_2 = (0.6, 0.4)$. The non-normalized utilitarian welfare values of the uniform allocation mechanism (both agents receive half of both items) and the welfare-maximizing mechanism (agent 1 receives good 1, agent 2 receives good 2) are 0.5 and 0.7, respectively. If a given mechanism achieves non-normalized utilitarian welfare of 0.6, it would achieve a normalized utilitarian welfare of 50.

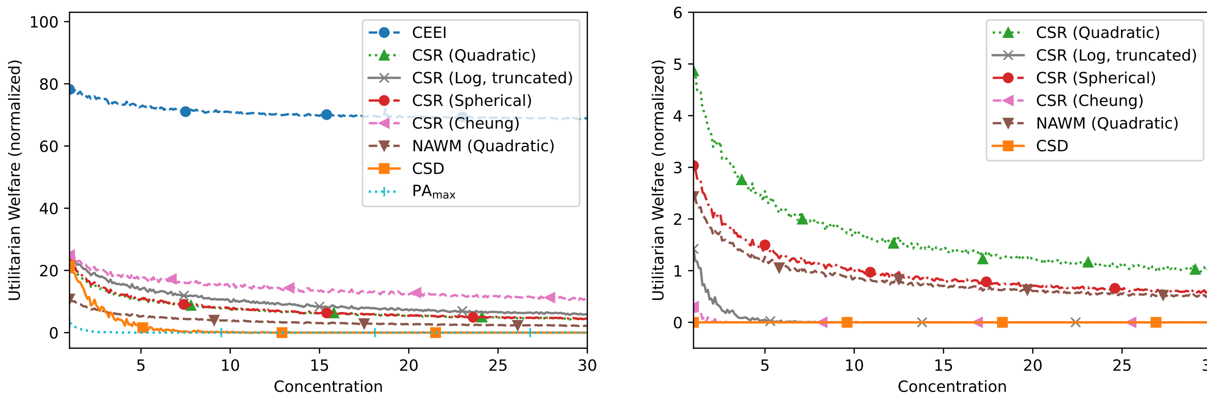


**Figure 3** **Utilitarian welfare of different mechanisms with $n = 2$ agents and $m = 2$ items for two different biases. Reports are centered around $(0.25, 0.75)$ (left side) and $(0.05, 0.95)$ (right side), and the x axis varies the level of concentration around those centers. On the right side, CEEI (welfare values between 52 and 55) and PA$_{max}$ (welfare value 0 everywhere) are left out for presentational reasons and the y axis is scaled differently than in the other plots we present.**

under utilitarian welfare, with the corresponding plots for Nash welfare to be found in the e-companion (Section EC.2). The mechanisms' similar performance across the two welfare measures is not surprising. For intuition, note that on allocations where some agents receive very low utility while others receive very high utility, Nash and utilitarian welfare would differ. However, since all mechanisms from Table 1 are proportional or envy-free (and, more generally, designed to be "fair"), one would not expect big differences between utilitarian and Nash welfare.

Second, observe the performance of CSD in the three different settings. As we discussed in Section 5.4, CSD is a good choice when the two agents are likely to disagree as to which of the two items is more valuable, i.e., when their reports lie on opposite sides of $(0.5, 0.5)$. For unbiased distributions (Figure 2), that kind of disagreement happens frequently and we see that CSD outperforms the other incentive-compatible mechanisms. Additionally, the performance is unaffected by the concentration since changing the concentration does not affect the probability that the agents disagree. This insight can in fact be used to compute CSD's normalized welfare in this setting: when the agents agree, CSD coincides with the uniform allocation. When the agents disagree, CSD coincides with the allocation that maximizes utilitarian welfare (i.e., giving every item to the agent who values it higher). The probability of each of these two cases is 0.5, so the normalized utilitarian welfare is $0.5 \cdot 0 + 0.5 \cdot 100 = 50$, which matches CSD's behavior in Figure 2. The same normalized welfare of 50 holds with respect to Nash welfare as it is still the case that CSD obtains the maximal possible value with probability 0.5 and that it yields the uniform allocation otherwise.

For the distributions with mean $(0.25, 0.75)$, CSD continues to achieve some gains from trade for low concentrations because the distribution of reports has high enough variance for agents to sometimes value the first item higher than the second item (left side of Figure 3). As the concentration increases, however, the reports converge to $(0.25, 0.75)$, so that the agents almost never disagree on which item is more valuable. This yields normalized utilitarian welfare of zero as CSD then coincides with the uniform allocation. For the most biased setting with mean $(0.05, 0.95)$, even for low concentrations, agents almost never disagree about which of the two items is more valuable, so that CSD does not achieve any gains from trade (right side of Figure 3).

Let us now consider the performance of the other mechanisms. In Figure 2, we see that the Competitive Scoring Rule with the spherical rule, $\mathcal{C}^{R_s}$, outperforms the other non-CSD Competitive Scoring Rules at high concentrations. This is to be expected as per the discussion in Section 5.4: if we fix reports $y_1$ and $y_2$, we know from Equation 5 that the utilitarian welfare of $\mathcal{C}^R$ is proportional to the difference $|dG^R(y_1) - dG^R(y_2)|$. If both $y_1$ and $y_2$ are close to 0.5, as they are for high concentrations, then this difference is proportional to $d^2G^R(0.5)$, the second derivative of $G^R$ at 0.5. For the spherical scoring rule we have $d^2G^{R_s}(0.5) = 2.828$, for the quadratic scoring rule and the truncated log scoring rule (with $a = 2$), we have $d^2G^{R_q}(0.5) = d^2G^{R_{t-l}}(0.5) = 2$, and for $R_{\text{Cheung}}$ we have $d^2G^{R_{\text{Cheung}}}(0.5) = 1.33$. At high concentrations, it can be checked that the relative performance of the mechanisms is consistent with these second derivatives. Note that the welfare values in the figures are not exactly proportional to the numbers given here, both because we normalize the welfare values in the plots and because the concentration values are not high enough to observe fully limiting behavior.

The relative performance of the different Competitive Scoring Rules in Figure 3 can be explained analogously, now examining the second derivatives of the corresponding expected score functions at 0.75 and 0.95, respectively. For a bias of $(0.25, 0.75)$, $\mathcal{C}^{R_{\text{Cheung}}}$ achieves the highest welfare as can be predicted from visually inspecting its high curvature at $p = 0.75$ in Figure 1. As before, we can also verify the relative performances at $(0.25, 0.75)$ mathematically by twice differentiating the expected score functions at 0.75. For a bias of $(0.05, 0.95)$, $\mathcal{C}^{R_q}$ achieves the highest welfare, with $\mathcal{C}^{R_{t-l}}$, $\mathcal{C}^{R_{\text{Cheung}}}$, and $\mathcal{C}^{R_{\text{CSD}}}$ achieving no gains from trade at higher concentrations because the expected score functions of their respective scoring rules are linear around 0.95.[20]

Finally, observe that the $\text{PA}_{\max}$ mechanism of Cole et al. (2013a) performs poorly in all three settings. Arguably, this is not a surprising finding given that it is optimized for worst-case performance and for all values of $m$. In particular, it often reverts to the uniform allocation, sacrificing any gains from trade.

[20] We observe that normalized welfare tends to decrease as agents' reports become more similar (either due to higher concentration or increased bias). Note that while similar reports limit the potential gains from trade, it is not obvious that *normalized* welfare should decrease since this limitation applies to all mechanisms, whether incentive compatible or not. At a high level, the decrease in normalized welfare is observed because incentive-compatible mechanisms are forced to be less discriminative than the utilitarian welfare maximizing mechanism, which simply allocates the items to whoever claims to value them more highly.

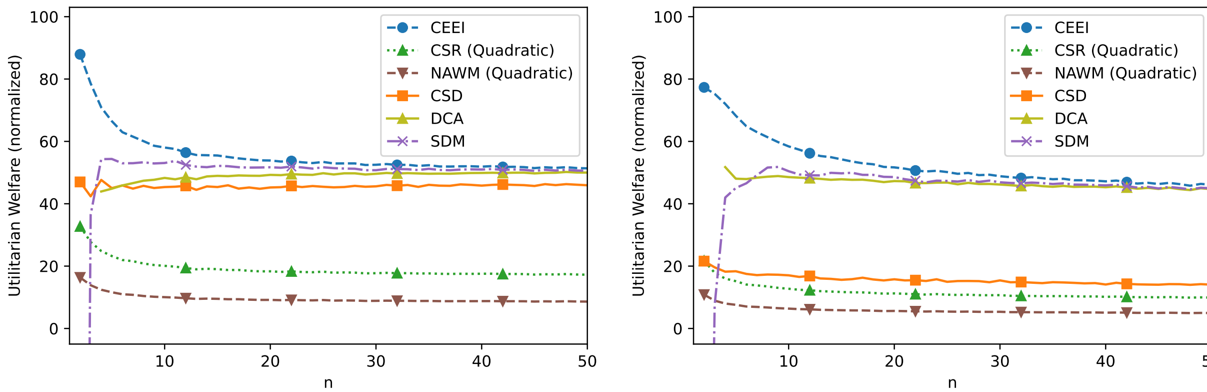


**Figure 4** **Utilitarian welfare of different mechanisms with $m = 2$ items for uniform (left side) and biased (right side) reports. The number of agents $n$ varies from 2 to 50 (x axis).**

### 6.2. Many Agents, Many Items

We now turn to settings with larger values of $n$ and $m$. As was the case in Section 6.1, the observed results between utilitarian and Nash welfare are similar, so we focus here on the results for utilitarian welfare. The corresponding plots for Nash welfare can be found in the e-companion (Section EC.2). In Figure 4, we present results for $m = 2$ items and vary $n$ from 2 to 50. As before, we generate both uniform (left side) and biased reports (right side). On the left side of Figure 5, we present results for $n = 10$ agents and vary the number of items between $m = 2$ and $m = 40$ under uniform sampling. On the right side of Figure 5, the number of agents is also $n = 10$ but the number of items is fixed to $m = 20$ and we use biased sampling. For uniform sampling, we draw reports uniformly at random from the simplex $\Delta_m$ using a Dirichlet distribution with parameters $(\alpha_1, \ldots, \alpha_m) = (1, \ldots, 1)$. For biased sampling, generalizing the Beta distribution from Section 6.1, we use a Dirichlet distribution with parameters $(\alpha_1, \alpha_2, \alpha_3, \ldots, \alpha_m) = (\alpha_1, 3\alpha_1, 5\alpha_1, \ldots, (2m-1)\alpha_1)$. For example, for $m = 5$, this yields reports with mean $(0.04, 0.12, 0.2, 0.28, 0.36)$.

The mechanisms we evaluate in Figures 4 and 5 are, with the exception of DCA, those from Table 1 that are defined for all values of $n \geq 2$ and $m \geq 2$. PA only appears on the left side of Figure 5 because it had welfare values below 0 for all parameter choices in the other settings. For the same reason, SDM does not appear on the right side of Figure 5. For visual clarity, we only include the Competitive Scoring Rule with the quadratic scoring rule. The performance of CSRs with other scoring rules is similar. Also note that while CSD is a Competitive Scoring Rule when restricted to $n = 2$, it is not a member of that family for higher values of $n$. Because computing the CSD allocation is #P-complete (Aziz et al. 2013), we use an approximate version of CSD for all $n \geq 8$, where we take the average allocation over 1000 randomly sampled permutations.

Considering the setting with many agents and few items as depicted in Figure 4, several themes stand out. SDM and DCA outperform all other incentive-compatible mechanisms, which can be

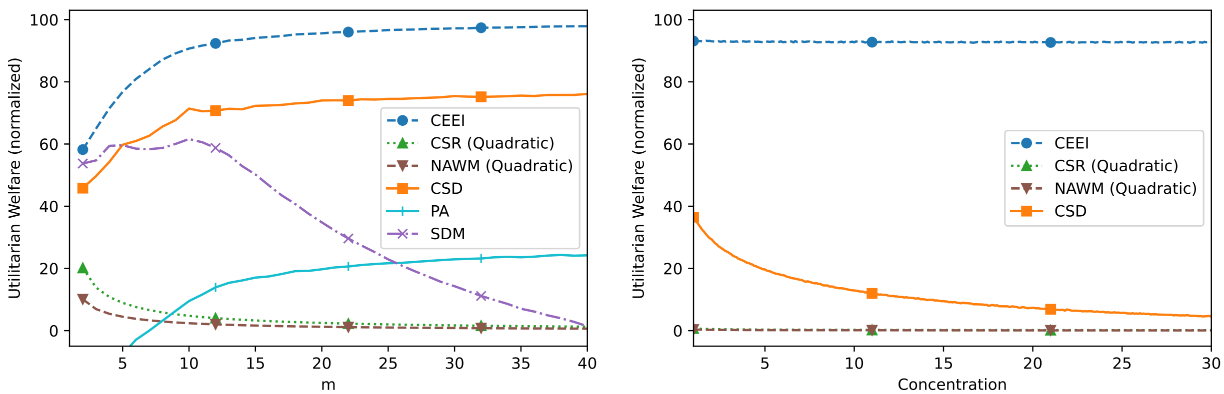


**Figure 5** **Utilitarian welfare of different mechanisms with $n = 10$ agents for uniform reports varying $m$ from 2 to 40 (left side) and biased report with $m = 20$ varying concentration (right side).**

explained by the fact that both of these mechanisms are designed to approximate CEEI. Note that, in contrast to SDM, DCA satisfies proportionality. This is crucial in settings where stakeholders can only accept mechanisms that guarantee every participant to be at least as well off as uniformly distributing the resources at hand. On the other hand, DCA's restriction to $m = 2$ items limits its applicability, so that SDM significantly outperforms all other incentive-compatible mechanisms for settings with many agents and $m \geq 3$ items. CSD performs well under uniform sampling, but its performance degrades significantly under biased reports, which is consistent with the $n = m = 2$ results from Section 6.1. CSRs do badly for higher $n$, which is unsurprising because it is known from wagering that CSR payments never exceed twice the agent's wager, which, in the fair division setting, translates to saying that no agent can receive more than a $2/n$ fraction of any item.[21]

Turning to the setting with many items relative to agents as depicted in Figure 5, we see that CSD achieves the highest welfare for almost all values of $m$ under uniform sampling (left side) and all values of concentration in the biased-report setting (right side). No other incentive-compatible mechanism achieves any noticeable improvement over the uniform allocation for biased reports. The robust performance of CSD is particularly notable considering the results observed in Section 6.1, where it is substantially outperformed by other mechanisms for biased reports and high concentrations. The intuition for this result is that, in our setup, as $m$ increases, agents more frequently disagree with respect to the items' ordinal valuations, which is known to increase welfare from CSD. Conversely, we see the performance of SDM declining as $m$ increases, which is to be expected because SDM is designed to perform well when $n$ is large relative to $m$. For uniform reports, PA does achieve some gains from trade as $m$ increases. For intuition, note that PA is known to perform well when no agent imposes too much externality on the other agents. This is

[21] This follows from Equation 1 and the fact that $R$ is bounded in $[0, 1]$.

the case when $m$ is large relative to $n$ and reports are uniform because one would not expect much agreement between the agents' valuations. Finally, note that even for uniform reports, CSRs and NAWMs both achieve very little gain from trade.

## 7. Conclusion

In this paper, we show that the fair division and wagering settings are, in fact, the same. Given that the corresponding research communities have previously studied these settings in isolation, it is no surprise that their equivalence has immediate implications. In particular, applying the correspondence to Competitive Scoring Rules, arguably the most prominent class of wagering mechanisms, immediately leads to the first incentive-compatible fair division mechanism that is both fair (proportional and envy-free) and responsive to agent preferences. Moreover, leveraging a result from the wagering literature, we can characterize, subject to mild technical conditions, all non-wasteful, incentive-compatible fair division mechanisms for two agents as Competitive Scoring Rules. Because they are parameterized by proper scoring rules, this reduces the mechanism design problem in this setting to one of proper scoring rule selection.

In comparing the existing mechanisms from the fair division literature with the new mechanisms gained through the correspondence, several rules of thumb for application become apparent. For the base setting of two agents and two items, both axiomatic properties and our simulation results suggest that Competitive Scoring Rules are superior to other, wasteful mechanisms. An appropriate choice of parameter (i.e., scoring rule) is determined by the area in which the agents' reports are expected to lie. When the number of agents is large relative to the number of items, our simulation results suggest that Strong Demand Matching should be used. (For two items, the Double Clinching Auction performs similarly well in simulation but additionally satisfies proportionality.) Conversely, when the number of items is large relative to the number of agents, Constrained Serial Dictatorship proves surprisingly robust.

Going forward, 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.

## Acknowledgments

The authors thank Vasilis Gkatzelis, Mirko Kremer, Jochen Schlapp, and the anonymous reviewers for helpful feedback. This paper is a significantly extended version of Freeman et al. (2019). This full paper contains the following contributions that do not appear in the preliminary version. First, it establishes a formal connection between normality, a property considered in the wagering literature, and envy-freeness, a notion of fairness in the fair division literature (Theorem 3). Second, Competitive Scoring Rules, a class of mechanisms originally developed in the wagering literature are shown to characterize all incentive-compatible and non-wasteful fair division

mechanisms for two agents, subject to two mild technical conditions (Lemma 1). Leveraging this characterization, Section 5 casts known fair division mechanisms into this scoring rule framework and presents improved bounds on the approximation to optimal utilitarian welfare that any fair division mechanism can achieve (Corollary 2). Finally, Section 6 contains new and significantly expanded simulations comparing all known incentive-compatible mechanisms.

## References

Georgios Amanatidis, Georgios Birmpas, George Christodoulou, and Evangelos Markakis. Truthful allocation mechanisms without payments: Characterization and implications on fairness. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 545–562, Cambridge, MA, USA, 2017.

Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, and Alexandros A. Voudouris. Fair Division of Indivisible Goods: A Survey. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI'22)*, pages 5385–5393, 2022. Survey Track.

Kenneth J Arrow and Gerard Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, pages 265–290, 1954.

Kenneth J. Arrow and Michael D. Intriligator, editors. *Handbook of Mathematical Economics*. North-Holland, 1982.

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.

Xiaohui Bei, Guangda Huzhang, and Warut Suksompong. Truthful fair division without free disposal. *Social Choice and Welfare*, 55(3):523–545, 2020.

Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. The price of fairness. *Operations Research*, 59 (1):17–31, 2011.

Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. On the efficiency-fairness trade-off. *Management Science*, 58(12):2234–2250, 2012.

Steven J. Brams and Alan D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.

Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

Simina Brânzei and Peter Bro Miltersen. A dictatorship theorem for cake cutting. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 482–488, 2015.

Simina Brânzei, Vasilis Gkatzelis, and Ruta Mehta. Nash social welfare approximation for strategic agents. *Operations Research*, 70(1):402–415, 2022.

Glenn W. Brier. Verification of Forecasts Expressed in Terms of Probability. *Monthly Weather Review*, 78 (1):1–3, 1950.

Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. On low-envy truthful allocations. In *Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT)*, pages 111–119, 2009.

Yiling Chen, John K. Lai, David C. Parkes, and Ariel D. Procaccia. Truth, justice, and cake cutting. *Games and Economic Behavior*, 77:284–297, 2013.

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)*, pages 377–394, 2014.

Yiling Chen, Yang Liu, and Juntao Wang. Randomized wagering mechanisms. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1845–1852, 2019.

Yun Kuen Cheung. Better strategyproof mechanisms without payments or prior - an analytic approach. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 194–200, 2016.

Yann Chevaleyre, Paul E. Dunne, Ulle Endriss, Jérôme Lang, Michel Lemaître, Nicolas Maudet, Julian Padget, Steve Phelps, Juan A. Rodríguez-Aguilar, and Paulo Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.

Richard Cole, Vasilis Gkatzelis, and Gagan Goel. Positive results for mechanism design without money. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1165–1166, 2013a.

Richard Cole, Vasilis Gkatzelis, and Gagan Goel. Mechanism design for fair division. In *Proceedings of the 14th ACM Conference on Economics and Computation (EC)*, pages 251–268, 2013b.

Richard Cole, Vasilis Gkatzelis, and Gagan Goel. Mechanism design for fair division. Working Paper, 2022.

Xiaotie Deng, Qi Qi, and Amin Saberi. Algorithmic solutions for envy-free cake cutting. *Operations Research*, 60(6):1461–1476, 2012.

Nikhil R Devanur, Bach Q Ha, and Jason D Hartline. Prior-free auctions for budgeted agents. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, pages 287–304, 2013.

Shahar Dobzinski, Ron Lavi, and Noam Nisan. Multi-unit auctions with budget limits. In *Proceedings of the 49th Annual Symposium on Foundations of Computer Science (FOCS)*, 2008.

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, 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.

Rupert Freeman, David M. Pennock, and Jennifer Wortman Vaughan. An equivalence between wagering and fair-division mechanisms. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1957–1964, 2019.

Rafael M Frongillo and Ian A Kash. General truthfulness characterizations via convex analysis. *Games and Economic Behavior*, 130:636–662, 2021.

Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.

I.J. Good. Rational decisions. *Journal of the Royal Statistical Society, Series B (Methodological)*, 14:107–114, 1952.

Mingyu Guo and Vincent Conitzer. Strategy-proof allocation of multiple items between two agents without payments or priors. In *Proceedings of the 9th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 881–888, 2010.

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.

Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.

Victor Richmond Jose. A Characterization for the Spherical Scoring Rule. *Theory and Decision*, 66(3): 263–281, 2009.

D. Marc Kilgour and Yigal Gerchak. Elicitation of probabilities using competitive scoring rules. *Decision Analysis*, 1(2):108–113, 2004.

Nicolas Lambert, John Langford, Jennifer Wortman, Yiling Chen, Daniel 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 M. Reeves, Yoav Shoham, and David M. Pennock. An axiomatic characterization of wagering mechanisms. *Journal of Economic Theory*, 156:389–416, 2015.

Avishay Maya and Noam Nisan. Incentive compatible two player cake cutting. In *Proceedings of the 8th Conference on Web and Internet Economics (WINE)*, pages 170–183, 2012.

Vijay Menon and Kate Larson. Deterministic, strategyproof, and fair cake cutting. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 352–358, 2017.

Elchanan Mossel and Omer Tamuz. Truthful fair division. In *Proceedings of the 3rd International Symposium on Algorithmic Game Theory (SAGT)*, pages 288–299, 2010.

John Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.

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.

Thornton B. Roby. Belief states and the uses of evidence. *Behavioral Science*, 10(3):255–270, 1965.

Leonard J. Savage. Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association*, 66:783–801, 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.

Biaoshuai Tao. On existence of truthful fair cake cutting mechanisms. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*, pages 404–434, 2022.

Hal R. Varian. Equity, envy and efficiency. *Journal of Economic Theory*, 9:63–91, 1974.

# Appendix

## A. Proof of Theorem 1

THEOREM 1. *A fair division mechanism $\mathcal{A}$ is (weakly, strictly) incentive compatible if and only if the corresponding wagering mechanism $\mathcal{B}$ is (weakly, strictly) incentive compatible.*

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

$$\begin{aligned}
u_i(\mathcal{A}_i((\boldsymbol{y}_1,\dots,\boldsymbol{v}_i,\dots,\boldsymbol{y}_n)),\boldsymbol{v}_i) &= \boldsymbol{v}_i \cdot \mathcal{A}_i((\boldsymbol{y}_1,\dots,\boldsymbol{v}_i,\dots,\boldsymbol{y}_n)) \\
&= \sum_{k=1}^{m} v_{i,k} \cdot \mathcal{A}_{i,k}((\boldsymbol{y}_1,\dots,\boldsymbol{v}_i,\dots,\boldsymbol{y}_n)) \\
&= \sum_{k=1}^{m} v_{i,k} \cdot \mathcal{B}_i((\boldsymbol{y}_1,\dots,\boldsymbol{v}_i,\dots,\boldsymbol{y}_n),k) \\
&= \mathop{\mathbf{E}}_{X\sim \boldsymbol{v}_i} \mathcal{B}_i((\boldsymbol{y}_1,\dots,\boldsymbol{v}_i,\dots,\boldsymbol{y_n}),X) \\
&\geq \mathop{\mathbf{E}}_{X\sim \boldsymbol{v}_i} \mathcal{B}_i((\boldsymbol{y}_1,\dots,\boldsymbol{y}_i,\dots,\boldsymbol{y_n}),X) \\
&= \sum_{k=1}^{m} v_{i,k} \cdot \mathcal{B}_i((\boldsymbol{y}_1,\dots,\boldsymbol{y}_i,\dots,\boldsymbol{y}_n),k) \\
&= \sum_{k=1}^{m} v_{i,k} \cdot \mathcal{A}_{i,k}((\boldsymbol{y}_1,\dots,\boldsymbol{y}_i,\dots,\boldsymbol{y}_n)) \\
&= \boldsymbol{v}_i \cdot \mathcal{A}_i((\boldsymbol{y}_1,\dots,\boldsymbol{y}_i,\dots,\boldsymbol{y}_n)) \\
&= u_i(\mathcal{A}_i((\boldsymbol{y}_1,\dots,\boldsymbol{y}_i,\dots,\boldsymbol{y}_n)),\boldsymbol{v}_i).
\end{aligned}$$

The third and sixth 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 $\boldsymbol{y_i} \neq \boldsymbol{p_i}$. We have:

$$\begin{aligned}
\mathop{\mathbf{E}}_{X\sim \boldsymbol{p}_i} \mathcal{B}_i((\boldsymbol{y}_1,\dots,\boldsymbol{p}_i,\dots,\boldsymbol{y_n}),X) &= \sum_{k=1}^{m} p_{i,k}\cdot \mathcal{B}_i((\boldsymbol{y}_1,\dots,\boldsymbol{p}_i,\dots,\boldsymbol{y}_n),k) \\
&= \sum_{k=1}^{m} p_{i,k}\cdot \mathcal{A}_{i,k}((\boldsymbol{y}_1,\dots,\boldsymbol{p}_i,\dots,\boldsymbol{y}_n)) \\
&= \boldsymbol{p}_i\cdot \mathcal{A}_i((\boldsymbol{y}_1,\dots,\boldsymbol{p}_i,\dots,\boldsymbol{y}_n)) \\
&\geq \boldsymbol{p}_i\cdot \mathcal{A}_i((\boldsymbol{y}_1,\dots,\boldsymbol{y}_i,\dots,\boldsymbol{y}_n)) \\
&= \sum_{k=1}^{m} p_{i,k}\cdot \mathcal{A}_{i,k}((\boldsymbol{y}_1,\dots,\boldsymbol{y}_i,\dots,\boldsymbol{y}_n)) \\
&= \sum_{k=1}^{m} p_{i,k}\cdot \mathcal{B}_i((\boldsymbol{y}_1,\dots,\boldsymbol{y}_i,\dots,\boldsymbol{y}_n),k) \\
&= \mathop{\mathbf{E}}_{X\sim \boldsymbol{p}_i} \mathcal{B}_i((\boldsymbol{y}_1,\dots,\boldsymbol{y}_i,\dots,\boldsymbol{y_n}),X).
\end{aligned}$$

The second and fifth equalities follow 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. □

## B. Proof of Theorem 2

THEOREM 2. *An incentive-compatible fair division mechanism $\mathcal{A}$ is proportional if and only if the corresponding incentive-compatible wagering mechanism $\mathcal{B}$ is individually rational.*

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

$$u_i(\mathcal{A}_i(\boldsymbol{y}),\boldsymbol{y}_i) = \boldsymbol{y}_i\cdot \mathcal{A}_i(\boldsymbol{y}) = \sum_{k=1}^{m} y_{i,k}\cdot \mathcal{A}_{i,k}(\boldsymbol{y}) = \sum_{k=1}^{m} y_{i,k}\cdot \mathcal{B}_i(\boldsymbol{y},k) = \mathop{\mathbf{E}}_{X\sim \boldsymbol{y_i}} \mathcal{B}_i(\boldsymbol{y},X) \geq \frac{1}{n}.$$

The third equality follows from the definition of the correspondence. The inequality follows 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 and let $i \in [n]$. Assume that agent $i$ truthfully reports $\boldsymbol{y}_i = \boldsymbol{p}_i$ (this is sufficient since for individual rationality we must only show that *some* report yields an expected payoff of at least $\frac{1}{n}$).We have:

$$\mathop{\mathbf{E}}_{X\sim \boldsymbol{p}_i} \mathcal{B}_i(\boldsymbol{y},X) = \sum_{k=1}^{m} p_{i,k}\cdot \mathcal{B}_i(\boldsymbol{y},k) = \sum_{k=1}^{m} p_{i,k}\cdot \mathcal{A}_{i,k}(\boldsymbol{y}) = \boldsymbol{p}_i\cdot \mathcal{A}_i(\boldsymbol{y}) \geq \frac{1}{n}$$

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

□

## C. Proof of Theorem 3

THEOREM 3. *If an anonymous, incentive-compatible wagering mechanism is normal then the corresponding anonymous, incentive-compatible fair division mechanism is envy-free.*

Proof. We begin by showing that anonymity is preserved by the correspondence. Let $i \in [n]$, $k \in [m]$, and let $\sigma$ be a permutation of $[n]$. Suppose that fair division mechanism $\mathcal{A}$ is anonymous. Then, for the wagering mechanism $\mathcal{B}$ corresponding to $\mathcal{A}$, we have

$$\mathcal{B}_i((\boldsymbol{y}_1,\dots,\boldsymbol{y}_n),k)=\mathcal{A}_{i,k}((\boldsymbol{y}_1,\dots,\boldsymbol{y}_n))=\mathcal{A}_{\sigma(i),k}((\boldsymbol{y}_{\sigma^{-1}(1)},\dots,\boldsymbol{y}_{\sigma^{-1}(n)}))=\mathcal{B}_{\sigma(i)}((\boldsymbol{y}_{\sigma^{-1}(1)},\dots,\boldsymbol{y}_{\sigma^{-1}(n)}),k).$$

On the other hand, if $\mathcal{B}$ is anonymous, then

$$\mathcal{A}_{i,k}((\boldsymbol{y}_1,\dots,\boldsymbol{y}_n))=\mathcal{B}_i((\boldsymbol{y}_1,\dots,\boldsymbol{y}_n),k)=\mathcal{B}_{\sigma(i)}((\boldsymbol{y}_{\sigma^{-1}(1)},\dots,\boldsymbol{y}_{\sigma^{-1}(n)}),k)=\mathcal{A}_{\sigma(i),k}((\boldsymbol{y}_{\sigma^{-1}(1)},\dots,\boldsymbol{y}_{\sigma^{-1}(n)})).$$

Having established that anonymity is preserved, we proceed by showing the contrapositive. Let $\mathcal{A}$ be an anonymous, weakly incentive compatible fair division mechanism that is not envy-free. That is, there exist agents $i,j$ and profile of reports $\boldsymbol{y}$ such that

$$u_j(\mathcal{A}_i(\boldsymbol{y}),\boldsymbol{y}_j)=\boldsymbol{y}_j\cdot\mathcal{A}_i(\boldsymbol{y})>\boldsymbol{y}_j\cdot\mathcal{A}_j(\boldsymbol{y})=u_j(\mathcal{A}_j(\boldsymbol{y}),\boldsymbol{y}_j). \tag{6}$$

Now, consider the profile $\boldsymbol{y'}$ that is derived from $\boldsymbol{y}$ by changing $i$'s report to be $\boldsymbol{y}_j$ and leaving all other agents' reports unchanged, that is, $\boldsymbol{y}'_i=\boldsymbol{y}_j$ and $\boldsymbol{y}'_{j'}=\boldsymbol{y}_{j'}$ for all $j'\neq i$. By anonymity, it must be the case that

$$\mathcal{A}_j(\boldsymbol{y'})=\mathcal{A}_i(\boldsymbol{y'}). \tag{7}$$

Further, because $\boldsymbol{y'}$ and $\boldsymbol{y}$ differ only in the report of agent $i$, incentive compatibility says that

$$u_i(\mathcal{A}_i(\boldsymbol{y'}),\boldsymbol{y}'_i)=\boldsymbol{y}'_i\cdot\mathcal{A}_i(\boldsymbol{y'})\geq\boldsymbol{y}'_i\cdot\mathcal{A}_i(\boldsymbol{y})=u_i(\mathcal{A}_i(\boldsymbol{y}),\boldsymbol{y}'_i). \tag{8}$$

Since $\boldsymbol{y}'_i=\boldsymbol{y}_j$, Equation 8 is equivalent to

$$\boldsymbol{y}_j\cdot\mathcal{A}_i(\boldsymbol{y'})\geq\boldsymbol{y}_j\cdot\mathcal{A}_i(\boldsymbol{y}). \tag{9}$$

Combining Equations 7, 9, and 6 (in that order) yields

$$\boldsymbol{y}_j\cdot\mathcal{A}_j(\boldsymbol{y'})=\boldsymbol{y}_j\cdot\mathcal{A}_i(\boldsymbol{y'})\geq\boldsymbol{y}_j\cdot\mathcal{A}_i(\boldsymbol{y})>\boldsymbol{y}_j\cdot\mathcal{A}_j(\boldsymbol{y}).$$

Taking $\boldsymbol{\theta}=\boldsymbol{y_j}$ exhibits a violation of normality of the corresponding wagering mechanism: $\mathbf{E}_{X\sim\boldsymbol{\theta}}\,\mathcal{B}_i((\boldsymbol{y}_1,\dots,\boldsymbol{y}'_i,\dots,\boldsymbol{y_n}),X)\geq\mathbf{E}_{X\sim\boldsymbol{\theta}}\,\mathcal{B}_i((\boldsymbol{y}_1,\dots,\boldsymbol{y}_i,\dots,\boldsymbol{y_n}),X)$ but $\mathbf{E}_{X\sim\boldsymbol{\theta}}\,\mathcal{B}_j((\boldsymbol{y}_1,\dots,\boldsymbol{y}'_i,\dots,\boldsymbol{y_n}),X)>\mathbf{E}_{X\sim\boldsymbol{\theta}}\,\mathcal{B}_j((\boldsymbol{y}_1,\dots,\boldsymbol{y}_i,\dots,\boldsymbol{y_n}),X)$. That is, agent $i$ changing her report from $\boldsymbol{y}_i$ to $\boldsymbol{y}'_i$ weakly increased agent $i$'s expected payoff and strictly increased agent $j$'s expected payoff (both with respect to $\boldsymbol{y}_j$). This violates normality since it is either the case that $\mathbf{E}_{X\sim\boldsymbol{\theta}}\,\mathcal{B}_i((\boldsymbol{y}_1,\dots,\boldsymbol{y}'_i,\dots,\boldsymbol{y_n}),X)>\mathbf{E}_{X\sim\boldsymbol{\theta}}\,\mathcal{B}_i((\boldsymbol{y}_1,\dots,\boldsymbol{y}_i,\dots,\boldsymbol{y_n}),X)$ or that $\mathbf{E}_{X\sim\boldsymbol{\theta}}\,\mathcal{B}_i((\boldsymbol{y}_1,\dots,\boldsymbol{y}'_i,\dots,\boldsymbol{y_n}),X)=\mathbf{E}_{X\sim\boldsymbol{\theta}}\,\mathcal{B}_i((\boldsymbol{y}_1,\dots,\boldsymbol{y}_i,\dots,\boldsymbol{y_n}),X)$. In the former case normality dictates that $\mathbf{E}_{X\sim\boldsymbol{\theta}}\,\mathcal{B}_j((\boldsymbol{y}_1,\dots,\boldsymbol{y}'_i,\dots,\boldsymbol{y_n}),X)<\mathbf{E}_{X\sim\boldsymbol{\theta}}\,\mathcal{B}_j((\boldsymbol{y}_1,\dots,\boldsymbol{y}_i,\dots,\boldsymbol{y_n}),X)$, while in the latter case normality dictates that $\mathbf{E}_{X\sim\boldsymbol{\theta}}\,\mathcal{B}_j((\boldsymbol{y}_1,\dots,\boldsymbol{y}'_i,\dots,\boldsymbol{y_n}),X)=\mathbf{E}_{X\sim\boldsymbol{\theta}}\,\mathcal{B}_j((\boldsymbol{y}_1,\dots,\boldsymbol{y}_i,\dots,\boldsymbol{y_n}),X)$. In both cases normality is violated. $\square$

## D. Proof of Proposition 2

PROPOSITION 2. *Let $R$ be a strictly proper scoring rule and let $\boldsymbol{y}$ be a profile of reports for which there exist $j,\ell$ with $\boldsymbol{y}_j \neq \boldsymbol{y}_\ell$. Then the utility of every agent $i \in [n]$ under CSR $\mathcal{C}^R$ is*

$$u_i(\mathcal{C}_i^R(\boldsymbol{y}), \boldsymbol{y}_i) > 1/n.$$

Proof. We prove the following stronger statement. It says that if an agent $i$ receives strictly higher score from $R$ in expectation (according to $i$'s report) than some other agent $j$, then agent $i$ receives utility strictly higher than $1/n$ from Competitive Scoring Rule $\mathcal{C}^R$.

LEMMA 2. *Let $R$ be a proper scoring rule. Then, for every profile of reports $\boldsymbol{y} \in \Delta_m^n$ and for all agents $i \in [n]$ for which there exists an agent $j \neq i$ with $\mathbf{E}_{X\sim \boldsymbol{y}_i} R(\boldsymbol{y}_i, X) > \mathbf{E}_{X\sim \boldsymbol{y}_i} R(\boldsymbol{y}_j, X)$,*

$$u_i(\mathcal{C}_i^R(\boldsymbol{y}), \boldsymbol{y}_i) > 1/n.$$

Proof of Lemma 2. Let $\boldsymbol{y} \in \Delta_m^n$, $i \in [n]$, and suppose that $\mathbf{E}_{X\sim \boldsymbol{y}_i} R(\boldsymbol{y}_i, X) > \mathbf{E}_{X\sim \boldsymbol{y}_i} R(\boldsymbol{y}_j, X)$ for some $j \neq i$. Therefore,

$$\begin{aligned}
u_i(\mathcal{C}_i^R(\boldsymbol{y}), \boldsymbol{y}_i) &= \sum_{k=1}^m y_{i,k} \cdot \mathcal{C}_{i,k}^R(\boldsymbol{y}) \\
&= \sum_{k=1}^m y_{i,k} \cdot \left(\frac{1}{n} + \frac{1}{n}\left(R(\boldsymbol{y}_i, k) - \frac{1}{n-1}\sum_{\ell \neq i} R(\boldsymbol{y}_\ell, k)\right)\right) \\
&= \frac{1}{n} + \frac{1}{n}\left(\sum_{k=1}^m y_{i,k} \cdot R(\boldsymbol{y}_i, k) - \frac{1}{n-1}\sum_{\ell \neq i}\sum_{k=1}^m y_{i,k} \cdot R(\boldsymbol{y}_\ell, k)\right) \\
&> \frac{1}{n},
\end{aligned}$$

where the final inequality holds because $\sum_{k=1}^m y_{i,k} \cdot R(\boldsymbol{y}_i, k) \geq \sum_{k=1}^m y_{i,k} \cdot R(\boldsymbol{y}_\ell, k)$ for all $\ell \in [n]$ by properness of $R$, and $\sum_{k=1}^m y_{i,k} \cdot R(\boldsymbol{y}_i, k) > \sum_{k=1}^m y_{i,k} \cdot R(\boldsymbol{y}_j, k)$ by assumption. □

Proposition 2 follows as a corollary of Lemma 2. To see this, suppose that $R$ is strictly proper and that there exist agents $j,\ell \in [n]$ for which $\boldsymbol{y}_j \neq \boldsymbol{y}_\ell$. Let $i \in [n]$. Then it cannot be the case that $\boldsymbol{y}_i = \boldsymbol{y}_j$ and $\boldsymbol{y}_i = \boldsymbol{y}_\ell$. Suppose without loss of generality that $\boldsymbol{y}_i \neq \boldsymbol{y}_j$. By strict properness of $R$, $\mathbf{E}_{X\sim \boldsymbol{y}_i} R(\boldsymbol{y}_i, X) > \mathbf{E}_{X\sim \boldsymbol{y}_i} R(\boldsymbol{y}_j, X)$. Thus we can apply Lemma 2 to show that $u_i(\mathcal{C}_i^R(\boldsymbol{y}), \boldsymbol{y}_i) = \boldsymbol{y}_i \cdot \mathcal{C}_i^R(\boldsymbol{y}) > 1/n$. □

## E. Proof of Theorem 5

THEOREM 5. *For $n = 2$ agents, a fair division mechanism $\mathcal{A}$ is swap-dictatorial if and only if $\mathcal{A}$ is a Competitive Scoring Rule $\mathcal{C}^R$ for some proper scoring rule $R$, i.e., if and only if $\mathcal{A}_{i,k}(\boldsymbol{y}) = \frac{1}{2} + \frac{1}{2}\big(R(\boldsymbol{y}_i, k) - R(\boldsymbol{y}_{3-i}, k)\big)$ for all $i \in \{1,2\}$ and all $k \in [m]$.*

Proof. We first prove the backward direction, that all CSRs are swap-dictatorial. To see this, note that for $n = 2$, Equation 1 can be rewritten as

$$\mathcal{C}_{i,k}^R(\boldsymbol{y}) = \frac{1}{2}R(\boldsymbol{y}_i, k) + \frac{1}{2}\big(1 - R(\boldsymbol{y}_{3-i}, k)\big). \tag{10}$$

Equation 10 can be interpreted as selecting each agent $i \in \{1,2\}$ to be the dictator with probability 0.5. When agent $i$ is the dictator, she is allocated an $R(\boldsymbol{y}_i, k)$ fraction of every item $k$ with the remaining items being allocated to agent $3-i$. The set of bundles $D$ that is offered to the dictator is $\{(R(\boldsymbol{y}_i, k))_{k\in[m]} : \boldsymbol{y}_i \in \Delta_m\}$. By properness of $R$, the dictator is allocated a bundle that maximizes her utility. Therefore, $\mathcal{A}$ is swap-dictatorial.

We next show the forward direction, that all swap-dictatorial mechanisms are CSRs. Let $\mathcal{A}$ be a swap-dictatorial mechanism. Define scoring rule $R$ as follows:

$$R(\boldsymbol{y}_i, k) = \left(\arg\max_{\boldsymbol{d}\in D} u_i(\boldsymbol{d}, \boldsymbol{y}_i)\right)_k$$

That is, for report $\boldsymbol{y}_i$ and outcome $k$, $R$ pays an amount equal to the fraction of item $k$ that an agent receives when she is the dictator with report $\boldsymbol{y}_i$. $R$ is proper because, among available bundles/contingent payments, the agent receives the one that maximizes her reported utility function. To see that $\mathcal{A}$ is a CSR, note that, by definition, when agent $i$ is the dictator, she receives $R(\boldsymbol{y}_i, k)$ of every item $k \in [m]$. Similarly, when agent $3-i$ is the dictator, agent $i$ receives $1 - R(\boldsymbol{y}_{3-i}, k)$ of every item $k \in [m]$. In expectation over the two choices of dictator, agent $i$ receives

$$\mathcal{C}^R_{i,k}(\boldsymbol{y}) = \frac{1}{2}R(\boldsymbol{y}_i, k) + \frac{1}{2}\big(1 - R(\boldsymbol{y}_{3-i}, k)\big)$$

of every item $k \in [m]$. □

## F. Proof of Proposition 3

PROPOSITION 3. *For $m = 2$, the Linear Increasing Price Mechanism (LIP) with parameter $a$ of Guo and Conitzer (2010) is $\mathcal{C}^{R_{t-l}}$.*

Proof. Claim 5 of Guo and Conitzer (2010) describes the fraction of each item that an agent receives when she is the dictator (recall that LIPs are swap-dictatorial). Since every anonymous swap-dictatorial mechanism is a competitive scoring rule with the scoring rule defined by the menu of allocations offered to the dictator, the proposition follows from applying simple algebra to Claim 5 of Guo and Conitzer (2010). □

## G. Proof of Proposition 4

PROPOSITION 4. *The Sphere Mechanism of Han et al. (2011, Definition 5) is $\mathcal{C}^{R_s}$.*

Proof. The result follows directly from applying the equivalence to the Sphere Mechanism of Han et al. (2011, Definition 5). □

## H. Proof of Proposition 5

PROPOSITION 5. *The mechanism of Cheung (2016, Section 4) is $\mathcal{C}^{R_{\text{Cheung}}}$.*

Proof. The result follows from combining Cheung's Observation 1 (which gives the CSR form) and his Equation 6 (which gives the particular scoring rule), along with the equivalence. □

## I. Proof of Proposition 6

PROPOSITION 6. *For $n = 2$, the Constrained Serial Dictatorship (CSD) mechanism of Aziz and Ye (2014) is $\mathcal{C}^{R_{\text{CSD}}}$.*

Proof. For the purposes of this proof, let $\overline{S}_i$ denote the set of $\lfloor \frac{m}{2} \rfloor$ outcomes that have the highest valuation according to $\boldsymbol{y}_i$, with ties broken in favor of lower-indexed outcomes. Similarly, let $\underline{S}_i$ denote the set of $\lfloor \frac{m}{2} \rfloor$ outcomes that have the lowest valuation according to $\boldsymbol{y}_i$, with ties broken in favor of higher-indexed outcomes. In case $m$ is odd, let $S_i = [m] \setminus (\underline{S}_i \cup \overline{S}_i)$ denote the remaining item.

Let $k \in [m]$ and WLOG consider agent $i$. Since item $k$ can be a member of $\overline{S}_i$, $\underline{S}_i$, or $S_i$ for each agent $i$, there are nine mutually exclusive cases to consider. In each case, we verify that the fraction of item $k$ allocated to agent $i$ by CSD is $\mathcal{C}_i^{R\text{CSD}}(\boldsymbol{y}) = \frac{1}{2} + \frac{1}{2}(R_{\text{CSD}}(\boldsymbol{y}_i, k) + R_{\text{CSD}}(\boldsymbol{y}_{3-i}, k))$. Suppose that $k \in \overline{S}_i$. Then, whenever agent $i$ chooses first in the CSD ordering, she receives all of item $k$. If $k \in \overline{S}_{3-i}$, then agent $3-i$ also receives all of item $k$ when she chooses first. Thus, agent $i$ is allocated a $0.5 = \frac{1}{2} + \frac{1}{2}(R_{\text{CSD}}(\boldsymbol{y}_i, k) + R_{\text{CSD}}(\boldsymbol{y}_{3-i}, k))$ fraction of item $k$ by CSD. If $k \in \underline{S}_{3-i}$, then agent $3-i$ does not receive item $k$ when she is the dictator. So, in this case, agent $i$ receives all of item $k$, which again aligns with $\frac{1}{2} + \frac{1}{2}(R_{\text{CSD}}(\boldsymbol{y}_i, k) + R_{\text{CSD}}(\boldsymbol{y}_{3-i}, k))$. Likewise, if $k \in S_{3-i}$ then agent $3-i$ receives exactly half of item $k$ when she is the dictator, meaning that agent $i$ receives a $0.75 = \frac{1}{2} + \frac{1}{2}(R_{\text{CSD}}(\boldsymbol{y}_i, k) + R_{\text{CSD}}(\boldsymbol{y}_{3-i}, k))$ fraction of item $k$ in total. It is easy to check that the equivalence also holds in the remaining six cases. □

## J. Proof of Theorem 7

THEOREM 7. *Let $n = 2$ and consider Competitive Scoring Rule $\mathcal{C}^R$, defined using scoring rule $R$ with associated expected loss function $L^R$. Then, the utility of agent $i$ is*

$$u_i(\mathcal{C}_i^R(\boldsymbol{y}), \boldsymbol{y}_i) = \frac{1}{2} + \frac{1}{2} L^R(\boldsymbol{y}_i, \boldsymbol{y}_{3-i}). \tag{2}$$

*Hence, Competitive Scoring Rule $\mathcal{C}^R$ achieves utilitarian welfare*

$$\mathbf{U}(\mathcal{C}^R(\boldsymbol{y}), \boldsymbol{y}) = \frac{1}{2} + \frac{1}{4}\left(L^R(\boldsymbol{y}_1, \boldsymbol{y}_2) + L^R(\boldsymbol{y}_2, \boldsymbol{y}_1)\right) \tag{3}$$

*and Nash welfare*

$$\mathbf{N}(\mathcal{C}^R(\boldsymbol{y}), \boldsymbol{y}) = \frac{1}{2} \cdot \sqrt{(1 + L^R(\boldsymbol{y}_1, \boldsymbol{y}_2))\,(1 + L^R(\boldsymbol{y}_2, \boldsymbol{y}_1))}. \tag{4}$$

Proof. Looking at the definition of CSRs (Definition 13 on p. 12), we see that, for $n = 2$, the CSR score expected by agent $i$ for truthfully reporting $\boldsymbol{y}_i = \boldsymbol{p}_i$ can be written as:

$$\begin{aligned}
\mathop{\mathbf{E}}_{X \sim y_i}\left[\mathcal{C}_i^R(\boldsymbol{y}, X)\right] &= \mathop{\mathbf{E}}_{X \sim y_i}\left[\frac{1}{2} + \frac{1}{2}\big(R(\boldsymbol{y}_i, X) - R(\boldsymbol{y}_{3-i}, X)\big)\right] \\
&= \frac{1}{2} + \frac{1}{2} \mathop{\mathbf{E}}_{X \sim y_i}\left[R(\boldsymbol{y}_i, X) - R(\boldsymbol{y}_{3-i}, X)\right] \\
&= \frac{1}{2} + \frac{1}{2} L^R(\boldsymbol{y}_i, \boldsymbol{y}_{3-i})
\end{aligned}$$

Following the definitions of utilitarian and Nash welfare (p. 5), we thus obtain

$$\mathbf{U}(\mathcal{C}^R(\boldsymbol{y}), \boldsymbol{y}) = \frac{1}{2}\left(\frac{1}{2} + \frac{1}{2} L^R(\boldsymbol{y}_1, \boldsymbol{y}_2) + \frac{1}{2} + \frac{1}{2} L^R(\boldsymbol{y}_2, \boldsymbol{y}_1)\right) = \frac{1}{2} + \frac{1}{4}\left(L^R(\boldsymbol{y}_1, \boldsymbol{y}_2) + L^R(\boldsymbol{y}_2, \boldsymbol{y}_1)\right)$$

$$\mathbf{N}(\mathcal{C}^R(\boldsymbol{y}), \boldsymbol{y}) = \sqrt{\left(\frac{1}{2} + \frac{1}{2} L^R(\boldsymbol{y}_1, \boldsymbol{y}_2)\right) \cdot \left(\frac{1}{2} + \frac{1}{2} L^R(\boldsymbol{y}_2, \boldsymbol{y}_1)\right)} = \frac{1}{2} \cdot \sqrt{(1 + L^R(\boldsymbol{y}_1, \boldsymbol{y}_2))\,(1 + L^R(\boldsymbol{y}_2, \boldsymbol{y}_1))}$$

□

**This page is intentionally blank. Proper e-companion title page, with INFORMS branding and exact metadata of the main paper, will be produced by the INFORMS office when the issue is being assembled.**

# Electronic Companion

## EC.1. Description of Known Mechanisms

Table 1 on page 22 gives an overview of all known incentive-compatible mechanisms. Competitive Scoring Rules were introduced in Section 4; what follows is a description of the other mechanisms from that table.

**Competitive Equilibrium from Equal Incomes (CEEI).**

Competitive Equilibrium from Equal Incomes has been explored in both the fair division and the wagering settings. In the context of fair division, CEEI endows each agent $i$ with 1 unit of currency and simulates a market equilibrium in which each item has a price and all agents spend their entire budget on items that maximize their utility-to-price ratio (Foley 1967, Varian 1974). In the wagering context, this is equivalent to the parimutuel consensus mechanism of Eisenberg and Gale (1959).

In the fair division setting, CEEI is envy-free, proportional, and Pareto optimal, but not incentive compatible. In the wagering setting, the parimutuel consensus mechanism is individually rational and strictly budget balanced, but of course still not incentive compatible.

**No-Arbitrage Wagering Mechanisms (NAWMs).**

No-Arbitrage Wagering Mechanisms (Chen et al. 2014) are defined as follows. Let $R$ be a (strictly) proper scoring rule bounded by $[0,1]$, and for every agent $i$ let $\overline{\boldsymbol{y}}_{-i}$ denote the mean of the reports of all other agents.[22] NAWMs pay every agent $i$ an amount

$$\frac{1}{n} + \frac{n-1}{n^2}[R(\boldsymbol{y}_i, k) - R(\overline{\boldsymbol{y}}_{-i}, k)].$$

NAWMs are individually rational, strictly incentive compatible (when defined using a strictly proper scoring rule), and weakly (but not strictly) budget balanced. It therefore follows from the correspondence and the results in Section 3 that the corresponding fair division mechanisms are proportional, strictly incentive compatible, and wasteful (and therefore not Pareto optimal). Additionally, they are not envy-free, as the following example illustrates.

Example EC.1. Consider the NAWM defined using the quadratic scoring rule. Let $n=3$ and $m=2$, and suppose that reports are given by $\boldsymbol{y}_1 = (0,1)$, $\boldsymbol{y}_2 = (0.9, 0.1)$, and $\boldsymbol{y}_3 = (1,0)$. Agent 2 receives a $\frac{1}{3} + \frac{2}{9}(0.99 - 0.75) = 0.3867$ share of item 1, and a $\frac{1}{3} + \frac{2}{9}(0.19 - 0.75) = 0.2089$ share of item 2. For agent 3, the shares are $\frac{1}{3} + \frac{2}{9}(1 - 0.6975) = 0.4006$ and $\frac{1}{3} + \frac{2}{9}(0 - 0.7975) = 0.1561$, respectively. The utility that agent 2 achieves is therefore $0.9 \cdot 0.3867 + 0.1 \cdot 0.2089 = 0.3689$, while agent 3's allocation, according to agent 2's report, achieves utility of $0.9 \cdot 0.4006 + 0.1 \cdot 0.1561 = 0.3761$. Therefore, agent 2 envies agent 3.

[22] Different ways to aggregate the reports of the other agents are possible but we do not consider them here. See the paper of Chen et al. (2014) for details.

**Constrained Serial Dictatorship (CSD).**

Constrained Serial Dictatorship (Aziz and Ye 2014) is defined as follows. Imagine fixing a particular ordering of the agents, allocating, in order, to each agent her most preferred $m/n$ of the remaining items (allowing partial items to be allocated). The output of CSD is the expected allocation that would arise by selecting one of these orderings uniformly at random. CSD is known to be weakly (not strictly) incentive compatible, non-wasteful and proportional, but not envy-free or Pareto optimal. Moreover, computing the outcome of CSD is #P complete, which renders an exact algorithm infeasible for larger $n$. (Variants of CSD employ sampling to circumvent the computational complexity issue but they no longer guarantee proportionality.)

**The Double Clinching Auction (DCA).**

The Double Clinching Auction (Freeman et al. 2017) is a wagering mechanism that defines the outcome-dependent payments through an auction, namely the Adaptive Clinching Auction (Dobzinski et al. 2008). The details are beyond the scope of this paper and we refer the reader to the original work (Freeman et al. 2017) for additional details. Note that the Double Clinching Auction is only defined for the case of $m = 2$ and $n \geq 4$.

DCA is individually rational, weakly (not strictly) incentive compatible, and weakly (not strictly) budget balanced. Therefore, the corresponding fair division mechanism is proportional, weakly incentive compatible, and wasteful (and therefore not Pareto optimal). We additionally show that DCA is envy-free.

Theorem EC.1. *The Double Clinching Auction is envy-free.*

Proof Sketch. The Adaptive Clinching Auction, a crucial building block of DCA, is known to be envy-free, in the sense that no bidder prefers any other bidder's allocation and payment to her own (Devanur et al. 2013, Theorem 4.3). Since each security is an outcome-contingent payment under DCA, this is exactly equivalent to saying that every agent obtains higher expected payment from the payoffs assigned to her by the mechanism than she would obtain from any other agent's payoffs. When translated to the fair division setting, this guarantee gives envy-freeness of DCA. □

**Partial Allocation (PA).**

The Partial Allocation mechanism (Cole et al. 2013b, 2022) is designed to approximate CEEI, providing each agent at least a $1/e$ fraction of the utility she would receive under CEEI. 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 the work of Cole et al. (2013b, 2022) for further details. PA is known to be weakly (not strictly) incentive compatible and envy-free. However, it is wasteful (and therefore not Pareto optimal) and not proportional.

**Strong Demand Matching (SDM).**

Strong Demand Matching (Cole et al. 2013b, 2022) is also designed to approximate CEEI. Its approximation factor is particularly good when all items are highly demanded, for instance, when there are many more agents than items and no item is uniformly disliked. SDM 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 item. We refer the reader to the work of Cole et al. (2013b, 2022) for further details. SDM is weakly (not strictly) incentive-compatible and envy-free, but also wasteful (and therefore not Pareto optimal) and not proportional. Moreover, note that SDM violates normality, as the following example shows.

EXAMPLE EC.2. Consider an instance defined by $\boldsymbol{y}_1 = (0.5, 0.5)$ and $\boldsymbol{y}_2 = (1, 0)$. SDM allocates all of item 2 to agent 1, and all of item 1 to agent 2. According to $\boldsymbol{\theta} = (0.6, 0.4)$, agent 1 receives utility 0.4 and agent 2 receives utility 0.6. Now consider another instance defined by $\boldsymbol{y}'_1 = (1, 0)$ and $\boldsymbol{y}_2 = (1, 0)$. SDM allocates half of item 1 to agent 1, and the other half of item 1 to agent 2. According to $\boldsymbol{\theta}$, agent 1 receives utility $0.6/2 = 0.3 < 0.4$ and agent 2 receives utility $0.6/2 = 0.3 < 0.6$. Both agents have had their utility strictly reduced as a result of agent 1 changing her report from $\boldsymbol{y}_1$ to $\boldsymbol{y}'_1$, a violation of normality.

**Partial Allocation Max ($\text{PA}_{\text{max}}$).**

The $\text{PA}_{\text{max}}$ mechanism of Cole et al. (2013a)[23] is defined for $n = 2$ agents. On a given instance, it selects either the uniform allocation or the outcome of the PA mechanism, whichever gives higher utilitarian welfare. $\text{PA}_{\text{max}}$ guarantees a $\frac{2}{3}$ approximation to the optimal utilitarian welfare. $\text{PA}_{\text{max}}$ is weakly (not strictly) incentive compatible, envy-free, and proportional, but is wasteful (and therefore not Pareto optimal).

PROPOSITION EC.1. *$PA_{max}$ is envy-free and proportional.* [24]

Proof. For $n = 2$, the original PA mechanism gives both agents the same utility (Cheung 2016, Footnote 6). Denote this common utility by $x$. If $x < 0.5$, then $\text{PA}_{\text{max}}$ outputs the uniform allocation. If $x \geq 0.5$, then $\text{PA}_{\text{max}}$ outputs the PA allocation. In either case, both agents receive utility at least 0.5 from the $\text{PA}_{\text{max}}$ allocation. Therefore, $\text{PA}_{\text{max}}$ is proportional.

To see that $\text{PA}_{\text{max}}$ is envy-free, note that both PA and the uniform allocation mechanism are envy-free. Therefore, any mechanism that only returns allocations that are the output of one of these two mechanisms must also be envy-free. □

[23] Cole et al. (2013a) refer to the mechanism simply as the Max mechanism.

[24] We thank Vasilis Gkatzelis for pointing this out.

To see that $\text{PA}_{\text{max}}$ is wasteful, note that when $\boldsymbol{y}_1 = (0.8, 0.2)$ and $\boldsymbol{y}_2 = (0.3, 0.7)$, $\text{PA}_{\text{max}}$ assigns a 0.7 fraction of item 1 to agent 1 and a 0.8 fraction of item 2 to agent 2.

Note that the extended version of $\text{PA}_{\text{max}}$ due to Cheung (2016) takes a carefully crafted weighted sum of $\text{PA}_{\text{max}}$ and two generalizations of the original PA mechanism. In doing so, it achieves an improved $0.67776 > \frac{2}{3}$ approximation to the optimal utilitarian welfare while sacrificing proportionality.

## EC.2. Additional Figures

We here present additional figures for Nash welfare, complementing the results from Section 6.

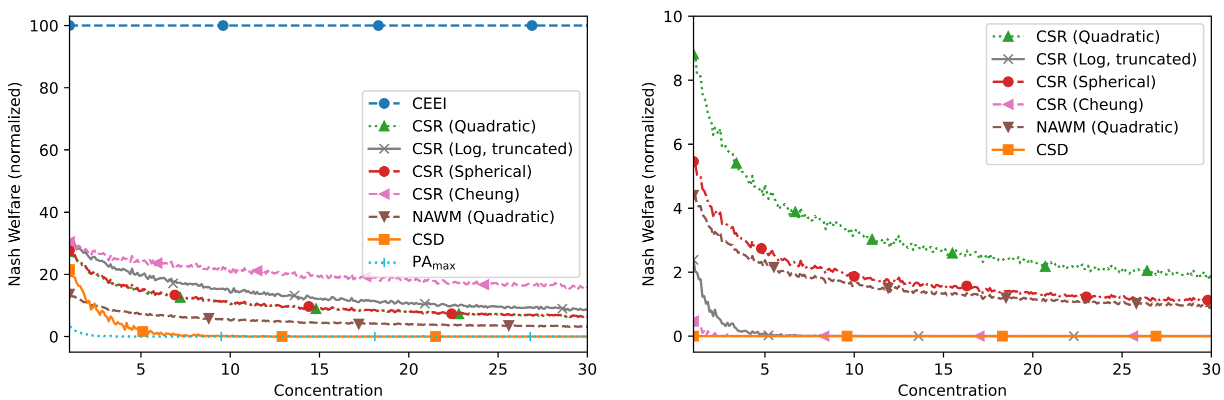


**Figure EC.1 Nash welfare of different mechanisms with $n = 2$ agents and $m = 2$ items for two different biases. Reports are centered around $(0.25, 0.75)$ (left side) and $(0.05, 0.95)$ (right side), and the x axis varies the level of concentration around those centers. On the right side, CEEI (welfare value 100) and $\text{PA}_{\text{max}}$ (welfare value 0 everywhere) are left out for presentational reasons and the y axis is scaled differently than in the other plots we present.**

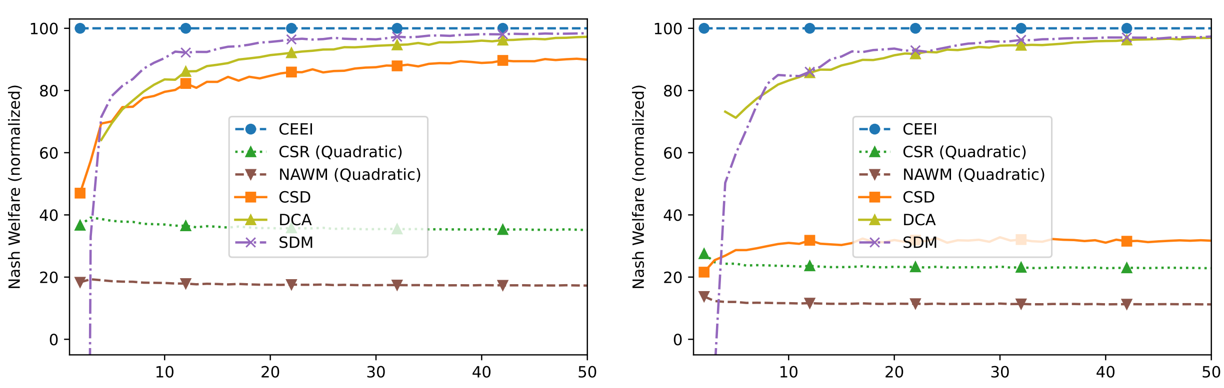


**Figure EC.2 Nash welfare of different mechanisms with $m = 2$ items with uniform reports (left side) and biased reports centered around $(0.25, 0.75)$ (right side). The number of agents $n$ varies from 2 to 50 (x axis).**

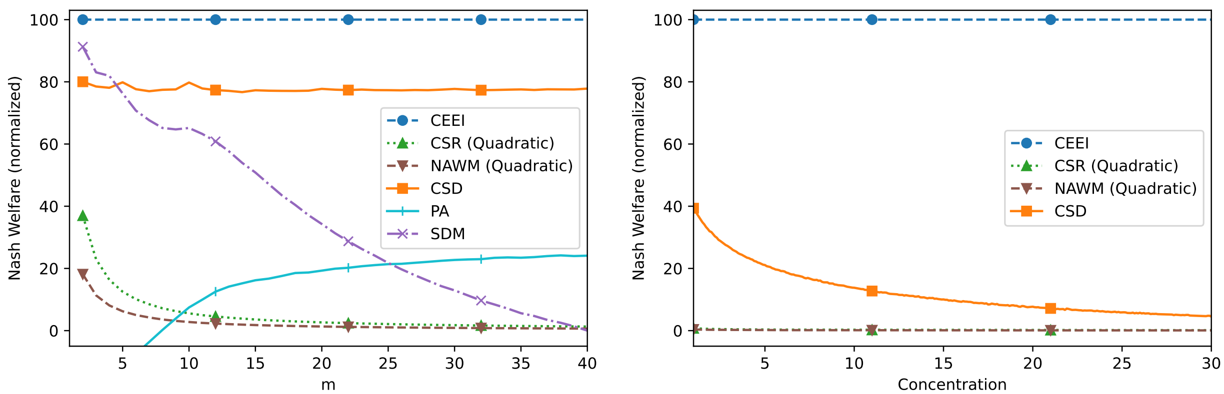


**Figure EC.3 Nash welfare of different mechanisms with $n = 10$ agents for uniform reports varying $m$ from 2 to 40 (left side) and biased report with $m = 20$ varying concentration (right side).**