

The Double Clinching Auction for Wagering

RUPERT FREEMAN, Duke University¹

DAVID M. PENNOCK, Microsoft Research

JENNIFER WORTMAN VAUGHAN, Microsoft Research

We develop the first incentive compatible and near-Pareto-optimal wagering mechanism. Wagering mechanisms can be used to elicit predictions from agents who reveal their beliefs by placing bets. Lambert et al. [20, 21] introduced weighted score wagering mechanisms, a class of budget-balanced wagering mechanisms under which agents with immutable beliefs truthfully report their predictions. However, we demonstrate that these and other existing incentive compatible wagering mechanisms are not Pareto optimal: agents have significant budget left over even when additional trade would be mutually beneficial. Motivated by this observation, we design a new wagering mechanism, the *double clinching auction*, a two-sided version of the adaptive clinching auction [9]. We show that no wagering mechanism can simultaneously satisfy weak budget balance, individual rationality, weak incentive compatibility, and Pareto optimality. However, we prove that the double clinching auction attains the first three and show in a series of simulations using real contest data that it comes much closer to Pareto optimality than previously known incentive compatible wagering mechanisms, in some cases almost matching the efficiency of the Pareto optimal (but not incentive compatible) parimutuel consensus mechanism. When the goal of wagering is to crowdsource probabilities, Pareto optimality drives participation and incentive compatibility drives accuracy, making the double clinching auction an attractive and practical choice. Our mechanism may be of independent interest as the first two-sided version of the adaptive clinching auction.

1 INTRODUCTION

Wagering mechanisms allow a principal to elicit the beliefs of a group of agents without paying them directly or taking on any risk. Each agent specifies a *belief*, her own subjective estimate of the likelihood of a future event, such as the Democratic nominee winning the 2020 U.S. Presidential election. She also specifies a monetary budget or *wager*, the maximum amount that she is willing to lose. These wagers are then collected by the principal and, after the truth is revealed, redistributed to agents in such a way that agents with more accurate predictions are more highly rewarded. Meanwhile, since agents directly report their beliefs, the principal is able to leverage the wisdom of crowds to obtain an accurate consensus forecast for the event, for example by computing an average [16], a budget-weighted average [3], a supra-Bayesian inference [22], or another aggregate measure of the forecasts [13].

Lambert et al. [20, 21] introduced the class of *weighted-score wagering mechanisms* (WSWMs), the unique wagering mechanisms to simultaneously satisfy a set of desirable properties including strict budget balance and incentive compatibility. Incentive compatibility is achieved through the use of *strictly proper scoring rules*, reward functions designed to incentivize truthful reports from risk-neutral agents. In particular, each agent’s payoff under a WSWM is proportional to the

¹Work conducted in part while R. Freeman was an intern at Microsoft Research. He is also thankful for support from NSF under award IIS-1527434 and ARO under grant W911NF-12-1-0550.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC’17, June 26–30, 2017, Cambridge, Massachusetts, USA. ACM ISBN 978-1-4503-4527-9/17/06...\$15.00.

© 2017 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

DOI: <http://dx.doi.org/10.1145/3033274.3085118>

difference between her own score and the budget-weighted score of the other agents. Chen et al. [6] later proposed the class of *no-arbitrage wagering mechanisms* (NAWMs), which are incentive compatible but only weakly budget balanced, allowing the principal to profit off of disagreement among agents. Under an NAWM, an agent’s payoff is proportional to her score minus the score of the budget-weighted average belief. To our knowledge, these mechanisms and their derivatives (such as the randomized, private WSWM of Cummings et al. [8]) are the only known incentive compatible wagering mechanisms.

As an artifact of their use of proper scoring rules, these mechanisms have one undesirable property: In general, it is not possible for any agent to lose her full wager, even if all other agents are perfectly informed. In other words, the mechanisms are not *Pareto optimal*, in the sense that agents have significant budget left over even when additional trade would be mutually beneficial. This is a serious concern in practice since agents typically gravitate to venues where they have the opportunity for large gains. If these mechanisms yield badly suboptimal allocations, agents may question the rules or simply go elsewhere. Indeed, all widely deployed wagering mechanisms, including parimutuels, bookmakers, and double auctions, feature Pareto optimality. Additionally, wagers effectively lose their meaning as budgets. This has a surprising implication on the quality of reports. Because an agent can never lose her full wager, she may be able to artificially inflate her reported budget risk-free. It turns out that when agents misreport their budgets, they can also have incentive to misreport their beliefs. (See Example 3.3.)

Motivated by this observation, we ask whether it is possible to design an incentive compatible wagering mechanism that achieves Pareto optimality without sacrificing other key properties. Unfortunately, the answer is no. We prove that no weakly incentive compatible wagering mechanism can achieve Pareto optimality along with individual rationality (meaning agents have incentive to participate) and weak budget balance. If the principal cannot force agents to participate and does not wish to subsidize the market, he must compromise on Pareto optimality. Given that, we seek an incentive compatible mechanism that is near-Pareto-optimal in practice.

Our mechanism is inspired by the observation that the output of a wagering mechanism has a natural interpretation as an allocation of securities. An agent who wins ρ_1 dollars if the Democrat is elected and loses ρ_0 dollars otherwise can equivalently be viewed as paying ρ_0 dollars up front for $\rho_0 + \rho_1$ shares of an Arrow-Debreu security worth \$1 if and only if the Democrat is elected. Thus wagering mechanisms can be viewed as allocating items (the securities) to agents, and it is natural to ask whether techniques from the auctions literature can be used. The *clinging auction* [2] produces VCG allocations and payments for multiple identical items, but VCG-style approaches cannot be applied when agents have budgets. Instead, we build on the *adaptive clinging auction* [4, 9], an extension of the clinging auction that incorporates budget constraints.

Our mechanism, the *double clinging auction* (DCA), is a two-sided version of the adaptive clinging auction. It elicits truthful reports by selling a variable number of securities to the agents via two simultaneous instances of the adaptive clinging auction, one which sells securities that pay off \$1 only if the event of interest happens (yes securities), and one which sells securities that pay off \$1 only if it does not (no securities). The principal can always sell a pair of yes and no securities for \$1 or more without risk, since he will owe exactly \$1 to the agents regardless of the outcome. Our key technical contribution is determining the number of security pairs that the principal can sell via adaptive clinging auctions in such a way that he never loses money, without incentivizing agents to misreport their beliefs.

We also show that under the double clinging auction, each agent has at least some risk of losing her entire budget, making the budget declaration risky to inflate and restoring the semantics of the wager as the largest acceptable worst-case loss.

To evaluate the efficiency of the DCA, we run a series of simulations using thousands of probability judgments about hundreds of events, collected from an online forecasting contest called ProbabilitySports [12]. We compare the performance of the DCA with WSWMs, NAWMs, and the *parimutuel consensus mechanism* [11], which is Pareto optimal but not incentive compatible. Our simulations show that the DCA is indeed significantly closer to Pareto optimal than the other incentive compatible mechanisms, sometimes approaching the efficiency of the parimutuel consensus mechanism, which was specifically designed to maximize trade. Given the results, we are optimistic that the DCA can serve as a practical wagering mechanism that both satisfies agent demand and encourages honest revelations.

We follow previous authors [6, 17, 19, 21], assuming that agents have *immutable beliefs* that do not update during wagering. Our agents “agree to disagree”, unlike Bayesian agents. While immutable beliefs and perfect Bayesian reasoning are both idealizations, the former is arguably closer to reality. In practice, overconfident opponents, each expecting to gain, trade all the time [10, 18], contradicting the no-trade theorems implied in the Bayesian setting. Other authors have explored incentive properties of wagering mechanisms with Bayesian [20] or boundedly rational [23] agents.

2 WAGERING MECHANISMS

Let X be a binary random variable or *event* with value or *outcome* in $\{0, 1\}$. For example, imagine $X = 1$ is the outcome that the Democratic nominee wins the 2020 U.S. Presidential election and $X = 0$ is the outcome that he or she loses. We consider a setting in which a principal is interested in eliciting the beliefs of a set of agents $\mathcal{N} = \{1, \dots, N\}$ about the likelihood that $X = 1$. Following the line of work initiated by Lambert et al. [21], we assume that each agent $i \in \{1, \dots, N\}$ has a private, subjective, immutable belief p_i about the probability that $X = 1$, and that agents are risk neutral up to some budget limitation. That is, each agent budgets for the largest loss that she is willing to tolerate, then maximizes her expected wealth subject to the budget constraint.

The principal operates a *wagering mechanism* in which each agent i submits a report $\hat{p}_i \in [0, 1]$, capturing her subjective belief about the likelihood that $X = 1$, and a wager $w_i \geq 0$, representing the maximum amount that she is prepared to lose. After observing the realized value of X , denoted x , the principal redistributes the agents’ wagers, rewarding agents based on their wagers and the accuracy of their reports. We denote by $\Pi_i(\hat{\mathbf{p}}, \mathbf{w}, x)$ the net payoff to agent i under reports $\hat{\mathbf{p}}$ and wagers \mathbf{w} when $X = x$. For a wagering mechanism to be valid, it must be the case that no agent can lose more than her wager (i.e., for all i , $\hat{\mathbf{p}}$, \mathbf{w} , and x , we have $\Pi_i(\hat{\mathbf{p}}, \mathbf{w}, x) \geq -w_i$) and an agent can choose not to participate by wagering 0 (i.e., $\Pi_i(\hat{\mathbf{p}}, \mathbf{w}, x) = 0$ whenever $w_i = 0$). We denote by $\hat{\mathbf{p}}_{-i}$ the predictions of all agents other than i and by \mathbf{w}_{-i} the wagers of all agents other than i .

2.1 Examples of Wagering Mechanisms and Connections to Proper Scoring Rules

There is a close connection between wagering mechanisms and *proper scoring rules* used to elicit truthful predictions from individual agents [14, 25]. A scoring rule s maps a prediction $p \in [0, 1]$ and an outcome $x \in \{0, 1\}$ to a score or reward in $\mathbb{R} \cup \{-\infty\}$. We say s is *proper* if for all $p, q \in [0, 1]$, $ps(p, 1) + (1-p)s(p, 0) \geq ps(q, 1) + (1-p)s(q, 0)$, and *strictly proper* if this inequality is strict whenever $p \neq q$. An agent who is rewarded for her prediction using a proper scoring rule therefore maximizes her expected reward by reporting her true belief, uniquely if the scoring rule is strictly proper. A common example of a strictly proper scoring rule is the Brier score [5], $s(p, x) = 1 - (x - p)^2$.

For a wagering mechanism to elicit truthful reports about the likelihood of X , it must be the case that, fixing the wagers \mathbf{w} and reports $\hat{\mathbf{p}}_{-i}$ of other agents, agent i ’s payoff Π_i is a proper scoring rule. Building on this idea, Lambert et al. [20, 21] introduced the class of *weighted score wagering*

mechanisms (WSWMs). A WSWM has a payoff function of the form

$$\Pi_i(\hat{\mathbf{p}}, \mathbf{w}, x) = w_i \left(s(\hat{p}_i, x) - \frac{\sum_{j \in \mathcal{N}} w_j s(\hat{p}_j, x)}{\sum_{j \in \mathcal{N}} w_j} \right) \quad (1)$$

where s is any strictly proper scoring rule bounded in $[0, 1]$. WSWMs are the unique wagering mechanisms to simultaneously satisfy a set of desirable axioms that includes strict budget balance (the principal neither makes nor loses money), individual rationality (all agents have incentive to participate), strict incentive compatibility (agents have incentive to truthfully reveal their beliefs about X), anonymity (all agents are treated the same), sybilproofness (agents cannot profit by creating false identities), and a normality property (loosely, if agent i changes her report to improve her own expected payoff, the expected payoffs of other agents can't increase).

Chen et al. [6] pointed out that under a WSWM, it can be possible for an agent to risklessly profit: there exist reports $\hat{\mathbf{p}}$ and wagers \mathbf{w} such that for some agent i , both $\Pi_i(\hat{\mathbf{p}}, \mathbf{w}, 1)$ and $\Pi_i(\hat{\mathbf{p}}, \mathbf{w}, 0)$ are positive. They proposed an alternative class of incentive compatible mechanisms called *no-arbitrage wagering mechanisms (NAWMs)*, in which this extra profit is instead collected by the principal. The payoff to each agent is proportional to the difference between the score of his own prediction and the score of a type of weighted average of the other agents' predictions. We will return to these mechanisms later in the paper.

2.2 Security Interpretation of Wagering Mechanisms

The output of a wagering mechanism has a natural interpretation as an allocation of Arrow-Debreu securities with payoffs that are contingent on the realization of X . We define a yes security to be a contract worth \$1 in the outcome $X = 1$ and \$0 if $X = 0$. Similarly, a no security is worth \$0 if $X = 1$ and \$1 if $X = 0$. A risk neutral agent with belief p about the likelihood that $X = 1$ would be willing to buy a yes security at any price up to p or a no security at any price up to $1 - p$. Since such trades reveal information about agents' beliefs, securities of this form are often considered in the context of prediction markets.

Suppose a wagering mechanism would yield a net payoff to agent i of $\rho_1 = \Pi_i(\hat{\mathbf{p}}, \mathbf{w}, 1)$ when $X = 1$ and $\rho_0 = \Pi_i(\hat{\mathbf{p}}, \mathbf{w}, 0)$ when $X = 0$. This is equivalent to the payoff that i would receive if she were sold $y_i = \max\{\rho_1 - \rho_0, 0\}$ yes securities and $n_i = \max\{\rho_0 - \rho_1, 0\}$ no securities for a total cost of $\pi_i = \max\{-\rho_0, -\rho_1\}$. For example, if $\rho_0 < \rho_1$, then agent i 's participation in the wagering mechanism is equivalent to agent i paying the principal $\pi_i = -\rho_0$ before X is realized and then receiving $y_i = \rho_1 - \rho_0$ from the principal in the outcome $X = 1$.

Therefore, the output of a wagering mechanism can be completely specified by a triple $(\mathbf{y}, \mathbf{n}, \boldsymbol{\pi})$, where for each agent i , $y_i \geq 0$ is the number of yes securities allocated to i , $n_i \geq 0$ is the number of no securities allocated to i , and π_i is the cost paid by i for these securities. To be a valid output, we require that for all i , either $y_i = 0$ or $n_i = 0$ (or both), and $\pi_i \leq w_i$. This requirement is without loss of generality since any (fraction of a) pair of yes and no securities can be precisely converted into (a fraction of) \$1. We rely on the securities-based interpretation of wagering mechanisms for the remainder of this paper.¹

2.3 Properties of Wagering Mechanisms

Lambert et al. [21] introduced several desirable properties for wagering mechanisms. We focus on three of these properties in our analysis: individual rationality, incentive compatibility, and budget balance. The definitions from Lambert et al. [21] are easily translated into our security-based

¹Note that in the case of WSWMs, the observation of Chen et al. [6] implies that it is possible to have $\pi_i < 0$ for some i , meaning that i is allocated securities and actually *receives* money from the principal.

representation. First, individual rationality requires that agents participate willingly; agents have nothing to lose (in expectation) by participating.

Definition 2.1. A wagering mechanism is *individually rational* if, for any player i and any subjective probability p_i , there exists a report \hat{p}_i such that for all \hat{p}_{-i}, \mathbf{w} ,

$$p_i y_i(\hat{\mathbf{p}}, \mathbf{w}) + (1 - p_i) n_i(\hat{\mathbf{p}}, \mathbf{w}) \geq \pi_i(\hat{\mathbf{p}}, \mathbf{w}).$$

Incentive compatibility requires that each agent maximizes her expected payoff by reporting truthfully, regardless of the reports and wagers of other agents.

Definition 2.2. A wagering mechanism is *weakly incentive compatible* if, for every agent i with belief p_i and all reports $\hat{\mathbf{p}}$ and wagers \mathbf{w} ,

$$\begin{aligned} p_i y_i((p_i, \hat{\mathbf{p}}_{-i}), \mathbf{w}) + (1 - p_i) n_i((p_i, \hat{\mathbf{p}}_{-i}), \mathbf{w}) - \pi_i((p_i, \hat{\mathbf{p}}_{-i}), \mathbf{w}) \\ \geq p_i y_i(\hat{\mathbf{p}}, \mathbf{w}) + (1 - p_i) n_i(\hat{\mathbf{p}}, \mathbf{w}) - \pi_i(\hat{\mathbf{p}}, \mathbf{w}). \end{aligned}$$

The mechanism satisfies *strict incentive compatibility* if the inequality is strict whenever $p_i \neq \hat{p}_i$.

Finally, a wagering mechanism is budget balanced if the principal never loses.

Definition 2.3. A wagering mechanism is *weakly budget balanced* if, for all $\hat{\mathbf{p}}$ and \mathbf{w} ,

$$\sum_{i \in N} y_i(\hat{\mathbf{p}}, \mathbf{w}) \leq \sum_{i \in N} \pi_i(\hat{\mathbf{p}}, \mathbf{w}) \quad \text{and} \quad \sum_{i \in N} n_i(\hat{\mathbf{p}}, \mathbf{w}) \leq \sum_{i \in N} \pi_i(\hat{\mathbf{p}}, \mathbf{w}).$$

The mechanism is *strictly budget balanced* if the inequalities hold with equality for all $\hat{\mathbf{p}}$ and \mathbf{w} .

3 A TRADEOFF BETWEEN EFFICIENCY AND INCENTIVE COMPATIBILITY

Our goal is to design a wagering mechanism not to maximize profit but to maximize the amount of useful and credible information gathered. In this context, both incentive compatibility and Pareto optimality are important. The former literally keeps agents honest, steering them to report their true best estimates and reassuring the principal that probabilities are not tainted by irrelevant strategic play. The latter keeps agents happy, earning them as much utility as possible without inexplicably leaving dollars on the table. Pareto optimality is standard in prediction markets, parimutuel markets, betting exchanges, and financial exchanges. A badly suboptimal allocation may confuse agents, discourage them from playing, or encourage them to inflate their budgets, as we shall see below, which may cause their probability reports to become untruthful too.

3.1 Inefficient Allocations and Budget Inflation

In this section, we consider the undesirable effects of Pareto inefficiency. We start with an example.

Example 3.1. There are $N = 4$ agents with reported beliefs $\hat{\mathbf{p}} = (0.1, 0.2, 0.5, 0.7)$ and wagers $\mathbf{w} = (1, 1, 1, 1)$. Under the Brier scoring rule WSWM, the outcome is $(\mathbf{y} = (0, 0, 0.25, 0.65), \mathbf{n} = (0.55, 0.35, 0, 0), \boldsymbol{\pi} = (0.36, 0.19, 0.05, 0.29))$.

Observe that in Example 3.1, no agent stands to lose her full wager, regardless of the outcome. Indeed, the closest is agent 1, who risks losing 36% of her wager in the worst case. The *total risk*—the sum of all the agents' worst-case losses—is less than 25% of the total wagers. Thus WSWM is facilitating much less trade than if the agents were left to trade amongst themselves. Further, consider the 0.9 yes and 0.9 no securities allocated in Example 3.1. Thinking of these securities as any other commodity, we see that their allocation is not efficient. Some yes securities are allocated to agent 3 even though agent 4 has both a higher valuation and leftover budget.

The example above is in no way an edge case or specially manufactured; we will see in Section 6 that, if anything, it shows higher-than-average efficiency compared to our real-data simulations. Indeed, the following observation, which was originally made by Lambert et al. [21], shows that under a WSWM, agents who report any uncertainty will never lose their entire wager.

PROPOSITION 3.2 (LAMBERT ET AL. 21). *For any weighted score wagering mechanism, for any $i \in \mathcal{N}$ and any reports $\hat{\mathbf{p}}$ and \mathbf{w} , if $\hat{p}_i \in (0, 1)$ and $w_i > 0$, then $\pi_i(\hat{\mathbf{p}}, \mathbf{w}) < w_i$.*

This observation has an important implication that goes beyond the desire to facilitate as much trade as possible. Because an agent who reports her true budget w_i can never lose it all, she may be able to report a higher budget w'_i such that her maximum loss is still bounded by w_i . It turns out that when agents misreport their budgets, they may also have incentive to misreport their beliefs.

Example 3.3. In Example 3.1, agent 4 derives utility $0.65 \cdot 0.7 - 0.29 = 0.17$ for being allocated 0.65 yes securities at a cost of 0.29, since she values each yes security at 0.7. However, since this does not exhaust her budget, she could inflate her budget to $w'_4 = 2.04$ and instead be allocated 1.05 yes securities at a cost of 0.47, deriving utility 0.26. This budget inflation is completely safe in the sense that she never loses more than her budget, regardless of the reports and wagers of the other agents, (even if all other agents have arbitrarily large budgets and perfectly predict the outcome).

However, if agent 4 lowers her probability report to $\hat{p}_4 = 0.6$, she is able to inflate her budget even further. Intuitively, this is because 0.6 is a more moderate report than 0.7, so that even if $X = 0$, her loss will be lower. Agent 4 can safely report $w'_4 = 2.78$ along with $\hat{p}_4 = 0.6$ without any risk of spending more than her budget, regardless of the reports and wagers of the other agents. She is then allocated 0.96 yes securities at a cost of 0.38. Her expected utility is now $0.96 \cdot 0.7 - 0.38 = 0.30$, which is higher than she could safely obtain by truthfully reporting $\hat{p}_i = 0.7$.

3.2 Pareto Optimality

In this section, we define a natural notion of Pareto optimality for wagering mechanisms. For a fixed number of securities, a Pareto optimal allocation is, as usual, any locally optimal allocation that cannot be improved for one agent without harming others. However, the number of pairs of securities is not fixed: the principal or the agents can always manufacture more yes-no pairs at the cost of \$1. Given this, we need a slightly expanded definition of Pareto optimality.

We say that a wagering mechanism is *Pareto optimal* if, treating agents' reports and wagers as their true beliefs and budgets, after all yes and no securities have been allocated and payments for these securities collected by the principal, there is no side bet that agents could make that would strictly benefit one without harming another, even if agents are allowed to create their own securities. We first define the notion of a profitable side bet.

Definition 3.4. Given reports $\hat{\mathbf{p}}$, wagers \mathbf{w} , allocations \mathbf{y} and \mathbf{n} of yes and no securities, and payments $\boldsymbol{\pi}$, a triple $(\Delta\mathbf{y}, \Delta\mathbf{n}, \Delta\boldsymbol{\pi})$ is a *profitable side bet* if the following three conditions hold:

- (1) $\sum_{i \in \mathcal{N}} \Delta y_i = \sum_{i \in \mathcal{N}} \Delta n_i = \sum_{i \in \mathcal{N}} \Delta \pi_i = 0$.
- (2) For all $i \in \mathcal{N}$, $\min\{y_i + \Delta y_i, n_i + \Delta n_i\} - (\pi_i + \Delta \pi_i) \geq -w_i$.
- (3) For all $i \in \mathcal{N}$, $\hat{p}_i \Delta y_i + (1 - \hat{p}_i) \Delta n_i \geq \Delta \pi_i$, with strict inequality for at least one i .

Let's examine this definition. The first condition ensures that $(\Delta\mathbf{y}, \Delta\mathbf{n}, \Delta\boldsymbol{\pi})$ is a valid exchange among the agents, that is, all cash or securities given to one agent must come from other agents. The second condition ensures that no agent's budget is violated. The third guarantees that the exchange is profitable for at least one agent without harming any other agent (assuming truthful reports). We can now formally define Pareto optimality.

Definition 3.5. A wagering mechanism is *Pareto optimal* if for all reports $\hat{\mathbf{p}}$ and wagers \mathbf{w} , the mechanism's output $(\mathbf{y}(\hat{\mathbf{p}}, \mathbf{w}), \mathbf{n}(\hat{\mathbf{p}}, \mathbf{w}), \boldsymbol{\pi}(\hat{\mathbf{p}}, \mathbf{w}))$ is such that there exists no profitable side bet.

This definition is difficult to work with directly. We show that there is an intuitive equivalent characterization of Pareto optimality in terms of allocations and costs: A mechanism is Pareto optimal if and only if there is some threshold price such that all agents with beliefs above the threshold spend their entire budget on yes securities while all agents with beliefs below the threshold spend their entire budget on no securities. This is formalized in the following theorem.

THEOREM 3.6. *A wagering mechanism is Pareto optimal if and only if for all reports $\hat{\mathbf{p}}$ and \mathbf{w} , there exists an agent $j \in \mathcal{N}$ such that*

$$\begin{aligned} \forall i : \hat{p}_i < \hat{p}_j, & \quad \pi_i(\hat{\mathbf{p}}, \mathbf{w}) = w_i \quad \text{and} \quad y_i(\hat{\mathbf{p}}, \mathbf{w}) = 0, \\ \forall i : \hat{p}_i > \hat{p}_j, & \quad \pi_i(\hat{\mathbf{p}}, \mathbf{w}) = w_i \quad \text{and} \quad n_i(\hat{\mathbf{p}}, \mathbf{w}) = 0. \end{aligned}$$

The first step of the proof, which appears in the appendix,² is to show that any time a profitable side bet exists, there is a profitable side bet with $\Delta\pi_i = 0$ for all agents i . This is because \$1 in cash is equivalent to a pair of yes and no securities. Thus we can limit attention to side bets that only involve the exchange of securities. The second step shows that, any time a profitable side bet exists, there exists a profitable side bet involving only two agents. The final step is to show that there is no profitable side bet between two agents if and only if the conditions in Theorem 3.6 hold.

Eisenberg and Gale [11] defined and analyzed the *parimutuel consensus mechanism* (PCM), a natural Pareto-optimal wagering mechanism. The outcome of the PCM is defined by a price p , such that all agents with $\hat{p}_i > p$ exhaust their entire wager buying yes securities at price p , and all agents with $\hat{p}_i < p$ exhaust their entire wager buying no securities at price $1 - p$. Any imbalance in demand for yes and no securities at price p is bridged by agents with report exactly p , who may buy either yes or no securities at the discretion of the mechanism. We can think of the PCM as a parimutuel mechanism with a proxy agent that switches agents' bets to the outcome most favorable to them, given the price. The PCM satisfies budget balance and individual rationality. However, the PCM does not satisfy incentive compatibility. An agent may affect the price p in a way that is favorable to them, as illustrated by the following example.

Example 3.7. Let $N = 2$, with $\mathbf{p} = (0.5, 0.75)$ and $\mathbf{w} = (1, 1)$. The outcome of the PCM with truthful reports is $(\mathbf{y} = (0, 2), \mathbf{n} = (2, 0), \boldsymbol{\pi} = (1, 1))$. Note that the price p is 0.5, so agent 1 achieves 0 utility. If agent 1 misreports $\hat{p}_1 = 2/3$, then the outcome becomes $(\mathbf{y} = (0, 1.5), \mathbf{n} = (1.5, 0), \boldsymbol{\pi} = (0.5, 1))$. Now, the price p is $2/3$, so agent 1 gets no securities at a price of $1/3$, gaining positive utility.

3.3 An Impossibility Result

We have shown that WSWM fails to produce Pareto optimal allocations and PCM fails to achieve incentive compatibility. In this section, we show that the tradeoff is unavoidable: no incentive compatible wagering mechanism can achieve Pareto optimality along with two other core properties.

The proof extends the intuition that for any two agents i and j with differing reports $\hat{p}_i < \hat{p}_j$, they must trade according to some intermediate price $p \in [\hat{p}_i, \hat{p}_j]$. It is therefore always in the interests of at least one of the agents to misreport her belief closer to that of the other agent, forcing the price further from her own true belief and thus achieving a higher payoff in expectation.

THEOREM 3.8. *No wagering mechanism simultaneously satisfies individual rationality, weak incentive compatibility, weak budget balance, and Pareto optimality. This holds even if the number of agents*

²An appendix containing all omitted proofs can be found in the full version of the paper on the authors' websites.

is arbitrarily large and all agents wager the same amount of money. Any three of the four properties are simultaneously attainable.

Individual rationality is hard to imagine giving up: We cannot force agents to participate. Weak incentive compatibility is key to ensuring the credibility of agents' reports. Although untruthful mechanisms like parimutuel wagering flourish in practice and do display an ability to aggregate useful information [1, 24], our goal is to create a mechanism that simplifies reasoning for the agents and principal and that offers some modicum of assurance that the reports the principal is seeing are accurate to the best abilities of the agents. Some wagering mechanisms, in particular automated market maker algorithms for prediction markets [7], do give up budget balance, subsidizing trade as a reward for information. However, most mechanisms seek profits if anything, not losses. When a subsidy is not possible or desired, we must relax Pareto optimality. In the remainder of this paper, we present and analyze our *double clinching auction* wagering mechanism which maintains individual rationality, weak incentive compatibility, and weak budget balance, while coming close to Pareto optimality in practice.

4 THE ADAPTIVE CLINCHING AUCTION

Since wagering mechanisms can be interpreted as allocating items (securities) to agents, it is natural to ask whether techniques from the auctions literature might be useful. Ausubel's clinching auction [2] produces VCG allocations and payments in the setting in which there are multiple identical items and each agent has a fixed valuation per item. However, VCG-style approaches cannot be applied in our setting since agents have budgets. Instead, we build on the *adaptive clinching auction* of Dobzinski et al. [9], which extends Ausubel's auction to handle budget constraints.

In this section, we review the adaptive clinching auction and state some known results that are used in our analysis. Many details are necessarily omitted. For a full description, we point the reader to Dobzinski et al. [9] and, for the divisible-items version, Bhattacharya et al. [4]. In describing the auction, we use notation that parallels that of the wagering mechanism setting, but the general description in this section is for arbitrary items.

Suppose that there are m identical, indivisible items for sale to a set of agents \mathcal{N} . Each agent i has a private value p_i for each item and a budget w_i , which we assume is known to the auctioneer.

The adaptive clinching auction is an ascending price auction. Each agent $i \in \mathcal{N}$ reports a bid \hat{p}_i . The price p per item starts at 0 and grows over time. Items are allocated as the price increases. As this happens, the auctioneer keeps track of the number of items $q_i(p)$ that have been allocated to each agent i at prices less than p along with the total cost $c_i(p)$ of those items and the agent's remaining budget $B_i(p) = w_i - c_i(p)$. Define the demand of agent i at price p to be

$$D_i(p) = \begin{cases} \infty & p = 0, \\ \lfloor \frac{B_i(p)}{p} \rfloor & 0 < p < \hat{p}_i, \\ 0 & p \geq \hat{p}_i \text{ and } p > 0. \end{cases} \quad (2)$$

The adaptive clinching auction allocates items to agent i at price p if the total demand of the other agents falls below the total supply. In particular, let $q(p) = m - \sum_{i \in \mathcal{N}} q_i(p)$ be the total number of items yet to be allocated. At any point, if $D_{-i}(p) = \sum_{j \neq i} D_j(p) < q(p)$, then $q(p) - D_{-i}(p)$ items are allocated to (or "clinched by") agent i at a price of p per item, and the relevant variables are updated accordingly.

The auction ends when the total demand no longer exceeds the total supply, that is, when $\sum_{i \in \mathcal{N}} D_i(p) \leq q(p)$. At this point, the price stops ascending and all agents with $D_i(p) > 0$ are allocated their full demand at a per-item price of p . If the total demand at price p is strictly less than the supply (i.e., $\sum_{i \in \mathcal{N}} D_i(p) < q(p)$), then the remaining $q(p) - \sum_{i \in \mathcal{N}} D_i(p)$ items are allocated to

agents i with $\hat{p}_i = p$. (We will see below that this is always possible to do.) A worked example is contained in the appendix.

The adaptive clinching auction can be extended to handle divisible items. While this extension is more complicated to write down, conceptually we simply view the auction as a continuous-time process. Bhattacharya et al. [4] give a formal description. We omit the details, but summarize the properties of the auction that we use to derive our results.³

First, agents have incentive to participate in the auction and to bid truthfully.

LEMMA 4.1 (DOBZINSKI ET AL. [9]). *The adaptive clinching auction for divisible items is individually rational. When budgets are known to the auctioneer, it is also incentive compatible: Every agent i maximizes expected utility by reporting $\hat{p}_i = p_i$.*

While Dobzinski et al. [9] only state incentive compatibility for the case of indivisible items, their proof carries through for the continuous version, and this fact is used heavily by Bhattacharya et al. [4]. It follows from the observation that the report \hat{p}_i only determines the price at which agent i drops out of the auction. While the price is below p_i , agent i can clinch (portions of) items at a per-item price below her value, thus deriving positive utility. After the price rises above p_i , any items she would clinch would cost more than her value, so she would derive negative utility. Thus, it is optimal to drop out of the auction exactly when the price reaches p_i .

We additionally use the fact that no agent is charged more than her budget.

LEMMA 4.2 (BHATTACHARYA ET AL. [4], DOBZINSKI ET AL. [9]). *The adaptive clinching auction for divisible items never charges an agent more than her budget.*

We also rely heavily on the following facts, which together imply that no agent (or the auctioneer) can be made better off without harming another agent.

LEMMA 4.3 (BHATTACHARYA ET AL. [4], DOBZINSKI ET AL. [9]). *The adaptive clinching auction for divisible items always allocates all m items.*

LEMMA 4.4 (BHATTACHARYA ET AL. [4], DOBZINSKI ET AL. [9]). *If an agent receives a non-zero allocation of items from the adaptive clinching auction for divisible items, then any player with a higher bid exhausts her entire budget.*

Finally, the utility of each agent is (weakly) increasing in the number of items sold.

LEMMA 4.5 (GOEL ET AL. [15]). *Fixing $\hat{\mathbf{p}}$ and \mathbf{w} , if $\hat{p}_i = p_i$ then i receives weakly greater expected utility from the adaptive clinching auction for divisible items when the number of items m increases.*

5 THE DOUBLE CLINCHING AUCTION

In this section, we present the double clinching auction. Motivated by the observation that existing incentive compatible wagering mechanisms do not even allocate securities efficiently, we turn to the adaptive clinching auction as a way to efficiently allocate any fixed number of securities. The principal runs two instances of the adaptive clinching auction for divisible items, deriving the agents' bids from their reports. The first instance, which we refer to as the yes auction, sells some number m^* of yes securities to the agents, fixing the bid of each agent i to equal her report \hat{p}_i . The second instance, which we refer to as the no auction, sells m^* no securities, fixing the bid of agent i to $1 - \hat{p}_i$. If m^* is chosen such that the payment collected for each pair of yes and no securities is at least \$1, then the principal never loses money, that is, the mechanism is weakly

³Both Dobzinski et al. [9] and Bhattacharya et al. [4] describe the divisible-items version in terms of a *single* divisible item. For our purposes, it is more convenient to view it as an auction over some number m of divisible items. This is equivalent and simply requires a rescaling of agent values.

budget balanced. While many values of m^* balance the budget, we define one particular value of m^* , carefully selected to ensure that agents cannot profit by misreporting their beliefs.

The primary technical contribution of this section is the derivation of m^* and the proof that the resulting auction is indeed (weakly) incentive compatible.

5.1 Definition of the Double Clinching Auction

To formally define the double clinching auction, we first describe the selection of m^* , the number of securities to be sold in each of the two instances of the adaptive clinching auction. We start by defining a pair of demand functions. These are similar to Equation (2), but do not take into account items that may have been allocated. Let D_i^y be the demand of agent i for (arbitrarily divisible) yes securities at price p assuming a per-item value of \hat{p}_i , and D_i^n her demand for no securities at price p assuming a per-item value of $1 - \hat{p}_i$, that is,

$$D_i^y(p) = \begin{cases} \infty & p = 0, \\ \frac{w_i}{p} & 0 < p < \hat{p}_i, \\ 0 & p \geq \hat{p}_i \text{ and } p > 0, \end{cases} \quad \text{and} \quad D_i^n(p) = \begin{cases} \infty & p = 0, \\ \frac{w_i}{p} & 0 < p < 1 - \hat{p}_i, \\ 0 & p \geq \hat{p}_i \text{ and } p > 0. \end{cases}$$

Let $D^y(p) = \sum_{i \in \mathcal{N}} D_i^y(p)$ be the total demand of all agents for yes securities at price p , and $D_{-i}^y(p) = \sum_{j \neq i} D_j^y(p)$ be the total excluding agent i . Define $D^n(p)$ and $D_{-i}^n(p)$ similarly.

The double clinching auction allocates securities only when there are 4 or more agents with positive wagers. (Agents with wagers of zero can simply be dropped since this is equivalent to not participating.) If there are fewer than 4, then no trade occurs. For the remainder of this section, assume that there are $N \geq 4$ agents who submit reports $\hat{p}_1 \leq \hat{p}_2 \leq \dots \leq \hat{p}_N$ and wagers $\mathbf{w} > \mathbf{0}$.

Fixing the number of securities m , define the *lowest clinching prices* as

$$c_y(m) = \begin{cases} \inf\{p : \min_{i \in \mathcal{N}} D_{-i}^y(p) < m\} & m > 0, \\ \hat{p}_{N-1} & m = 0, \end{cases}$$

and

$$c_n(m) = \begin{cases} \inf\{p : \min_{i \in \mathcal{N}} D_{-i}^n(p) < m\} & m > 0, \\ 1 - \hat{p}_2 & m = 0. \end{cases}$$

Here $c_y(m)$ can be thought of as the price at which the first (possibly infinitesimal) fraction of a security would be clinched in an adaptive clinching auction for m yes securities, and similarly, $c_n(m)$ the price at which the first fraction of a security would be clinched in an auction for m no securities. The $m = 0$ case is simply a technical definition that is required in our proofs. Both $c_y(m)$ and $c_n(m)$ are well-defined since they each take the infimum of a non-empty set that is bounded below by 0 since $D_{-i}^y(0) = D_{-i}^n(0) = \infty$ for all agents i . It is easy to see that for all m , $c_y(m) \in (0, \hat{p}_{N-1}]$ and $c_n(m) \in (0, 1 - \hat{p}_2]$. The following lemma gives additional useful properties of these functions. To show continuity, it is sufficient to show that the functions are surjective (onto), since a surjective, monotonic function is continuous.

LEMMA 5.1. *Fixing reports $\hat{\mathbf{p}}$ and wagers \mathbf{w} , c_y and c_n are continuous and weakly decreasing.*

Let $M = \{m : c_y(m) + c_n(m) > 1\}$. For any $m \in M$, auctioning off m yes and m no securities via two adaptive clinching auctions is guaranteed to collect more than m dollars total, or more than \$1 for each pair, guaranteeing no loss for the principal. We set m^* to be the largest m in M : the most pairs of securities such that every pair, even every fraction of a pair, costs more than \$1 per share

(i.e., every ϵ shares cost more than $\$ \epsilon$). Formally, the number of pairs of securities auctioned is

$$m^* = \begin{cases} \sup M & \hat{p}_2 < \hat{p}_{N-1}, \\ 0 & \hat{p}_2 = \hat{p}_{N-1}. \end{cases} \quad (3)$$

The following lemma guarantees that m^* is well-defined. This is clearly the case when $\hat{p}_2 = \hat{p}_{N-1}$. To show that m^* is well-defined when $\hat{p}_2 < \hat{p}_{N-1}$, it is sufficient to show that the set M is non-empty and bounded above, which implies the existence of a unique least upper bound. To show that $m^* > 0$ when $\hat{p}_{N-1} > \hat{p}_2$, we argue that $c_y(0) + c_n(0) > 1$, which implies there must exist some $m' > 0$ such that $c_y(m') + c_n(m') > 1$. This in turn implies that $m' \in M$ and therefore, $m^* = \sup M \geq m' > 0$.

LEMMA 5.2. *For any $\hat{\mathbf{p}}$ and \mathbf{w} , m^* is well-defined. Furthermore, $m^* > 0$ when $\hat{p}_{N-1} > \hat{p}_2$.*

With these definitions in place, we can formally define the double clinching auction; see Algorithm 1. The principal first sets m^* according to Equation 3. He then runs an auction for m^* yes securities (the yes auction) and an auction for m^* no securities (the no auction). A worked example of the double clinching auction on the reports from Example 3.1 is given in the appendix.

We have already shown that this procedure is well defined. However, to show that the double clinching auction is a valid wagering mechanism, we must also show that no agent ever loses more than her wager; that is, for any $\hat{\mathbf{p}}$ and \mathbf{w} , the double clinching auction produces output $(\mathbf{y}, \mathbf{n}, \boldsymbol{\pi})$ such that for all $i \in \mathcal{N}$, $\min\{y_i, n_i\} = 0$ and $\pi_i \leq w_i$. We show this in the following theorem.

Algorithm 1: The Double Clinching Auction. Here $\text{ClinchingAuction}(m, \hat{\mathbf{p}}, \mathbf{w})$ denotes the allocation and payments produced by an adaptive clinching auction for m arbitrarily divisible items on bids $\hat{\mathbf{p}}$ and budgets \mathbf{w} .

```

Input: reports  $\hat{\mathbf{p}}$  and wagers  $\mathbf{w} > \mathbf{0}$  of  $N$  agents
if  $N < 4$  or  $\hat{p}_2 = \hat{p}_{N-1}$  then
    Set  $(\mathbf{y}, \mathbf{n}, \boldsymbol{\pi}) = (\mathbf{0}, \mathbf{0}, \mathbf{0})$ 
else
    Set  $m^*$  as in Equation 3
    Let  $(\mathbf{y}, \boldsymbol{\pi}_y) = \text{ClinchingAuction}(m^*, \hat{\mathbf{p}}, \mathbf{w})$ 
    Let  $(\mathbf{n}, \boldsymbol{\pi}_n) = \text{ClinchingAuction}(m^*, \mathbf{1} - \hat{\mathbf{p}}, \mathbf{w})$ 
    Let  $\boldsymbol{\pi} = \boldsymbol{\pi}_y + \boldsymbol{\pi}_n$ 
end if
Output  $(\mathbf{y}, \mathbf{n}, \boldsymbol{\pi})$ 

```

THEOREM 5.3. *The double clinching auction is a valid wagering mechanism.*

From Lemma 4.2, we know that no agent can lose more than her wager in either the yes auction or the no auction alone. It is therefore sufficient to show that no agent is ever allocated a positive number of securities in both auctions. This follows immediately from the following lemma, taking p to be the report \hat{p}_i of any agent, and the definition of the clinching auction.

LEMMA 5.4. *Fixing any reports $\hat{\mathbf{p}}$ and wagers \mathbf{w} , for any $p \in [0, 1]$, either $\min_{i \in \mathcal{N}} D_{-i}^y(p) \geq m^*$, $\min_{i \in \mathcal{N}} D_{-i}^n(1-p) \geq m^*$, or both.*

PROOF. If $m^* = 0$ then this claim is trivially true, since for all p , $\min_{i \in \mathcal{N}} D_{-i}^y(p) \geq 0$ and $\min_{i \in \mathcal{N}} D_{-i}^n(p) \geq 0$. So suppose that $m^* > 0$. Suppose that $\min_{i \in \mathcal{N}} D_{-i}^y(p) < m^*$ and $\min_{i \in \mathcal{N}} D_{-i}^n(1-p) < m^*$. Then there exists an $m' < m^*$ such that $\min_{i \in \mathcal{N}} D_{-i}^y(p) < m'$ and $\min_{i \in \mathcal{N}} D_{-i}^n(1-p) < m'$.

Therefore, when m' securities are sold, clinching in the yes auction begins at (or before) p , and clinching in the no auction begins at (or before) $1 - p$. That is, $c_y(m') \leq p$ and $c_n(m') \leq 1 - p$. So $c_y(m') + c_n(m') \leq p + 1 - p = 1$. This implies that m' is a lower upper bound on the set $\{m : c_y(m) + c_n(m) > 1\}$ than m^* is, violating the definition of m^* . \square

5.2 Properties of the Double Clinching Auction

In this section, we discuss some desirable properties of the double clinching auction. We first observe that the double clinching auction is weakly budget balanced and individually rational.

PROPOSITION 5.5. *The double clinching auction is weakly budget balanced and individually rational.*

The proof of individual rationality invokes the incentive compatibility and individual rationality of the clinching mechanism. The proof of budget balance uses the fact that all m^* yes and no security pairs are allocated by the yes and no clinching auctions and the following lemma, which implies that each pair sells for more than \$1. The proof relies on the continuity and monotonicity of the functions c_y and c_n .

LEMMA 5.6. *For any reports \hat{p} and wagers \mathbf{w} , $c_y(m^*) + c_n(m^*) = 1$.*

Finally we state our main theoretical result: incentive compatibility of the double clinching auction. The proof is significantly more involved and we develop it in the next subsection.

THEOREM 5.7. *The double clinching auction is weakly incentive compatible.*

5.3 Proof of Incentive Compatibility

In this section, we prove Theorem 5.7, beginning with some useful lemmas. The first states that an agent cannot benefit from misreporting her belief unless it increases the number of securities.

LEMMA 5.8. *For any $i \in \mathcal{N}$, fix the wagers \mathbf{w} of all agents and reports \hat{p}_{-i} of all agents but i , and let $\hat{p}_i = p_i$. Agent i cannot increase her expected utility under the double clinching auction by reporting any $\hat{p}'_i \neq p_i$ unless this report increases the value of m^* .*

PROOF. Let m^* denote the number of security pairs allocated by the double clinching auction when i reports $\hat{p}_i = p_i$, and \hat{m}^* the number when i reports \hat{p}'_i .

First, observe that agent i cannot benefit from any misreport for which $m^* = \hat{m}^*$. This follows immediately from the incentive compatibility of the adaptive clinching auction (Lemma 4.1). Agent i maximizes the utility she gains from both the yes and no auctions individually when her bids in these auctions are truthful. Fixing m^* , the yes and no auctions are run independently, so agent i maximizes her total utility by reporting her true belief. Next, suppose that $\hat{m}^* < m^*$. Agent i 's utility for bidding untruthfully for \hat{m}^* securities is weakly less than her utility for bidding truthfully for \hat{m}^* securities, by incentive compatibility of the adaptive clinching auction, which is weakly less than her utility for bidding truthfully for m^* securities, by Lemma 4.5. Thus she would weakly prefer to bid truthfully than to make any misreport that reduces m^* . \square

Lemma 5.9 follows because, when any agent increases her report, the demand for yes (respectively, no) securities at a fixed price does not decrease (respectively, increase).

LEMMA 5.9. *For any $i \in \mathcal{N}$, fix the wagers \mathbf{w} of all agents and reports \hat{p}_{-i} of all agents but i . For any fixed m , as i 's report \hat{p}_i increases, the lowest clinching price $c_y(m)$ weakly increases, while the lowest clinching price $c_n(m)$ weakly decreases.*

The proof of Lemma 5.10 uses the fact that as the price moves from any value p to a sufficiently close higher value p' , no new agent will drop out of the auction, and so demand functions only change by a very small amount.

LEMMA 5.10. For any reports \hat{p} and wagers \mathbf{w} , and any m , $\min_{i \in N} D_{-i}^y(c_y(m)) \leq m$ and $\min_{i \in N} D_{-i}^n(c_n(m)) \leq m$.

We are now ready to complete the proof of Theorem 5.7. We start by observing that no agent can be allocated both yes and no securities. We treat two cases. If an agent's misreport does not change the type of security that she is allocated, then it cannot increase the number of securities sold, and so by Lemma 5.8, cannot be profitable. If her misreport does change the type of security that she is allocated, she may be able to increase the number of securities auctioned. However, in this case, the amount she pays would be higher than her value for the securities she gets.

PROOF OF THEOREM 5.7. Consider an agent i with belief p_i and let \hat{p}_{-i} denote the reports of all agents other than i . Let m^* be the number of pairs of securities auctioned if i truthfully reports p_i . Denote by \hat{c}_y and \hat{c}_n the lowest clinching price functions if i misreports $\hat{p}_i \neq p_i$, and by \hat{D} the demand functions in the misreported instance. Let \hat{m}^* denote the number of pairs of securities auctioned in the misreported instance.

Noting that $\hat{p}_i \neq p_i$, we can break the analysis into four cases:

- | | |
|--|--|
| (1) $p_i < c_y(m^*)$ and $\hat{p}_i \leq c_y(m^*)$ | (3) $p_i \leq c_y(m^*)$ and $\hat{p}_i > c_y(m^*)$ |
| (2) $p_i > c_y(m^*)$ and $\hat{p}_i \geq c_y(m^*)$ | (4) $p_i \geq c_y(m^*)$ and $\hat{p}_i < c_y(m^*)$ |

Case 1 and Case 2 are symmetric, since in Case 2 $1 - p_i < c_n(m^*)$ and $1 - \hat{p}_i \leq c_n(m^*)$, which is equivalent to Case 1 reversing the outcomes yes and no. Similarly, Case 3 and Case 4 are symmetric. Therefore, it is sufficient to show that i does not benefit from misreporting in Cases 1 or 3.

Case 1: $p_i < c_y(m^*)$ and $\hat{p}_i \leq c_y(m^*)$. To show that i can not benefit from this misreport, we prove that she does not change the clinching prices $c_y(m^*)$ and $c_n(m^*)$. We will show that if $\hat{p}_i < c_y(m^*)$ then this is true because the global demand can only change at prices between p_i and \hat{p}_i , and this interval does not contain $c_y(m^*)$. When $\hat{p}_i = c_y(m^*)$, some more care is necessary.

If $m^* = 0$ then $c_y(m^*) = p_{N-1}$. Further, we know that $p_{N-1} = p_2$, or else it would be the case that $m^* > 0$, by Lemma 5.2. And, since we have assumed that $p_i < c_y(m^*)$, we know that i 's report is the lowest of all agents. Since $\hat{p}_i \leq c_y(m^*) = p_{N-1} = p_2$, \hat{p}_i is still the (equal) lowest report, and therefore both the second highest and second lowest reports are unchanged. In particular, $\hat{p}_2 = \hat{p}_{N-1}$, so $\hat{m}^* = 0$. By Lemma 5.8, this misreport does not benefit i .

Now suppose that $m^* > 0$. We first show that $c_y(m^*) = \hat{c}_y(m^*)$ and $c_n(m^*) = \hat{c}_n(m^*)$. If $\hat{p}_i < c_y(m^*)$ then the demand locally around $c_y(m^*)$ and $c_n(m^*)$ is unchanged. Therefore, since $c_y(m^*)$ and $c_n(m^*)$ are the prices at which demand drops below m^* , these quantities remain unchanged in the misreported instance. If $\hat{p}_i = c_y(m^*)$ then, by Lemma 5.9, $c_y(m^*) \leq \hat{c}_y(m^*)$, since $p_i < \hat{p}_i = c_y(m^*)$. However, for all $p > c_y(m^*)$, the demand in the misreported instance is exactly the same as that in the truthful instance, and therefore $\min_i \hat{D}_{-i}^y(p) < m^*$ for all $p > c_y(m^*)$, which implies that $\hat{c}_y(m^*) \leq c_y(m^*)$. This, together with the earlier statement that $\hat{c}_y(m^*) \geq c_y(m^*)$, gives us $c_y(m^*) = \hat{c}_y(m^*)$. By similar reasoning, $c_n(m^*) = \hat{c}_n(m^*)$.

Therefore $\hat{c}_y(m^*) + \hat{c}_n(m^*) = c_y(m^*) + c_n(m^*) = 1$. Since \hat{c}_y and \hat{c}_n are decreasing functions, m^* is an upper bound on the set $M = \{m : \hat{c}_y(m) + \hat{c}_n(m) > 1\}$. Since the double clinching auction sells a number of securities equal to the least upper bound of M , it therefore sells at most m^* securities in the misreported instance. By Corollary 5.8, agent i does not profit from this misreport.

Case 3: $p_i \leq c_y(m^*)$ and $\hat{p}_i > c_y(m^*)$. In this case, i 's misreport can increase the number of securities sold. However, we show that to do so, i must be allocated some yes securities. But since i 's misreport is higher than her true value, it must also be the case that the price for yes securities is higher in the misreported instance than the truthful instance. Because all yes securities are sold at

a price higher than i 's valuation in the truthful instance, it must still be the case in the misreported instance. Therefore i does not get any positive utility from the securities she is allocated.

There are two possibilities. First is that $\hat{c}_y(m^*) + \hat{c}_n(m^*) \leq 1$, in which case we need to sell (weakly) fewer securities in the misreported instance than the truthful instance. That is, $\hat{m}^* \leq m^*$. By Lemma 5.8, the misreport can not be profitable for i in this case.

Second is that $\hat{c}_y(m^*) + \hat{c}_n(m^*) > 1$. In this case it may be possible to sell more securities, so assume that $\hat{m}^* > m^*$ (otherwise i 's misreport is not profitable, by Lemma 5.8). By Lemma 5.10, $\min_j \hat{D}_{-j}^y(\hat{c}_y(m^*)) \leq m^* < \hat{m}^*$. So, by Lemma 5.4,

$$\min_j \hat{D}_{-j}^n(1 - \hat{c}_y(m^*)) \geq \hat{m}^*. \quad (4)$$

In what remains of the proof, we show that holding the number of securities the same as in the truthful instance, i 's misreport cannot result in the clinching price rising all the way above \hat{p}_i . We can then use the fact that the clinching price decreases as we sell more securities to deduce that $\hat{p}_i \geq \hat{c}_y(\hat{m}^*)$, which (after addressing some details) says that i is allocated yes securities, and not no securities. By lower bounding the price of the yes securities by i 's true valuation p_i , this says that i can not derive positive utility from this misreport. We now prove this formally.

We first show that $\hat{p}_i \geq \hat{c}_y(m^*)$. In the case that $m^* = 0$, this is true because $\hat{p}_i > c_y(m^*) = p_{N-1}$, so therefore \hat{p}_i is one of the two highest reports in the misreported instance. And, since $m^* = 0$, it follows that $\hat{c}_y(m^*) = \hat{p}_{N-1} \leq \hat{p}_i$.

For the case that $m^* > 0$, note that the demand is unchanged from the truthful instance at all prices greater than or equal to \hat{p}_i . Therefore for all $p \geq \hat{p}_i$, we have that $\min_j \hat{D}_{-j}^y(p) = \min_j D_{-j}^y(p) < m^*$, where the inequality holds because $p \geq \hat{p}_i > c_y(m^*)$. In particular, $\min_j \hat{D}_{-j}^y(\hat{p}_i) < m^*$, which implies that $\hat{p}_i \geq \hat{c}_y(m^*)$.

From $\hat{p}_i \geq \hat{c}_y(m^*)$, we have that $1 - \hat{p}_i \leq 1 - \hat{c}_y(m^*)$, which implies that $\hat{D}_{-i}^n(1 - \hat{p}_i) \geq \hat{D}_{-i}^n(1 - \hat{c}_y(m^*))$. Combining this with Equation 4, $\hat{D}_{-i}^n(1 - \hat{p}_i) \geq \hat{D}_{-i}^n(1 - \hat{c}_y(m^*)) \geq \min_j \hat{D}_{-j}^n(1 - \hat{c}_y(m^*)) \geq \hat{m}^*$, which implies that i does not receive no securities in the misreported instance. There are two possibilities remaining. If i also does not receive yes securities, then agent i achieves zero overall payoff after misreporting, which is no better than her payoff from reporting truthfully. Otherwise, the average price paid per yes security is at least $\hat{c}_y(\hat{m}^*) = 1 - \hat{c}_n(\hat{m}^*) \geq 1 - \hat{c}_n(m^*) \geq 1 - c_n(m^*) = c_y(m^*) \geq p_i$, where the first inequality follows from the fact that \hat{c}_n is decreasing and $\hat{m}^* > m^*$, and the second inequality follows from Lemma 5.9, because $\hat{p}_i > p_i$. Therefore i is paying a price for the securities equal to or greater than they are worth to her, so she obtains non-positive expected payoff, which is no better than her (non-negative) truthful payoff. \square

5.4 Beyond Weak Incentive Compatibility

Theorem 5.7 proves weak incentive compatibility. Taken at face value, weak incentive compatibility is, well, extremely weak. Indeed, simply paying each agent a constant amount regardless of her report satisfies weak incentive compatibility.

We show that the double clinching auction actually satisfies a stronger property: If agent i makes any misreport $\hat{p}_i \neq p_i$, then, for some set of reports \hat{p}_{-i} of the other agents, agent i obtains strictly lower expected utility than she would by reporting truthfully. If agent i is sufficiently uncertain about other agents, she is strictly better off reporting her true belief.

THEOREM 5.11. *Fix any set of agents \mathcal{N} with $N \geq 4$ and any wagers w . For any agent i with belief p_i and any report $\hat{p}_i \neq p_i$, there exist reports \hat{p}_{-i} of the other agents such that under the double*

clinching auction

$$p_i y_i((p_i, \hat{p}_{-i}), \mathbf{w}) + (1 - p_i) n_i((p_i, \hat{p}_{-i}), \mathbf{w}) - \pi_i((p_i, \hat{p}_{-i}), \mathbf{w}) \\ > p_i y_i(\hat{p}, \mathbf{w}) + (1 - p_i) n_i(\hat{p}, \mathbf{w}) - \pi_i(\hat{p}, \mathbf{w}).$$

5.5 Budget Inflation Under the Double Clinching Auction

As discussed in Section 3, even a wagering mechanism that satisfies incentive compatibility may give an agent incentive to misreport her belief if she can safely inflate her budget. Since the double clinching auction is not Pareto optimal, a bidder with complete knowledge of the reports and wagers of other agents could have incentive to inflate her budget. Further, there exist examples analogous to Example 3.3 where the potential for an agent to inflate her budget may also affect her incentive to report truthfully.

However, in reality agents operate with only limited knowledge about the reports of other agents. While the budget misreport in Example 3.3 was safe in the sense that the budget inflation could not lead to the misreporting agent overspending her true budget, we can show that completely safe manipulations are not possible under the double clinching auction. An agent cannot inflate her wager without at least some risk of losing more than her true budget. This is in stark contrast to Proposition 3.2 for the WSWM.

THEOREM 5.12. *Fix any set of agents \mathcal{N} with $N \geq 4$. For any agent i with report \hat{p}_i and wager w_i , there exist reports \hat{p}_{-i} and wagers \mathbf{w}_{-i} of the other agents such that $\pi_i = w_i$ under the double clinching auction.*

6 SIMULATIONS

For a fixed number of yes securities, the adaptive clinching auction is efficient, so we had reason to suspect that the double clinching auction, selling m^* yes and no securities, would be near efficient. In this section, in a series of simulations based on real probability reports, we show that indeed the DCA is much more efficient than the WSWM or the NAWM, in some cases coming remarkably close to Pareto optimality.

We compare the performance of the double clinching auction to the parimutuel consensus mechanism (PCM), the Brier scoring rule version of the weighted score wagering mechanism (WSWM), and the Brier-score no-arbitrage wagering mechanism (NAWM). The PCM is known to be Pareto optimal, serving as the gold standard with respect to the amount of trade generated, though is not incentive compatible. WSWMs and NAWMs provide a natural comparison as the only other known, non-trivial wagering mechanisms that are individually rational, incentive compatible, and budget balanced. We chose the Brier scoring rule since it is commonly used in practice.

We tested each wagering mechanism on a data set consisting of probability reports collected from an online prediction contest called ProbabilitySports [12].^{4,5} The data set consists of probabilistic predictions about the outcomes of 1643 U.S. National Football League games between the start of the 2000 preseason and the end of the 2004 season. For each match, between 64 and 1574 people reported their subjective probability of the home team winning the game. After each game, they earned points in the contest according to a Brier scoring rule.

The ProbabilitySports users provided probabilities but not wager amounts. We simulate wagers in two ways. First, we generate uniform wagers: we fix all wagers at 1, modeling a scenario where

⁴We thank Brian Galebach for providing us with this data.

⁵We also conducted simulations with random probability reports, generated both uniformly at random and according to a beta distribution. The results, and in particular the relative performance of the mechanisms, are very similar to those obtained using the ProbabilitySports data set.

agents are equal or cannot vary the default wager amount. Second, we generate wagers according to a Pareto distribution, reflecting the typical distribution of wealth in a population. We used a Pareto distribution with shape parameter 1.16 and scale parameter 1, which is often described as “20% of the population has 80% of the wealth.” Each random set of wagers was scaled so that the average wager for any single match is exactly 1, allowing a comparison to the uniform wager case.

6.1 Notes on Implementation

A perfectly faithful implementation of the double clinching auction, as defined in Section 5, would require running an adaptive clinching auction for arbitrarily divisible goods with continuously increasing price function and allocations. In practice, it is necessary to discretize the price increases, thus computing allocations and prices that approximately match the double clinching auction.

One might be concerned that this discretization could adversely affect the nice properties of the double clinching auction. In particular, it might now be possible for an agent to profit by misreporting her probability. To check whether this was the case, we empirically tested incentive compatibility on the 1643 matches from the ProbabilitySports data set. For each match, we chose a random agent i and ran the double clinching auction 101 times to calculate the expected payoff i would receive reporting each value in the set $\{0, 0.01, 0.02, \dots, 0.99, 1\}$. We found a single profitable misreport for only a single one of these matches, with the misreporting agent able to increase her expected utility from 5.1611 to 5.1612. This suggests that the mechanism retains (at least approximate) incentive compatibility when discretized.

6.2 Results

The results are summarized in Table 1. The top table shows various statistics averaged across all 1643 matches, with wagers for each match drawn from a Pareto distribution. The Risk/Wagers column reports the total risk, summed across all agents, divided by the total wager, summed across all agents, or $\sum_{i \in N} \pi_i / \sum_{i \in N} w_i$. A value of 1 means that every agent risks losing her entire wager for one outcome; a value of 0 means that no trade occurs. The %Full Stakes column reports the percentage of agents that risk losing their entire wager under one outcome (i.e., $\pi_i = w_i$). The #Securities column gives the total number of pairs of securities sold to the agents, or $\sum_{i \in N} y_i = \sum_{i \in N} n_i$. The Principal Profit column shows the principal’s net profit. Finally, the Agent Utility column gives the sum of the agents’ expected utilities, assuming immutable beliefs and truthful reports.

As expected, the PCM facilitates the most trade, in terms of both the risk:budget ratio and the number of securities sold. However, there is a notably large gap in these metrics between the double clinching auction and the NAWM and WSWM, with the double clinching auction selling almost five times as many securities as the NAWM and WSWM. Additionally, under the double clinching auction, over 80% of agents risk their entire wagers, compared with no agents under NAWM and WSWM. This is further evidence that falsely inflating a wager amount under the double clinching auction, while possibly beneficial in theory, would be extremely risky in practice, with a high chance of the manipulating agent losing more than her true budget.

We note that, while our objective is not to make a profit for the principal, the double clinching auction does yield a reasonable profit without sacrificing agent welfare.

The bottom table reports the same metrics when the agents’ wagers are equal. While the three other mechanisms exhibit very similar performance in this case, the double clinching auction displays a marked increase in the amount of trade facilitated, under all metrics, and a drop in profit. For the objective of maximizing trade, this is a particularly compelling argument to use the double clinching auction in cases when equal wagers are natural.

Note that all matches in the ProbabilitySports data set have a relatively large number of agents participating. However, in many cases we are interested in instances with smaller numbers of agents.

	Risk/Wagers	%Full Stakes	#Securities	Principal Profit	Agent Utility
PCM	0.94	97.2	601.1	0	200.1
DCA	0.80	82.6	489.3	28.4	152.1
NAWM	0.20	0	98.4	25.8	27.9
WSWM	0.16	0	101.2	0	53.7

	Risk/Wagers	%Full Stakes	#Securities	Principal Profit	Agent Utility
PCM	0.97	96.9	618.5	0	208.1
DCA	0.97	96.4	616.4	1.3	206.3
NAWM	0.21	0	102.5	28.7	28.8
WSWM	0.17	0	105.7	0	57.4

Table 1. Evaluation of wagering mechanisms using reports from the ProbabilitySports data set. The top table was generated using Pareto distribution wagers, the bottom with uniform wagers.

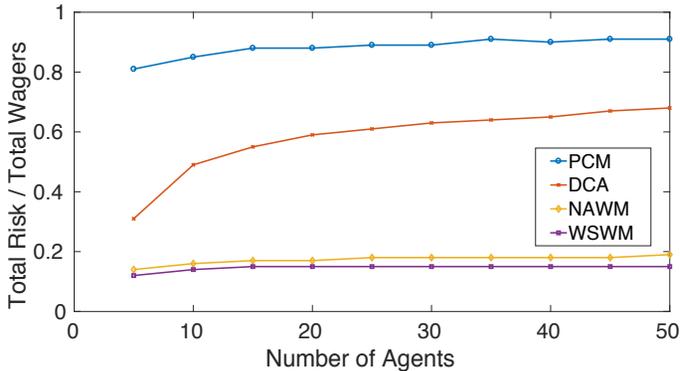


Fig. 1. The risk:budget ratio for each of the four wagering mechanisms, plotted as a function of N .

To investigate this behavior, we generated smaller instances by subsampling reports from the full set of reports for each match. Figure 1 plots the ratio of total risk to total budget, $(\sum_{i \in N} \pi_i) / (\sum_{i \in N} w_i)$, for the four mechanisms for values of N ranging from 5 to 50, with wagers randomly drawn from a Pareto distribution. We see that while the PCM, NAWM, and WSWM exhibit only minimal change as N increases, the double clinching auction facilitates more trade for larger values of N . However, even for $N = 5$, the double clinching auction facilitates approximately twice the trade as the WSWM and the NAWM, suggesting that the double clinching auction is the best truthful mechanism when maximizing trade is a primary objective.

7 CONCLUSION

We have defined and analyzed the double clinching auction, proving that it simultaneously satisfies incentive compatibility, budget balance, and individual rationality. While we showed that no wagering mechanism can simultaneously achieve these three properties along with Pareto optimality, our simulations suggest that the DCA comes close to Pareto optimality in practice, making it the first known incentive compatible wagering mechanism to do so.

It would be valuable, but apparently non-trivial, to extend the DCA to settings with non-binary outcomes. The DCA crucially exploits the fact that agents can be ordered by their reports in one dimension, allowing us to guarantee that no agent is allocated both yes and no securities. With

larger outcome spaces, this property no longer holds, and designing a mechanism in which the principal auctions off three or more types of securities would require novel techniques.

Even in the binary-outcome setting, a number of interesting problems remain. While our simulations suggest that the DCA comes close to achieving Pareto optimality, we have not established any formal approximation guarantee. An additional particularly compelling question is whether our choice of m^* is the largest number of securities that can be sold via a pair of adaptive clinching auctions while satisfying incentive compatibility and budget balance.

REFERENCES

- [1] Mukhtar M. Ali. 1977. Probability and utility estimates for racetrack bettors. *J. Political Economy* 85, 4 (1977), 803–816.
- [2] Lawrence M Ausubel. 2004. An efficient ascending-bid auction for multiple objects. *The American Economic Review* 94, 5 (2004), 1452–1475.
- [3] Alina Beygelzimer, John Langford, and David M. Pennock. 2012. Learning performance of prediction markets with Kelly bettors. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- [4] Sayan Bhattacharya, Vincent Conitzer, Kamesh Munagala, and Lirong Xia. 2010. Incentive compatible budget elicitation in multi-unit auctions. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- [5] Glenn W. Brier. 1950. Verification of forecasts expressed in terms of probability. *Monthly Weather Review* 78, 1 (1950), 1–3.
- [6] Yiling Chen, Nikhil R. Devanur, David M. Pennock, and Jennifer Wortman Vaughan. 2014. Removing arbitrage from wagering mechanisms. In *ACM Conference on Economics and Computation (EC)*.
- [7] Yiling Chen and David M. Pennock. 2007. A utility framework for bounded-loss market makers. In *Conference on Uncertainty in Artificial Intelligence (UAI)*.
- [8] Rachel Cummings, David M. Pennock, and Jennifer Wortman Vaughan. 2016. The possibilities and limitations of private prediction markets. In *ACM Conference on Economics and Computation (EC)*.
- [9] Shahar Dobzinski, Ron Lavi, and Noam Nisan. 2012. Multi-unit auctions with budget limits. *Games and Economic Behavior* 74, 2 (2012), 486–503.
- [10] Pradeep Dubey, John Geanakoplos, and Martin Shubik. 1987. The revelation of information in strategic market games: A critique of rational expectations equilibrium. *Journal of Mathematical Economics* 16 (1987), 3105–3137.
- [11] Edmund Eisenberg and David Gale. 1959. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics* 30, 1 (1959), 165–168.
- [12] Brian Galebach. 2017. ProbabilitySports. <http://probabilitiesports.com/>. (2017).
- [13] Christian Genest and James V. Zidek. 1986. Combining probability distributions: A critique and an annotated bibliography. *Statist. Sci.* 1, 1 (1986), 114–148.
- [14] Tilmann Gneiting and Adrian E. Raftery. 2007. Strictly proper scoring rules, prediction, and estimation. *J. Amer. Statist. Assoc.* 102, 477 (2007), 359–378.
- [15] Gagan Goel, Vahab Mirrokni, and Renato Paes Leme. 2015. Clinching auctions with online supply. *Games and Economic Behavior* (2015).
- [16] Robert A. Jacobs. 1995. Methods for combining experts' probability assessments. *Neural Computation* 7, 5 (1995), 867–888.
- [17] David J. Johnstone. 2007. The Parimutuel Kelly Probability Scoring Rule. *Decision Analysis* 4, 2 (2007), 66–75.
- [18] James Jordan and R. Radner. 1979. The nonexistence of Rational expectations equilibrium: A robust example. (1979). Working paper, Department of Economics, University of Minnesota.
- [19] D. Marc Kilgour and Yigal Gerchak. 2004. Elicitation of Probabilities Using Competitive Scoring Rules. *Decision Analysis* 1, 2 (2004), 108–113.
- [20] Nicolas S. Lambert, John Langford, Jennifer Wortman Vaughan, Yiling Chen, Daniel Reeves, Yoav Shoham, and David M. Pennock. 2015. An axiomatic characterization of wagering mechanisms. *Journal of Economic Theory* 156 (2015), 389–416.
- [21] Nicolas S. Lambert, John Langford, Jennifer Wortman, Yiling Chen, Daniel Reeves, Yoav Shoham, and David M. Pennock. 2008. Self-financed wagering mechanisms for forecasting. In *ACM Conference on Electronic Commerce (EC)*.
- [22] Peter A. Morris. 1977. Combining expert judgments: A Bayesian approach. *Management Science* 23, 7 (1977), 679–693.
- [23] David M. Pennock, Vasilis Syrgkanis, and Jennifer Wortman Vaughan. 2016. Bounded rationality in wagering mechanisms. In *Conference on Uncertainty in Artificial Intelligence (UAI)*.
- [24] Charles R. Plott, Jorgen Wit, and Winston C. Yang. 1997. *Parimutuel betting markets as information aggregation devices: Experimental results*. Technical Report Social Science Working Paper 986. California Institute of Technology.
- [25] Leonard J. Savage. 1971. Elicitation of personal probabilities and expectations. *J. Amer. Statist. Assoc.* 66, 336 (1971), 783–801.