

The Possibilities and Limitations of Private Prediction Markets

RACHEL CUMMINGS, California Institute of Technology
DAVID M. PENNOCK, Microsoft Research
JENNIFER WORTMAN VAUGHAN, Microsoft Research

We consider the design of *private prediction markets*, financial markets designed to elicit predictions about uncertain events without revealing too much information about market participants' actions or beliefs. Our goal is to design market mechanisms in which participants' trades or wagers influence the market's behavior in a way that leads to accurate predictions, yet no single participant has too much influence over what others are able to observe. We study the possibilities and limitations of such mechanisms using tools from differential privacy. We begin by designing a private one-shot wagering mechanism in which bettors specify a belief about the likelihood of a future event and a corresponding monetary wager. Wagers are redistributed among bettors in a way that more highly rewards those with accurate predictions. We provide a class of wagering mechanisms that are guaranteed to satisfy truthfulness, budget balance on expectation, and other desirable properties while additionally guaranteeing ϵ -joint differential privacy in the bettors' reported beliefs, and analyze the trade-off between the achievable level of privacy and the sensitivity of a bettor's payment to her own report. We then ask whether it is possible to obtain privacy in dynamic prediction markets, focusing our attention on the popular cost-function framework in which securities with payments linked to future events are bought and sold by an automated market maker. We show that under general conditions, it is impossible for such a market maker to simultaneously achieve bounded worst-case loss and ϵ -differential privacy without allowing the privacy guarantee to degrade extremely quickly as the number of trades grows, making such markets impractical in settings in which privacy is valued. We conclude by suggesting several avenues for potentially circumventing this lower bound.

1. INTRODUCTION

Betting markets of various forms—including the stock exchange [Grossman 1976], futures markets [Roll 1984], sports betting markets [Gandar et al. 1999], and markets at the racetrack [Thaler and Ziemba 1988]—have been shown to successfully collect and aggregate information. Over the last few decades, *prediction markets* designed specifically for the purpose of elicitation and aggregation, have yielded useful predictions in domains as diverse as politics [Berg et al. 2001], disease surveillance [Polgreen et al. 2007], and entertainment [Pennock et al. 2002].

The desire to aggregate and act on the strategically valuable information dispersed among employees has led many companies to experiment with internal prediction markets. An internal corporate market could be used to predict the launch date of a new product or the product's eventual success. Among the first companies to experiment with internal markets was Hewlett-Packard, which endowed each trader with a small budget of real money [Plott and Chen 2002]. Microsoft [Berg and Proebsting 2009] and Google [Cowgill and Zitzewitz 2015] began experimental markets using their own in-

A full version of this paper which includes the appendix is available online at <http://arxiv.org/abs/1602.07362>. The first author was supported in part by a Simons Award for Graduate Students in Theoretical Computer Science, NSF grant CNS-1254169, and US-Israel Binational Science Foundation grant 2012348. Much of this research was done while R. Cummings was at Microsoft Research.

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 ACM 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'16, July 24–28, 2016, Maastricht, The Netherlands.

ACM 978-1-4503-3936-0/16/07 ...\$15.00.

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

<http://dx.doi.org/10.1145/2940716.2940721>

ternal currencies in 2003 and 2005, respectively. Intel, Ford, GE, Siemens, and others have engaged in similar experiments [Charette 2007].

Proponents of internal corporate markets often argue that the market structure helps in part because, without it, “business practices... create incentives for individuals not to reveal their information” [Plott and Chen 2002]. However, even with a formal market structure in place, an employee might be hesitant to bet against the success of their team for fear of insulting her coworkers or angering management. If an employee has information that is unfavorable to the company, she might choose not to report it, leading to predictions that are overly optimistic for the company and ultimately contributing to an “optimism bias” in the market similar to the bias in Google’s corporate markets discovered by Cowgill and Zitzewitz [2015].

To address this issue, we consider the problem of designing *private* prediction markets. A private market would allow participants to engage in the market and contribute to the accuracy of the market’s predictions without fear of having their information or beliefs revealed. The goal is to provide participants with a form of “plausible deniability.” Although participants’ trades or wagers should together influence the market’s behavior and predictions, no single participant’s actions should have too much influence over what others can observe. We formalize this idea using the popular notion of *differential privacy* [Dwork et al. 2006; Dwork and Roth 2014], which can be used to guarantee that any participant’s actions cannot be inferred from observations.

We begin by designing a private analog of the *weighted score wagering mechanisms* first introduced by Lambert et al. [2008]. A *wagering mechanism* allows bettors to each specify a belief about the likelihood of a future event and a corresponding monetary wager. These wagers are then collected by a centralized operator and redistributed among bettors in such a way that more accurate bettors receive higher rewards. Lambert et al. [2008] showed that the class of weighted score wagering mechanisms, which are built on the machinery of proper scoring rules [Gneiting and Raftery 2007], is the unique set of wagering mechanisms to satisfy a set of desired properties such as budget balance, truthfulness, and anonymity. We design a class of wagering mechanisms with randomized payments that maintain the nice properties of weighted score wagering mechanisms in expectation while additionally guaranteeing ϵ -joint differential privacy in the bettors’ reported beliefs. We discuss the trade-offs that exist between the privacy of the mechanism (captured by the parameter ϵ) and the sensitivity of a bettor’s payment to her own report, and show how to set the parameters of our mechanisms to achieve a reasonable level of the plausible deniability desired in practice.

We next address the problem of running private dynamic prediction markets. We consider the setting in which traders buy and sell securities with values linked to future events. For example, a market might offer a security worth \$1 if Microsoft Bing’s market share increases in 2016 and \$0 otherwise. A risk neutral trader who believes that the probability of Bing’s market share increasing is p would profit from buying this security at any price less than $\$p$ or (short) selling it at any price greater than $\$p$. The market price of the security is thought to reflect traders’ collective beliefs about the likelihood of this event. We focus on *cost-function prediction markets* [Abernethy et al. 2013; Chen and Pennock 2007] such as Hanson’s popular logarithmic market scoring rule [Hanson 2003]. In a cost-function market, all trades are placed through an *automated market maker*, a centralized algorithmic agent that is always willing to buy or sell securities at some current market price that depends on the history of trade via a potential function called the cost function. We ask whether it is possible for a market maker to price trades according to a noisy cost function in a way that maintains traders’ privacy without allowing traders to make unbounded profit off of the noise. Unfortunately, we show that under general assumptions, it is impossible for a market maker to achieve bounded loss and ϵ -differential privacy without allowing the

privacy guarantee to degrade very quickly as the number of trades grows. In particular, the quantity e^ϵ must grow faster than linearly in the number of trades, making such markets impractical in settings in which privacy is valued. We suggest several avenues for future research aimed at circumventing this lower bound.

There is very little prior work on the design of private prediction markets, and to the best of our knowledge, we are the first to consider privacy for one-shot wagering mechanisms. Most closely related to our work is the recent paper of Waggoner et al. [2015] who consider a setting in which each of a set of self-interested agents holds a private data point consisting of an observation x and corresponding label y . A firm would like to purchase this data in order to learn a function to accurately predict the labels of new observations. Waggoner et al. propose a mechanism that provides incentives for the agents to reveal their data in such a way that the firm is able to solve its prediction task while maintaining the agents' privacy. The authors mention that similar ideas can be applied to produce privacy-preserving prediction markets, but their construction requires knowing the number of trades that will occur in advance to set parameters. The most straightforward way of applying their techniques to prediction markets results in a market maker falling in the class covered by our impossibility result, suggesting that such techniques cannot be used to derive a privacy-preserving market with bounded loss when the number trades is not bounded.

2. TOOLS FROM DIFFERENTIAL PRIVACY

We formalize privacy using the now-standard notion of *differential privacy*, which was introduced by Dwork et al. [2006]. The most basic version of differential privacy is used to measure the privacy of a randomized algorithm's output when given as input a database D with n entries from some input domain \mathcal{I} . Differential privacy is often studied in settings in which the n entries are provided by n agents, each of whom would like to keep their entry private. Two databases D and D' are said to be *neighboring* if they differ only in a single entry. Differential privacy requires that the distribution of the algorithm's output given D is close to the distribution of its output given any neighboring database D' . For these definitions, it is enough to view an algorithm as a randomized function mapping inputs to outputs; we are not concerned with the precise way in which the outputs are computed. In the following definitions, we restrict to real-valued outputs for consistency with the algorithms used in this paper.

DEFINITION 1 (DIFFERENTIAL PRIVACY [DWORK ET AL. 2006]). *For any $\epsilon, \delta \geq 0$, an algorithm $\mathcal{M} : \mathcal{I}^n \rightarrow \mathbb{R}$ is (ϵ, δ) -differentially private if for every pair of neighboring databases $D, D' \in \mathcal{I}^n$ and every subset $S \subseteq \mathbb{R}$,*

$$\Pr[\mathcal{M}(D) \in S] \leq e^\epsilon \Pr[\mathcal{M}(D') \in S] + \delta.$$

If $\delta = 0$, we say that \mathcal{M} is ϵ -differentially private.

As ϵ approaches 0, it becomes increasingly more difficult to distinguish neighboring databases, leading to a higher level of privacy. As ϵ grows large, the privacy guarantee grows increasingly weak. There is generally no consensus about what constitutes a good value of ϵ , and the strength of the guarantee needed may depend on the application. We discuss this point more later in the context of our results.

We sometimes abuse terminology and say that a random variable (such as a bettor's profit) is differentially private. This should be taken to mean that the algorithm used to compute the value of the random variable is differentially private.

In the context of mechanism design, differential privacy is often too strong of a notion. Suppose, for example, that the algorithm \mathcal{M} outputs a vector of prices that each of n agents will pay based on their joint input. While we may want the price that agent i pays to be differentially private in the input of the *other* agents, it is natural to allow

it to be more sensitive to changes in i 's own input. To capture this idea, Kearns et al. [2014] defined the notion of *joint differential privacy*. Call two neighboring databases D and D' i -neighbors if they differ only in the i th entry. Suppose that \mathcal{M} now outputs one element for each agent $i \in \{1, \dots, n\}$ and let $\mathcal{M}(D)_{-i}$ denote the vector of outputs to all agents excluding agent i . Then joint differential privacy is defined as follows.

DEFINITION 2 (JOINT DIFFERENTIAL PRIVACY [KEARNS ET AL. 2014]). *For any $\epsilon, \delta \geq 0$, an algorithm $\mathcal{M} : \mathcal{I}^n \rightarrow \mathbb{R}^n$ is (ϵ, δ) -joint differentially private if for every $i \in \{1, \dots, n\}$, for every pair of i -neighbors $D, D' \in \mathcal{I}^n$, and for every subset $S \subseteq \mathbb{R}^{n-1}$,*

$$\Pr[\mathcal{M}(D)_{-i} \in S] \leq e^\epsilon \Pr[\mathcal{M}(D')_{-i} \in S] + \delta.$$

If $\delta = 0$, we say that \mathcal{M} is ϵ -joint differentially private.

Joint differential privacy is still a strong requirement. It protects the privacy of any agent i from arbitrary coalitions; even if all other agents shared their private output, they would still not be able to learn too much about the input of agent i .

One useful tool for proving joint differential privacy is the *billboard lemma* [Hsu et al. 2014]. The idea behind the billboard lemma is quite intuitive and simple. Imagine that we display some message publicly so that it is viewable by all n agents, as if posted on a billboard, and suppose that the algorithm to compute this message is ϵ -differentially private. If each agent i 's output $\mathcal{M}(D)_i$ is computable from this public message along with i 's own private input, then \mathcal{M} is ϵ -joint differentially private. A more formal statement and proof are given in Hsu et al. [2014].

The definitions above assume the input database D is fixed. Differential privacy has also been considered for streaming algorithms [Chan et al. 2011; Dwork et al. 2010]. Let $\mathbb{N} = \{1, 2, 3, \dots\}$. Following Chan et al. [2011], a stream $\sigma \in \mathcal{I}^{\mathbb{N}}$ is a string of countable length of elements in \mathcal{I} , where $\sigma_t \in \mathcal{I}$ denotes the element at position or *time* t and $\sigma_{1, \dots, t} \in \mathcal{I}^t$ is the length t prefix of the stream σ . Two streams σ and σ' are said to be *neighbors* if they differ at exactly one time t .

A streaming algorithm \mathcal{M} is said to be *unbounded* if it accepts streams of indefinite length, that is, if for any stream $\sigma \in \mathcal{I}^{\mathbb{N}}$, $\mathcal{M}(\sigma) \in \mathbb{R}^{\mathbb{N}}$. In contrast, a streaming algorithm is T -bounded if it accepts only streams of length at most T . Dwork et al. [2010] consider only T -bounded streaming algorithms. Since we consider unbounded streaming algorithms, we use a more appropriate definition of differential privacy for streams adapted from Chan et al. [2011]. For unbounded streaming algorithms, it can be convenient to let the privacy guarantee degrade as the input stream grows in length. Chan et al. [2011] implicitly allow this in some of their results; see, for example, Corollary 4.5 in their paper. For clarity and preciseness, we explicitly capture this in our definition. Here and throughout the paper we use \mathbb{R}_+ to denote the nonnegative reals.

DEFINITION 3 (DIFFERENTIAL PRIVACY FOR STREAMS). *For any non-decreasing function $\epsilon : \mathbb{N} \rightarrow \mathbb{R}_+$ and any $\delta \geq 0$, a streaming algorithm $\mathcal{M} : \mathcal{I}^{\mathbb{N}} \rightarrow \mathbb{R}^{\mathbb{N}}$ is $(\epsilon(t), \delta)$ -differentially private if for every pair of neighboring streams $\sigma, \sigma' \in \mathcal{I}^{\mathbb{N}}$, for every $t \in \mathbb{N}$, and for every subset $S \subseteq \mathbb{R}^t$,*

$$\Pr[\mathcal{M}(\sigma_{1, \dots, t}) \in S] \leq e^{\epsilon(t)} \Pr[\mathcal{M}(\sigma'_{1, \dots, t}) \in S] + \delta.$$

If $\delta = 0$, we say that \mathcal{M} is $\epsilon(t)$ -differentially private.

Note that we allow ϵ to grow with t , but require that δ stay constant. In principle, one could also allow δ to depend on the length of the stream. However, allowing δ to increase would likely be unacceptable in scenarios in which privacy is considered important. In fact, it is more typical to require *smaller* values of δ for larger databases since for a database of size n , an algorithm could be considered (ϵ, δ) -private for δ on

the order of $1/n$ even if it fully reveals a small number of randomly chosen database entries [Dwork and Roth 2014]. Since we use this definition only when showing an impossibility result, allowing δ to decrease in t would not strengthen our result.

3. PRIVATE WAGERING MECHANISMS

We begin with the problem of designing a one-shot wagering mechanism that incentivizes bettors to truthfully report their beliefs while maintaining their privacy. A wagering mechanism allows a set of bettors to each specify a belief about a future event and a monetary wager. Wagers are collected by a centralized operator and redistributed to bettors in such a way that bettors with more accurate predictions are more highly rewarded. Lambert et al. [2008] showed that the class of *weighted score wagering mechanisms* (WSWMs) is the unique class of wagering mechanisms to satisfy a set of desired axioms such as budget balance and truthfulness. In this section, we show how to design a randomized wagering mechanism that achieves ϵ -joint differential privacy while maintaining the nice properties of WSWMs in expectation.

3.1. Standard wagering mechanisms

Wagering mechanisms, introduced by Lambert et al. [2008], are mechanisms designed to allow a centralized operator to elicit the beliefs of a set of bettors without taking on any risk. In this paper we focus on *binary* wagering mechanisms, in which each bettor i submits a report $p_i \in [0, 1]$ specifying how likely she believes it is that a particular event will occur, along with a wager $m_i \geq 0$ specifying the maximum amount of money that she is willing to lose. After all reports and wagers have been collected, all parties observe the realized outcome $\omega \in \{0, 1\}$ indicating whether or not the event occurred. Each bettor i then receives a payment that is a function of the outcome and the reports and wagers of all bettors. This idea is formalized as follows.

DEFINITION 4 (WAGERING MECHANISM [LAMBERT ET AL. 2008]). A wagering mechanism for a set of bettors $\mathcal{N} = \{1, \dots, n\}$ is specified by a vector Π of (possibly randomized) profit functions, $\Pi_i : [0, 1]^n \times \mathbb{R}_+^n \times \{0, 1\} \rightarrow \mathbb{R}$, where $\Pi_i(\mathbf{p}, \mathbf{m}, \omega)$ denotes the total profit to bettor i when the vectors of bettors' reported probabilities and wagers are \mathbf{p} and \mathbf{m} and the realized outcome is ω . It is required that $\Pi_i(\mathbf{p}, \mathbf{m}, \omega) \geq -m_i$ for all \mathbf{p}, \mathbf{m} , and ω , which ensures that no bettor loses more than her wager.

There are two minor differences between the definition presented here and that of Lambert et al. [2008]. First, for convenience, we use Π_i to denote the *total* profit to bettor i (i.e., her payment from the mechanism minus her wager), unlike Lambert et al. [2008], who use Π_i to denote the payment only. While this difference is inconsequential, we mention it to avoid confusion. Second, all previous work on wagering mechanisms has restricted attention to *deterministic* profit functions Π_i . Since randomization is necessary to attain privacy, we open up our study to *randomized* profit functions.

Lambert et al. [2008] defined a set of desirable properties or axioms that deterministic wagering mechanisms should arguably satisfy. Here we adapt those properties to potentially randomized wagering mechanisms, making the smallest modifications possible to maintain the spirit of the axioms. Four of the properties (truthfulness, individual rationality, normality, and monotonicity) were originally defined in terms of expected profit with the expectation taken over some true or believed distribution over the outcome ω . We allow the expectation to be over the randomness in the profit function as well. Sybilproofness was not initially defined in expectation; we now ask that this property hold in expectation with respect to the randomness in the profit function. We define anonymity in terms of the distribution over all bettors' profits, and ask that budget balance hold for any realization of the randomness in Π .

- (a) **Budget balance:** The operator makes no profit or loss, i.e., $\forall \mathbf{p} \in [0, 1]^n, \forall \mathbf{m} \in \mathbb{R}_+^n, \forall \omega \in \{0, 1\}$, and for any realization of the randomness in Π , $\sum_{i=1}^n \Pi_i(\mathbf{p}, \mathbf{m}, \omega) = 0$.
- (b) **Anonymity:** Profits do not depend on the identify of the bettors. That is, for any permutation of the bettors σ , $\forall \mathbf{p} \in [0, 1]^n, \forall \mathbf{m} \in \mathbb{R}_+^n, \forall \omega \in \{0, 1\}$, the joint distribution over profit vectors $\{\Pi_i(\mathbf{p}, \mathbf{m}, \omega)\}_{i \in \mathcal{N}}$ is the same as the joint distribution over profit vectors $\{\Pi_{\sigma(i)}((p_{\sigma^{-1}(i)}), (m_{\sigma^{-1}(i)}), \omega)\}_{i \in \mathcal{N}}$.
- (c) **Truthfulness:** Bettors uniquely maximize their expected profit by reporting the truth. That is, $\forall i \in \mathcal{N}, \forall \mathbf{p}_{-i} \in [0, 1]^{n-1}, \forall \mathbf{m} \in \mathbb{R}_+^n, \forall p^*, p_i \in [0, 1]$ with $p_i \neq p^*$,

$$\mathbb{E}_{\omega \sim p^*} [\Pi_i((p^*, \mathbf{p}_{-i}), \mathbf{m}, \omega)] > \mathbb{E}_{\omega \sim p^*} [\Pi_i((p_i, \mathbf{p}_{-i}), \mathbf{m}, \omega)].$$

- (d) **Individual rationality:** Bettors prefer participating to not participating. That is, $\forall i \in \mathcal{N}, \forall m_i > 0$, for all $p^* \in [0, 1]$, there exists some $p_i \in [0, 1]$ such that $\forall \mathbf{p}_{-i} \in [0, 1]^{n-1}, \forall \mathbf{m}_{-i} \in \mathbb{R}_+^{n-1}, \mathbb{E}_{\omega \sim p^*} [\Pi_i((p_i, \mathbf{p}_{-i}), \mathbf{m}, \omega)] \geq 0$.
- (e) **Normality:**¹ If any bettor j changes her report, the change in the expected profit to any other bettor i with respect to a fixed belief p^* is the opposite sign of the change in expected payoff to j . That is, $\forall i, j \in \mathcal{N}, i \neq j, \forall \mathbf{p}, \mathbf{p}' \in [0, 1]^n$ with $p'_k = p_k$ for all $k \neq j, \forall p^* \in [0, 1], \forall \mathbf{m} \in \mathbb{R}_+^n$,

$$\mathbb{E}[\Pi_j(\mathbf{p}, \mathbf{m}, \omega)] < \mathbb{E}[\Pi_j(\mathbf{p}', \mathbf{m}, \omega)] \implies \mathbb{E}[\Pi_i(\mathbf{p}, \mathbf{m}, \omega)] \geq \mathbb{E}[\Pi_i(\mathbf{p}', \mathbf{m}, \omega)].$$

All expectations are taken w.r.t. $\omega \sim p^*$ and the randomness in the mechanism.

- (f) **Sybilproofness:** Profits remain unchanged as any subset of players with the same reports manipulate user accounts by merging accounts, creating fake identities, or transferring wagers. That is, $\forall \mathcal{S} \subset \mathcal{N}, \forall \mathbf{p}$ with $p_i = p_j$ for all $i, j \in \mathcal{S}, \forall \mathbf{m}, \mathbf{m}' \in \mathbb{R}_+^n$ with $m_i = m'_i$ for $i \notin \mathcal{S}$ and $\sum_{i \in \mathcal{S}} m_i = \sum_{i \in \mathcal{S}} m'_i, \forall \omega \in \{0, 1\}$, two conditions hold:

$$\mathbb{E}[\Pi_i(\mathbf{p}, \mathbf{m}, \omega)] = \mathbb{E}[\Pi_i(\mathbf{p}, \mathbf{m}', \omega)] \quad \forall i \notin \mathcal{S},$$

$$\sum_{i \in \mathcal{S}} \mathbb{E}[\Pi_i(\mathbf{p}, \mathbf{m}, \omega)] = \sum_{i \in \mathcal{S}} \mathbb{E}[\Pi_i(\mathbf{p}, \mathbf{m}', \omega)].$$

- (g) **Monotonicity** The magnitude of a bettor's expected profit (or loss) increases as her wager increases. That is, $\forall i \in \mathcal{N}, \forall \mathbf{p} \in [0, 1]^n, \forall \mathbf{m} \in \mathbb{R}_+^n, \forall M_i > m_i, \forall p^* \in [0, 1]$, either $0 < \mathbb{E}_{\omega \sim p^*} [\Pi_i(\mathbf{p}, (m_i, \mathbf{m}_{-i}), \omega)] < \mathbb{E}_{\omega \sim p^*} [\Pi_i(\mathbf{p}, (M_i, \mathbf{m}_{-i}), \omega)]$ or $0 > \mathbb{E}_{\omega \sim p^*} [\Pi_i(\mathbf{p}, (m_i, \mathbf{m}_{-i}), \omega)] > \mathbb{E}_{\omega \sim p^*} [\Pi_i(\mathbf{p}, (M_i, \mathbf{m}_{-i}), \omega)]$.

Previously studied wagering mechanisms [Chen et al. 2014; Lambert et al. 2015, 2008] achieve truthfulness by incorporating *strictly proper scoring rules* [Savage 1971] into their profit functions. Scoring rules reward individuals based on the accuracy of their predictions about random variables. For a binary random variable, a scoring rule s maps a prediction or report $p \in [0, 1]$ and an outcome $\omega \in \{0, 1\}$ to a score. A strictly proper scoring rule incentivizes a risk neutral agent to report her true belief.

DEFINITION 5 (STRICTLY PROPER SCORING RULE [SAVAGE 1971]). A function $s : [0, 1] \times \{0, 1\} \rightarrow \mathbb{R} \cup \{-\infty\}$ is a strictly proper scoring rule if for all $p, q \in [0, 1]$ with $p \neq q, \mathbb{E}_{\omega \sim p} [s(p, \omega)] > \mathbb{E}_{\omega \sim p} [s(q, \omega)]$.

One common example is the Brier scoring rule [Brier 1950], defined as $s(p, \omega) = 1 - (p - \omega)^2$. Note that for the Brier scoring rule, $s(p, x) \in [0, 1]$ for all p and ω . Any strictly proper scoring rule with a bounded range can be scaled to have range $[0, 1]$.

¹Lambert et al. [2015] and Chen et al. [2014] used an alternative definition of normality for wagering mechanisms that essentially requires that if, from some agent i 's perspective, the prediction of agent j improves, then i 's expected profit decreases. This form of normality also holds for our mechanism.

The WSWMs incorporate proper scoring rules, assigning each bettor a profit based on how her score compares to the wager-weighted average score of all bettors, as in Algorithm 1. Lambert et al. [2008] showed that the set of WSWMs satisfy the seven axioms above and is the *unique* set of deterministic mechanisms that simultaneously satisfy budget balance, anonymity, truthfulness, normality, and sybilproofness.

ALGORITHM 1: Weighted-score wagering mechanisms [Lambert et al. 2008]

Parameters: number of bettors n , strictly proper scoring rule s with range in $[0, 1]$

Solicit reports \mathbf{p} and wagers \mathbf{m}

Realize state ω

for $i = 1, \dots, n$ **do**

 Pay bettor i

$$\Pi_i(\mathbf{p}, \mathbf{m}, \omega) = m_i \left(s(p_i, \omega) - \frac{\sum_{j \in \mathcal{N}} m_j s(p_j, \omega)}{\sum_{j \in \mathcal{N}} m_j} \right)$$

end for

3.2. Adding privacy

We would like our wagering mechanism to protect the privacy of each bettor i , ensuring that the $n - 1$ other bettors cannot learn too much about i 's report from their own realized profits, even if they collude. Note that paying each agent according to an independent scoring rule would easily achieve privacy, but would fail budget balance and sybilproofness. We formalize our desire to add privacy to the other good properties of weighted score wagering mechanisms using joint differential privacy.

- (h) **ϵ -joint differential privacy:** The vector of profit functions satisfies ϵ -joint differential privacy, i.e., $\forall i \in \mathcal{N}, \forall \mathbf{p} \in [0, 1]^n, \forall p'_i \in [0, 1], \forall \mathbf{m} \in \mathbb{R}_+^n, \forall \omega \in \{0, 1\}$, and $\forall \mathcal{S} \subset \mathbb{R}_+^{n-1}, \Pr[\Pi_{-i}((p_i, \mathbf{p}_{-i}), \mathbf{m}, \omega) \in \mathcal{S}] \leq e^\epsilon \Pr[\Pi_{-i}((p'_i, \mathbf{p}_{-i}), \mathbf{m}, \omega) \in \mathcal{S}]$.

This definition requires only that the report p_i of each bettor i be kept private, not the wager m_i . Private wagers would impose more severe limitations on the mechanism, even if wagers are restricted to lie in a bounded range; see Section 3.3.2 for a discussion. Note that if bettor i 's report p_i is correlated with his wager m_i , as might be the case for a Bayesian agent [Lambert et al. 2015], then just knowing m_i could reveal information about p_i . In this case, differential privacy would guarantee that other bettors can infer no more about p_i after observing their profits than they could from observing m_i alone. If bettors have immutable beliefs as assumed by Lambert et al. [2008], then reports and wagers are not correlated and m_i reveals nothing about p_i .

Unfortunately, it is not possible to jointly obtain properties (a)–(h) with any reasonable mechanism. This is due to an inherent tension between budget balance and privacy. This is easy to see. Budget balance requires that a bettor i 's profit is the negation of the sum of profits of the other $n - 1$ bettors, i.e., $\Pi_i(\mathbf{p}, \mathbf{m}, \omega) = -\sum_{j \neq i} \Pi_j(\mathbf{p}, \mathbf{m}, \omega)$. Therefore, under budget balance, the other $n - 1$ bettors could always collude to learn bettor i 's profit exactly. In order to obtain privacy, it would therefore be necessary for bettor i 's profit to be differentially private in her own report, resulting in profits that are almost entirely noise. This is formalized in the following theorem. We omit a formal proof since it follows immediately from the argument described here.

THEOREM 3.1. *Let Π be the vector of profit functions for any wagering mechanism that satisfies both budget balance and ϵ -joint differential privacy for any $\epsilon > 0$. Then for all $i \in \mathcal{N}$, Π_i is ϵ -differentially private in bettor i 's report p_i .*

Since it is unsatisfying to consider mechanisms in which a bettor's profit is not sensitive to her own report, we require only that budget balance hold in expectation over

the randomness of the profit function. An operator who runs many markets may be content with such a guarantee as it implies that he will not lose money on average.

(a') **Budget balance in expectation:** The operator neither makes a profit nor a loss in expectation, i.e., $\forall \mathbf{p} \in [0, 1]^n, \forall \mathbf{m} \in \mathbb{R}_+^n, \forall \omega \in \{0, 1\}, \sum_{i=1}^n \mathbb{E} [\Pi_i(\mathbf{p}, \mathbf{m}, \omega)] = 0$.

3.3. Private weighted score wagering mechanisms

Motivated by the argument above, we seek a wagering mechanism that simultaneously satisfies properties (a') and (b)–(h). Keeping Theorem 3.1 in mind, we would also like the wagering mechanism to be defined in such a way that each bettor i 's profit is sensitive to her own report p_i . Sensitivity is difficult to define precisely, but loosely speaking, we would like it to be the case that 1) the magnitude of $\mathbb{E} [\Pi_i(\mathbf{p}, \mathbf{m}, \omega)]$ varies sufficiently with the choice of p_i , and 2) there is not too much noise or variance in a bettor's profit, i.e., $\Pi_i(\mathbf{p}, \mathbf{m}, \omega)$ is generally not too far from $\mathbb{E} [\Pi_i(\mathbf{p}, \mathbf{m}, \omega)]$.

A natural first attempt would be to employ the standard Laplace Mechanism [Dwork and Roth 2014] on top of a WSWM, adding independent Laplace noise to each bettor's profit. The resulting profit vector would satisfy ϵ -joint differential privacy, but since Laplace random variables are unbounded, a bettor could lose more than her wager. Adding other forms of noise does not help; to obtain differential privacy, the noise must be unbounded [Dwork et al. 2006]. Truncating a bettor's profit to lie within a bounded range *after* noise is added could achieve privacy, but would result in a loss of truthfulness as the bettor's expected profit would no longer be a proper scoring rule.

ALGORITHM 2: Private wagering mechanism

Parameters: num bettors n , privacy param ϵ , strictly proper scoring rule s with range in $[0, 1]$

Fix $\alpha = 1 - e^{-\epsilon}$ and $\beta = e^{-\epsilon}$

Solicit reports \mathbf{p} and wagers \mathbf{m}

Realize state ω

for $i = 1, \dots, n$ **do**

Independently draw random variable $x_i(p_i, \omega)$ such that

$$x_i(p_i, \omega) = \begin{cases} 1 & \text{w.p. } \frac{\alpha s(p_i, \omega) + \beta}{1 + \beta} \\ -\beta & \text{w.p. } \frac{1 - \alpha s(p_i, \omega)}{1 + \beta} \end{cases}$$

end for

for $i = 1, \dots, n$ **do**

Pay bettor i

$$\Pi_i(\mathbf{p}, \mathbf{m}, \omega) = m_i \left(\alpha s(p_i, \omega) - \frac{\sum_{j \in \mathcal{N}} m_j x_j(p_j, \omega)}{\sum_{j \in \mathcal{N}} m_j} \right)$$

end for

Instead, we take a different approach. Like the WSWM, our *private wagering mechanism*, formally defined in Algorithm 2, rewards each bettor based on how good his score is compared with an aggregate measure of how good bettors' scores are on the whole. However, this aggregate measure is now calculated in a noisy manner. That is, instead of comparing a bettor's score to a weighted average of all bettors' scores, the bettor's score is compared to a weighted average of random variables that are equal to bettors' scores in expectation. As a result, each bettor's profit is, in expectation, equal to the profit she would receive using a WSWM, scaled down by a parameter α to ensure that no bettor ever loses more than her wager, as stated in the following lemma. The proof, which simply shows that for each i , $\mathbb{E}[x_i(p_i, \omega)] = \alpha s(p_i, \omega)$, is in the appendix.

LEMMA 3.2. *For any number of bettors $n > 0$ with reports $\mathbf{p} \in [0, 1]^n$ and wagers $\mathbf{m} \in \mathbb{R}_+^n$, for any setting of the privacy parameter $\epsilon > 0$, for any outcome $\omega \in \{0, 1\}$,*

the expected value of bettor i 's profit $\Pi_i(\mathbf{p}, \mathbf{m}, \omega)$ under the private wagering mechanism with scoring rule s is equal to bettor i 's profit under a WSWM with scoring rule αs .

Using this lemma, we show that this mechanism does indeed satisfy joint differential privacy as well as the other desired properties.

THEOREM 3.3. *The private wagering mechanism satisfies (a') budget balance in expectation, (b) anonymity, (c) truthfulness, (d) individual rationality, (e) normality, (f) sybilproofness, (g) monotonicity, and (h) ϵ -joint differential privacy.*

PROOF. Any WSWM satisfies budget balance in expectation (by satisfying budget balance), truthfulness, individual rationality, normality, sybilproofness, and monotonicity [Lambert et al. 2008]. Since these properties are defined in terms of expected profit, Lemma 3.2 implies that the private wagering mechanism satisfies them too.

Anonymity is easily observed since profits are defined symmetrically for all bettors.

Finally we show ϵ -joint differential privacy. We first prove that each random variable $x_i(p_i, \omega)$ is ϵ -differentially private in bettor i 's report p_i which implies that the noisy aggregate of scores is private in all bettors' reports. We then apply the billboard lemma (see Section 2) to show that the profit vector Π satisfies joint differential privacy.

To show that $x_i(p_i, \omega)$ is differentially private in p_i , for each of the two values that $x_i(p_i, \omega)$ can take on we must ensure that the ratio of the probability it takes this value under any report p and the probability it takes this value under any alternative report p' is bounded by e^ϵ . Fix any $\omega \in \{0, 1\}$. Since s has range in $[0, 1]$,

$$\frac{\Pr(x_i(p, \omega) = 1)}{\Pr(x_i(p', \omega) = 1)} = \frac{\alpha s(p, \omega) + \beta}{\alpha s(p', \omega) + \beta} \leq \frac{\alpha + \beta}{\beta} = \frac{1 - e^{-\epsilon} + e^{-\epsilon}}{e^{-\epsilon}} = e^\epsilon,$$

$$\frac{\Pr(x_i(p, \omega) = -\beta)}{\Pr(x_i(p', \omega) = -\beta)} = \frac{1 - \alpha s(p, \omega)}{1 - \alpha s(p', \omega)} \leq \frac{1}{1 - \alpha} = \frac{1}{1 - (1 - e^{-\epsilon})} = e^\epsilon.$$

Thus $x_i(p_i, \omega)$ is ϵ -differentially private in p_i . By Theorem 4 of McSherry [2009], the vector $(x_1(p_1, \omega), \dots, x_n(p_n, \omega))$ (and thus any function of this vector) is ϵ -differentially private in the vector \mathbf{p} , since each $x_i(p_i, \omega)$ does not depend on the reports of anyone but i . Since we view the wagers m_i as constants, the quantity $\sum_{j \in \mathcal{N}} m_j x_j(p_j, \omega) / \sum_{j \in \mathcal{N}} m_j$ is also ϵ -differentially private in the reports \mathbf{p} . Call this quantity X .

To apply the billboard lemma, we can imagine the operator publicly announcing the quantity X to the bettors. Given access to X , each bettor is able to calculate her own profit $\Pi_i(\mathbf{p}, \mathbf{m}, \omega)$ using only her own input and the values α and ω . The billboard lemma implies that the vector of profits is ϵ -joint differentially private. \square

3.3.1. Sensitivity of the mechanism. Having established that our mechanism satisfies properties (a') and (b)–(h), we next address the sensitivity of the mechanism in terms of the two facets described above: range of achievable expected profits and the amount of noise in the profit function. This discussion sheds light on how to set ϵ in practice.

The first facet is quantified by Lemma 3.2. As α grows, the magnitude of bettors' expected profits grows, and the range of expected profits grows as well. When α approaches 1, the range of expected profits achievable through the private wagering mechanism approaches that of a standard WSWM with the same proper scoring rule.

Unfortunately, since $\alpha = 1 - e^{-\epsilon}$, larger values of α imply larger values of the privacy parameter ϵ . This gives us a clear tradeoff between privacy and magnitude of expected payments. Luckily, in practice, it is probably unnecessary for ϵ to be very small for most markets. A relatively large value of ϵ can still give bettors plausible deniability. For example, setting $\epsilon = 1$ implies that a bettor's report can only change the probability

of another bettor receiving a particular profit by a factor of roughly 2.7 and leads to $\alpha \approx 0.63$, a tradeoff that may be considered acceptable in practice.

The second facet is quantified in the following theorem, which states that as more money is wagered by more bettors, each bettor's realized profit approaches its expectation. The bound depends on $\|\mathbf{m}\|_2/\|\mathbf{m}\|_1$. If all wagers are equal, this quantity is equal to $1/\sqrt{n}$ and bettors' profits approach their expectations as n grows. This is not the case at the other extreme, when there are a small number of bettors with wagers much larger than the rest. The proof, which uses Hoeffding's inequality to bound the difference between the quantity $m_j x_j(p_j, \omega)$ and its expectation, is in the appendix.

THEOREM 3.4. *For any $\delta \in [0, 1]$, any $\epsilon > 0$, any number of bettors $n > 0$, any vectors of reports $\mathbf{p} \in [0, 1]^n$ and wagers $\mathbf{m} \in \mathbb{R}_+^n$, with probability at least $1 - \delta$, for all $i \in \mathcal{N}$, the profit Π_i output by the private wagering mechanism satisfies*

$$|\Pi_i(\mathbf{p}, \mathbf{m}, \omega) - \mathbb{E}[\Pi_i(\mathbf{p}, \mathbf{m}, \omega)]| \leq m_i \left(\frac{\|\mathbf{m}\|_2}{\|\mathbf{m}\|_1} (1 + \beta) \sqrt{\frac{\ln(2/\delta)}{2}} \right).$$

The following corollary shows that if all wagers are bounded in some range $[L, U]$, profits approach their expectations as the number of bettors grows.

COROLLARY 3.5. *Fix any L and U , $0 < L < U$. For any $\delta \in [0, 1]$, any $\epsilon > 0$, any $n > 0$, any vectors of reports $\mathbf{p} \in [0, 1]^n$ and wagers $\mathbf{m} \in [L, U]^n$, with probability at least $1 - \delta$, for all $i \in \mathcal{N}$, the profit Π_i output by the private wagering mechanism satisfies*

$$|\Pi_i(\mathbf{p}, \mathbf{m}, \omega) - \mathbb{E}[\Pi_i(\mathbf{p}, \mathbf{m}, \omega)]| \leq m_i \left(\frac{U}{\sqrt{n}L} (1 + \beta) \sqrt{\frac{\ln(2/\delta)}{2}} \right).$$

3.3.2. Keeping wagers private. Property (h) requires that bettors' reports be kept private but does not guarantee private wagers. The same tricks used in our private wagering mechanism could be applied to obtain a privacy guarantee for both reports and wagers if wagers are restricted to lie in a bounded range $[L, U]$, but this would come with a great loss in sensitivity. Under the most straightforward extension, the parameter α would need to be set to $(L/U)(1 - e^{-\epsilon/n})$ rather than $(1 - e^{-\epsilon})$, greatly reducing the scale of achievable profits and thus making the mechanism impractical in most settings.

Loosely speaking, the extra factor of L/U stems from the fact that a bettor's effect on the profit of any other bettor must be roughly the same whether he wagers the maximum amount or the minimum. The poor dependence on n is slightly more subtle. We created a private-belief mechanism by replacing each bettor j 's score $s(p_j, \omega)$ in the WSWM with a random variable $x_j(p_j, \omega)$ that is ϵ -differentially private in p_j . To obtain private wagers, we would instead need to replace the full term $m_j s(p_j, \omega) / \sum_{k \in \mathcal{N}} m_k$ with a random variable for each j . This term depends on the wagers of *all* n bettors in addition to p_j . Since each bettor's profit would depend on n such random variables, achieving ϵ -joint differential privacy would require that each random variable be ϵ/n -differentially private in each bettor's wager.

We believe that sacrifices in sensitivity are unavoidable and not merely an artifact of our techniques and analysis, but leave a formal lower bound to future work.

4. LIMITS OF PRIVACY WITH COST-FUNCTION MARKET MAKERS

In practice, prediction markets are often run using dynamic mechanisms that update in real time as new information surfaces. We now turn to the problem of adding privacy guarantees to continuous-trade markets. We focus our attention on cost-function prediction markets, in which all trades are placed through an automated market maker [Abernethy et al. 2013; Chen and Pennock 2007; Hanson 2003]. The market maker

can be viewed as a streaming algorithm that takes as input a stream of trades and outputs a corresponding stream of market states from which trade prices can be computed. Therefore, the privacy guarantees we seek are in the form of Definition 3. We ask whether it is possible for the automated market maker to price trades according to a cost function while maintaining $\epsilon(t)$ -differential privacy without opening up the opportunity for traders to earn unbounded profits, leading the market maker to experience unbounded loss. We show a mostly negative result: to achieve bounded loss, the privacy term $e^{\epsilon(t)}$ must grow faster than linearly in t , the number of rounds of trade.

For simplicity, we state our results for markets over a single binary security, though we believe they extend to cost-function markets over arbitrary security spaces.

4.1. Standard cost-function market makers

We consider a setting in which there is a single binary security that traders may buy or sell. After the outcome $\omega \in \{0, 1\}$ has been revealed, a share of the security is worth \$1 if $\omega = 1$ and \$0 otherwise. A cost-function prediction market for this security is fully specified by a convex function C called the *cost function*. Let x_t be the number of shares that are bought or sold by a trader in the t th transaction; positive values of x_t represent purchases while negative values represent (short) sales. The market state after the first $t - 1$ trades is summarized by a single value $q_t = \sum_{\tau=1}^{t-1} x_\tau$, and the t th trader is charged $C(q_t + x_t) - C(q_t) = C(q_{t+1}) - C(q_t)$. Thus the cost function can be viewed as a potential function, with $C(q_{t+1}) - C(0)$ capturing the amount of money that the market maker has collected from the first t trades. The *instantaneous price* at round t , denoted p_t , is the price per share of purchasing an infinitesimally small quantity of shares: $p_t = C'(q_t)$. This framework is summarized in Algorithm 3.

ALGORITHM 3: Cost-function market maker (parameters: cost function C)

```

Initialize:  $q_1 = 0$ 
for  $t = 1, 2, \dots$  do
    Update instantaneous price  $p_t = C'(q_t)$ 
    A trader buys  $x_t \in \mathbb{R}$  shares and pays  $C(q_t + x_t) - C(q_t)$ 
    Update market state  $q_{t+1} = q_t + x_t$ 
end for
Realize outcome  $\omega$ 
if  $\omega = 1$  then
    for  $t = 1, 2, \dots$  do
        Market maker pays  $x_t$  to the trader from round  $t$ 
    end for
end if

```

The most common cost-function market maker is Hanson’s logarithmic market scoring rule (LMSR) [Hanson 2003]. The cost function for the single-security version of LMSR can be written as $C(q) = b \log(e^{(q+a)/b} + 1)$ where $b > 0$ is a parameter controlling the rate at which prices change as trades are made and a controls the initial market price at state $q = 0$. The instantaneous price at any state q is $C'(q) = e^{(q+a)/b} / (e^{(q+a)/b} + 1)$.

Under mild conditions on C , all cost-function market makers satisfy several desirable properties, including natural notions of no-arbitrage and information incorporation [Abernethy et al. 2013]. We refer to any cost function C satisfying these mild conditions as a *standard cost function*. Although the market maker subsidizes trade, crucially its worst-case loss is bounded. This ensures that the market maker does not go bankrupt, even if traders are perfectly informed. Formally, there exists a finite bound B such that for any T , any sequence of trades x_1, \dots, x_T , and any outcome $\omega \in \{0, 1\}$,

$$q_{T+1} \cdot \mathbf{1}(\omega = 1) - (C(q_{T+1}) - C(0)) \leq B,$$

where $\mathbf{1}$ is the indicator function that is 1 if its argument is true and 0 otherwise. The first term on the left-hand side is the amount that the market maker must pay (or collect from) traders when ω is revealed. The second is the amount collected from traders. For the LMSR with initial price $p_1 = 0.5$ ($a = 0$), the worst-case loss is $b \log(2)$.

4.2. The noisy cost-function market maker

Clearly the standard cost-function market maker does not ensure differential privacy. The amount that a trader pays is a function of the market state, the sum of all past trades. Thus anyone observing the stream of market prices could infer the exact sequence of past trades. To guarantee privacy while still approximating cost-function pricing, the market maker would need to modify the sequence of published prices (or equivalently, market states) to ensure that such information leakage does not occur.

In this section, we define and analyze a *noisy* cost-function market maker. The noisy market maker prices trades according to a cost function, but uses a noisy version of the market state in order to mask the effect of past trades. In particular, the market maker maintains a noisy market state $q'_t = q_t + \eta_t$, where q_t is the true sum of trades and η_t is a (random) noise term. The cost of trade x_t is $C(q'_t + x_t) - C(q'_t)$, with the instantaneous price now $p_t = C'(q'_t)$. Since the noise term η_t must be large enough to mask the trade x_t , we limit trades to be some maximum size k . A trader who would like to buy or sell more than k shares must do this over multiple rounds. The full modified framework is shown in Algorithm 4. For now we allow the noise distribution \mathcal{D} to depend arbitrarily on the history of trade. This framework is general; the natural adaptation of the privacy-preserving data market of Waggoner et al. [2015] to the single security prediction market setting would result in a market maker of this form, as would a cost-function market that used existing private streaming techniques for bit counting [Chan et al. 2011; Dwork et al. 2010] to keep noisy, private counts of trades.

ALGORITHM 4: Noisy cost-function market maker (parameters: cost function C , distribution \mathcal{D} over noise $\{\eta_t\}$, maximum trade size k)

```

Initialize:  $q_1 = 0$ 
Draw  $\eta_1$  and set  $q'_1 = \eta_1$ 
for  $t = 1, 2, \dots$  do
  Update instantaneous price  $p_t = C'(q'_t)$ 
  A trader buys  $x_t \in [-k, k]$  shares and pays  $C(q'_t + x_t) - C(q'_t)$ 
  Update true market state  $q_{t+1} = q_t + x_t$ 
  Draw  $\eta_{t+1}$  and update noisy market state  $q'_{t+1} = q_{t+1} + \eta_{t+1}$ 
end for
Realize outcome  $\omega$ 
if  $\omega = 1$  then
  for  $t = 1, 2, \dots$  do
    Market maker pays  $x_t$  to the trader from round  $t$ 
  end for
end if

```

In this framework, we can interpret the market maker as implementing a noise trader in a standard cost-function market. Under this interpretation, after a (real) trader purchases x_t shares at state q'_t , the market state momentarily moves to $q'_t + x_t = q_t + \eta_t + x_t = q_{t+1} + \eta_t$. The market maker, acting as a noise trader, then effectively “purchases” $\eta_{t+1} - \eta_t$ shares at this state for a cost of $C((q_{t+1} + \eta_t) + (\eta_{t+1} - \eta_t)) - C(q_{t+1} + \eta_t) = C(q_{t+1} + \eta_{t+1}) - C(q_{t+1} + \eta_t)$, bringing the market state to $q_{t+1} + \eta_{t+1} = q'_{t+1}$. The market maker makes this trade regardless of the impact on its own loss. These noise trades obscure the trades made by real traders, opening up the possibility of privacy.

However, these noisy trades also open up the opportunity for traders to profit off of the noise. It is important to ensure that bounded worst-case loss is maintained. For the noisy cost-function market maker, for any sequence of T trades x_1, \dots, x_T , any outcome $\omega \in \{0, 1\}$, and any *fixed* noise values η_1, \dots, η_T , the market maker's loss is

$$L_T(x_1, \dots, x_T, \eta_1, \dots, \eta_T, \omega) \equiv q_{T+1} \cdot \mathbf{1}(\omega = 1) - \sum_{t=1}^T (C(q'_t + x_t) - C(q'_t)).$$

As before, the first term is the (possibly negative) amount that the market maker pays to traders when ω is revealed, and the second is the amount collected from traders (which no longer telescopes). Unfortunately, we cannot expect this loss to be bounded for *any* noise values; the market maker could always get extremely unlucky and draw noise values that traders can exploit. Instead, we consider a relaxed version of bounded loss which holds in expectation with respect to the noise values η_t .

In addition to this relaxation, one more modification is necessary. Note that traders can (and should) base their actions on the current market price. Therefore, if our loss guarantee only holds in expectation with respect to noise values η_t , then it is no longer sufficient to give a guarantee that is worst case over any sequences of trades. Instead, we allow the sequence of trades to depend on the realized noise, introducing a game between traders and the market maker. To formalize this, we imagine allowing an adversary to control the traders. We define the notion of a *strategy* for this adversary.

DEFINITION 6 (TRADER STRATEGY). *A trader strategy \mathbf{s} is a set of (possibly randomized) functions $\mathbf{s} = \{s_1, s_2, \dots\}$, with each s_t mapping a history of trades and noisy market states $(x_1, \dots, x_{t-1}, q'_1, \dots, q'_t)$ to a new trade x_t for the trader at round t .*

Let S be the set of all strategies. With this definition in place, we can formally define what it means for a noisy cost-function market maker to have bounded loss.

DEFINITION 7 (BOUNDED LOSS FOR A NOISY COST-FUNCTION MARKET MAKER). *A noisy cost-function market maker with cost function C and distribution \mathcal{D} over noise values η_1, η_2, \dots is said to have bounded loss if there exists a finite B such that for all strategies $\mathbf{s} \in S$, all times $T \geq 1$, and all $\omega \in \{0, 1\}$,*

$$\mathbb{E}[L_T(x_1, \dots, x_T, \eta_1, \dots, \eta_T, \omega)] \leq B,$$

where the expectation is taken over the market's noise values η_1, η_2, \dots distributed according to \mathcal{D} and the (possibly randomized) actions x_1, x_2, \dots of a trader employing strategy \mathbf{s} . In this case, the loss of the market maker is said to be bounded by B . The noisy cost-function market maker has unbounded loss if no such B exists.

If the noise values were deterministic, this definition of worst-case loss would correspond to the usual one, but because traders react intelligently to the specific realization of noise, we must define worst-case loss in game-theoretic terms.

4.3. Limitations on privacy

By effectively acting as a noise trader, a noisy cost-function market maker can partially obscure trades. Unfortunately, the amount of privacy achievable through this technique is limited. In this section, we show that in order to simultaneously maintain bounded loss and achieve $\epsilon(t)$ -differential privacy, the quantity $e^{\epsilon(t)}$ must grow faster than linearly as a function of the number of rounds of trade.

Before stating our result, we explain how to frame the market maker setup in the language of differential privacy. Recall from Section 2 that a differentially private unbounded streaming algorithm \mathcal{M} takes as input a stream σ of arbitrary length and outputs a stream of values that depend on σ in a differentially private way. In the

market setting, the stream σ corresponds to the sequence of trades $\mathbf{x} = (x_1, x_2, \dots)$. We think of the noisy cost-function market maker (Algorithm 4) as an algorithm \mathcal{M} that, on any stream prefix (x_1, \dots, x_t) , outputs the noisy market states (q'_1, \dots, q'_{t+1}) .² The goal is to find a market maker such that \mathcal{M} is $\epsilon(t)$ -differentially private.

One might ask whether it is necessary to allow the privacy guarantee to diminish as the the number of trades grows. When considering the problem of calculating noisy sums of bit streams, for example, Chan et al. [2011] are able to maintain a fixed privacy guarantee as their stream grows in length by instead allowing the accuracy of their counts to diminish. This approach doesn't work for us; we cannot achieve bounded loss yet allow the market maker's loss to grow with the number of trades.

Our result relies on one mild assumption on the distribution \mathcal{D} over noise. In particular, we require that the noise η_{t+1} be chosen independent of the current trade x_t . We refer to this as the *trade-independent noise assumption*.³ The distribution of η_{t+1} may still depend on the round t , the history of trade x_1, \dots, x_{t-1} , and the realizations of past noise terms, η_1, \dots, η_t . This assumption is needed in the proof only to rule out unrealistic market makers that are specifically designed to monitor and infer the behavior of the specific adversarial trader that we consider, and the result likely holds even without it. However, it is not a terribly restrictive assumption as most standard ways of generating noise could be written in this form. For example, Chan et al. [2011] and Dwork et al. [2010] show how to maintain a noisy count of the number of ones in a stream of bits by computing the exact count and adding noise that is correlated across time but independent of the data. If similar ideas were used to choose the noise term in our setting, the trade-independent noise assumption would be satisfied. The noise employed in the mechanism of Waggoner et al. [2015] also satisfies this assumption. Our impossibility result then implies that their market would have unbounded loss if a limit on the number of rounds of trade were not imposed. To obtain privacy guarantees, Waggoner et al. must assume that the number of trades is known in advance and can therefore be used to set relevant market parameters.

We now state the main result.

THEOREM 4.1. *Consider any noisy cost-function market maker using a standard convex cost function C that is nonlinear in some region, a noise distribution \mathcal{D} satisfying the trade-independent noise assumption, and a bound $k > 0$ on trade size. If the market maker satisfies bounded loss, then it cannot satisfy $(\epsilon(t), \delta)$ -differential privacy for any function ϵ such that $e^{\epsilon(t)} = O(t)$ with any constant $\delta \in [0, 1)$.*

This theorem rules out bounded loss with $\epsilon(t) = \log(mt)$ for any constant $m > 0$. It is open whether it is possible to achieve $\epsilon(t) = m \log(t)$ (and therefore $e^{\epsilon(t)} = t^m$) for some $m > 1$, but such a guarantee would likely be insufficient in most practical settings.

Note that with unbounded trade size (i.e., $k = \infty$), our result would be trivial. A trader could move the market price an arbitrary amount in one trade. To provide privacy, the noisy market state would need to be independent of past trades. The price would not be reflective of trader beliefs, and the noise could be exploited by traders for profit. By imposing a bound on trade size, we only strengthen our negative result.

While the proof of Theorem 4.1 is quite technical, the intuition is simple. We consider the behavior of the noisy cost-function market maker when there is a single trader trading in the market repeatedly using a simple trading strategy. This trader chooses a *target state* q^* . Whenever the noisy market state q'_t is less than q^* (and so $p_t <$

²Announcing q'_t allows traders to infer the instantaneous price $p_t = C'(q'_t)$. It is equivalent to announcing p_t in terms of information revealed as long as C is strictly convex in the region around q'_t .

³The proof can be extended easily to the more general case in which the calculation of η_{t+1} is differentially private in x_t ; we make the slightly stronger assumption to simplify presentation.

$p^* \equiv C'(q^*)$, the trader purchases shares, pushing the market state as close to q^* as possible. When the noisy state q'_t is greater than q^* (so $p_t > p^*$), the trader sells shares, again pushing the state as close as possible to q^* . Each trade makes a profit for the trader *in expectation* if it were the case that $\omega = 1$ with probability p^* . Since there is only a single trader, this means that each such trade would result in an expected loss with respect to p^* for the market maker. Unbounded expected loss for any p^* implies unbounded loss in the worst case—either when $\omega = 0$ or $\omega = 1$. The crux of the proof involves showing that in order to achieve bounded loss against this trader, the amount of added noise η_t cannot be too big as t grows, resulting in a sacrifice of privacy.

To formalize this intuition, we first give a more precise description of the strategy \mathbf{s}^* employed by the single trader we consider.

DEFINITION 8 (TARGET STRATEGY). *The target strategy \mathbf{s}^* with target $q^* \in \mathbb{R}$ chosen from a region in which C is nonlinear is defined as follows. For all rounds t ,*

$$s_t^*(x_1, \dots, x_{t-1}, q'_1, \dots, q'_t) = \begin{cases} \min\{q^* - q'_t, k\} & \text{if } q'_t \leq q^*, \\ -\min\{q'_t - q^*, k\} & \text{otherwise.} \end{cases}$$

As described above, if $\omega = 1$ with probability p^* , a trader following this target strategy makes a non-negative expected profit on every round of trade. Furthermore, this trader makes an expected profit of at least some constant $\chi > 0$ on each round in which the noisy market state q'_t is more than a constant distance γ from q^* . The market maker must subsidize this profit, taking an expected loss with respect to p^* on each round. These ideas are formalized in Lemma 4.2, which lower bounds the expected loss of the market maker in terms of the probability of q'_t falling far from q^* . In this statement, D_C denotes the Bregman divergence⁴ of C . The proof is in the appendix.

LEMMA 4.2. *Consider a noisy cost-function market maker satisfying the conditions in Theorem 4.1 with a single trader following the target strategy \mathbf{s}^* with target q^* . Suppose $\omega = 1$ with probability $p^* = C'(q^*)$. Then for any γ such that $0 < \gamma \leq k$,*

$$\mathbb{E}[L_T(x_1, \dots, x_T, \eta_1, \dots, \eta_T, \omega)] \geq \chi \sum_{t=1}^T \Pr(|q'_t - q^*| \geq \gamma)$$

where the expectation and probability are taken over the randomness in the noise values η_1, η_2, \dots , the resulting actions x_1, x_2, \dots of the trader, and the random outcome ω , and where $\chi = \min\{D_C(q^* + \gamma, q^*), D_C(q^* - \gamma, q^*)\} > 0$.


We now complete the proof.

PROOF OF THEOREM 4.1. We will show that bounded loss implies that $(\epsilon(t), \delta)$ -differential privacy cannot be achieved with $e^{\epsilon(t)} = O(t)$ for any constant $\delta \in [0, 1)$.

Throughout the proof, we reason about the probabilities of various events conditioned on there being a single trader playing a particular strategy. All strategies we consider are deterministic, so all probabilities are taken just with respect to the randomness in the market maker's added noise (η_1, η_2, \dots) .

As described above, we focus on the case in which a single trader plays the target strategy \mathbf{s}^* with target q^* . Define R^* to be the open region of radius $k/4$ around q^* , that is, $R^* = (q^* - k/4, q^* + k/4)$. Let $\hat{q} = q^* + k/2$ and let $\hat{R} = (\hat{q} - k/4, \hat{q} + k/4)$. Notice that R^* and \hat{R} do not intersect, but from any market state $q \in R^*$ a trader could move the market state to \hat{q} with a purchase or sale of no more than k shares.

⁴The Bregman divergence of a convex function F of a single variable is defined as $D_F(p, q) = F(p) - F(q) - F'(q)(p - q)$. The Bregman divergence is always non-negative. If F is strictly convex, it is strictly positive when the arguments are not equal.



For any round t , let \mathbf{s}^t be the strategy in which $s_\tau^t = s_\tau^*$ for all rounds $\tau \neq t$, but $s_t^t(x_1, \dots, x_{t-1}, q'_1, \dots, q'_t) = \hat{q} - q'_t$ if $|\hat{q} - q'_t| \leq k$ (otherwise, s_t^t can be defined arbitrarily). In other words, a trader playing strategy \mathbf{s}^t behaves identically to a trader playing strategy \mathbf{s}^* on all rounds except round t . On round t , the trader instead attempts to move the market state to \hat{q} .

For any t , the behavior of a trader playing strategy \mathbf{s}^* and a trader playing strategy \mathbf{s}^t are indistinguishable through round $t-1$, and therefore the behavior of the market maker is indistinguishable as well. At round t , if it is the case that $q'_t \in R^*$ (and therefore $|q'_t - q^*| \leq k/4 < k$ and also $|q'_t - \hat{q}| \leq 3k/4 < k$), then a trader playing strategy \mathbf{s}^* would purchase $q^* - q'_t$ shares, while a trader playing strategy \mathbf{s}^t would purchase $\hat{q} - q'_t$. Differential privacy tells us that conditioned on such a state being reached, the probability that q'_{t+1} lies in any range (and in particular, in R^*) should not be too different depending on which of the two actions the trader takes. More formally, if the market maker satisfies $\epsilon(t)$ -differential privacy, then for all rounds t ,

$$\begin{aligned} e^{\epsilon(t)} &\geq \frac{\Pr(q'_{t+1} \in R^* | \mathbf{s} = \mathbf{s}^*, q'_t \in R^*) - \delta}{\Pr(q'_{t+1} \in R^* | \mathbf{s} = \mathbf{s}^t, q'_t \in R^*)} \geq \frac{\Pr(q'_{t+1} \in R^* | \mathbf{s} = \mathbf{s}^*, q'_t \in R^*) - \delta}{\Pr(q'_{t+1} \notin \hat{R} | \mathbf{s} = \mathbf{s}^t, q'_t \in R^*)} \\ &= \frac{\Pr(q'_{t+1} \in R^* | \mathbf{s} = \mathbf{s}^*, q'_t \in R^*) - \delta}{\Pr(q'_{t+1} \notin R^* | \mathbf{s} = \mathbf{s}^*, q'_t \in R^*)}. \end{aligned}$$

The first inequality follows from the definition of $(\epsilon(t), \delta)$ -differential privacy. The second follows from the fact that R^* and \hat{R} are disjoint. The last line is a consequence of the trade-independent noise assumption. By simple algebraic manipulation, for all t ,

$$\Pr(q'_{t+1} \notin R^* | \mathbf{s} = \mathbf{s}^*, q'_t \in R^*) \geq \frac{1 - \delta}{1 + e^{\epsilon(t)}}. \quad (1)$$

We now further investigate the term on the left-hand side of this equation. For the remainder of the proof, we assume that $\mathbf{s} = \mathbf{s}^*$ and implicitly condition on this.

Applying Lemma 4.2 with $\gamma = k/4$, we find that the expected value of the market maker's loss after T rounds if $\omega = 1$ with probability $p^* = C'(q^*)$ is lower bounded by $\chi \sum_{t=1}^T \Pr(q'_t \notin R^*)$ for the appropriate constant χ . This implies that for at least one of $\omega = 1$ or $\omega = 0$, $\mathbb{E}[L_T(x_1, \dots, x_T, \eta_1, \dots, \eta_T, \omega)] \geq \chi \sum_{t=1}^T \Pr(q'_t \notin R^*)$ where the expectation is just over the random noise of the market maker and the resulting actions of the trader. Since we have assumed that the market maker's loss is bounded, this implies there must exist some loss bound B such that

$$\frac{B}{\chi} \geq \sum_{t=1}^{\infty} \Pr(q'_t \notin R^*). \quad (2)$$

Fix any constant $\alpha \in (0, 1)$. Equation 2 implies that for all but finitely many t , $\Pr(q'_t \notin R^*) < \alpha$, or equivalently, for all but finitely many t , $\Pr(q'_t \in R^*) \geq 1 - \alpha$. Call the set of t for which this holds \mathcal{T} . Equation 2 also implies that

$$\begin{aligned} \frac{B}{\chi} &\geq \sum_{t=1}^{\infty} [\Pr(q'_{t+1} \notin R^* | q'_t \in R^*) \Pr(q'_t \in R^*) + \Pr(q'_{t+1} \notin R^* | q'_t \notin R^*) \Pr(q'_t \notin R^*)] \\ &\geq \sum_{t=1}^{\infty} \Pr(q'_{t+1} \notin R^* | q'_t \in R^*) \Pr(q'_t \in R^*) \geq (1 - \alpha) \sum_{t \in \mathcal{T}} \Pr(q'_{t+1} \notin R^* | q'_t \in R^*). \end{aligned}$$

Combining this with Equation 1 yields

$$\sum_{t \in \mathcal{T}} \frac{1 - \delta}{1 + e^{\epsilon(t)}} \leq \frac{B}{\chi(1 - \alpha)}. \quad (3)$$

Now suppose for contradiction that $e^{\epsilon(t)} = O(t)$. Then by definition, for some constant $m > 1$ there exists a round τ such that for all $t > \tau$, $e^{\epsilon(t)} \leq mt$. Then

$$\sum_{t \in \mathcal{T}} \frac{1 - \delta}{1 + e^{\epsilon(t)}} \geq \sum_{t \in \mathcal{T}, t > \tau} \frac{1 - \delta}{1 + e^{\epsilon(t)}} \geq \sum_{t \in \mathcal{T}, t > \tau} \frac{1 - \delta}{1 + mt} > \frac{1 - \delta}{m} \sum_{t \in \mathcal{T}, t > \tau} \frac{1}{1 + t}.$$

Since this sum is over all natural numbers t except a finite number, it must diverge, and therefore Equation 3 cannot hold. Therefore, we cannot have $e^{\epsilon(t)} = O(t)$. \square

5. DISCUSSION

We designed a class of randomized wagering mechanisms that keep bettors' reports private while maintaining truthfulness, budget balance in expectation, and other desirable properties of weighted score wagering mechanisms. The parameters of our mechanisms can be tuned to achieve a tradeoff between the level of privacy guaranteed and the sensitivity of a bettor's payment to her own report. Determining how to best make this tradeoff in practice (and more generally, what level of privacy is acceptable in differentially private algorithms) is an open empirical question.

While our results in the dynamic setting are negative, there are several potential avenues for circumventing our lower bound. The lower bound shows that it is not possible to obtain reasonable privacy guarantees using a noisy cost-function market maker when traders may buy or sell fractional security shares, as is typically assumed in the cost function literature. Indeed, the adversarial trader we consider buys and sells arbitrarily small fractions when the market state is close to its target. This behavior could be prevented by enforcing a minimum unit of purchase. Perhaps cleverly designed noise could allow us to avoid the lower bound with this additional restriction. However, based on preliminary simulations of a noisy cost-function market based on Hanson's LMSR [2003] with noise drawn using standard binary streaming approaches [Chan et al. 2011; Dwork et al. 2010], it appears an adversary can still cause a market maker using these techniques to have unbounded loss by buying one unit when the noisy market state is below the target and selling one unit when it is above.

One could also attempt to circumvent the lower bound by adding a transaction fee for each trade that is large enough that traders cannot profit off the market's noise. While the fee could always be set large enough to guarantee bounded loss, a large fee would discourage trade in the market and limit its predictive power. A careful analysis would be required to ensure that the fee could be set high enough to maintain bounded loss without rendering the market predictions useless.

REFERENCES

- J. Abernethy, Y. Chen, and J. W. Vaughan. 2013. Efficient market making via convex optimization, and a connection to online learning. *ACM Transactions on Economics and Computation* 1, 2 (2013).
- H. Berg and T. A. Proebsting. 2009. Hanson's automated market maker. *Journal of Prediction Markets* 3, 1 (2009), 45–59.
- J. E. Berg, R. Forsythe, F. D. Nelson, and T. A. Rietz. 2001. Results from a dozen years of election futures markets research. In *Handbook of Experimental Economic Results*, C. A. Plott and V. Smith (Eds.).

- G. W. Brier. 1950. Verification of forecasts expressed in terms of probability. *Monthly Weather Review* 78, 1 (1950), 1–3.
- T.-H. H. Chan, E. Shi, and D. Song. 2011. Private and continual release of statistics. *ACM Transactions on Information and System Security* 14, 3 (2011), 26.
- R. Charette. 2007. An internal futures market. *Information Management* (2007).
- Y. Chen, N. R. Devanur, D. M. Pennock, and J. W. Vaughan. 2014. Removing arbitrage from wagering mechanisms. In *15th ACM Conf. on Economics and Computation*.
- Y. Chen and D. M. Pennock. 2007. A utility framework for bounded-loss market makers. In *Conference on Uncertainty in Artificial Intelligence*.
- B. Cowgill and E. Zitzewitz. 2015. Corporate prediction markets: Evidence from Google, Ford, and Firm X. *Review of Economic Studies* 82, 4 (2015), 1309–1341.
- C. Dwork, F. McSherry, K. Nissim, and A. Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *3rd Conference on Theory of Cryptography*.
- C. Dwork, M. Naor, T. Pitassi, and G. N. Rothblum. 2010. Differential privacy under continual observation. In *42nd ACM Symposium on Theory of Computing*.
- C. Dwork and A. Roth. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Comp. Sci.* 9, 34 (2014), 211–407.
- J. M. Gandar, W. H. Dare, C. R. Brown, and R. A. Zuber. 1999. Informed traders and price variations in the betting market for professional basketball games. *Journal of Finance* LIII, 1 (1999), 385–401.
- T. Gneiting and A. E. Raftery. 2007. Strictly proper scoring rules, prediction, and estimation. *J. Amer. Statist. Assoc.* 102, 477 (2007), 359–378.
- S. J. Grossman. 1976. On the efficiency of competitive stock markets where traders have diverse information. *The Journal of Finance* 31, 2 (1976), 573–585.
- R. Hanson. 2003. Combinatorial information market design. *Information Systems Frontiers* 5, 1 (2003), 105–119.
- J. Hsu, Z. Huang, A. Roth, T. Roughgarden, and Z. S. Wu. 2014. Private matchings and allocations. In *46th ACM Symposium on Theory of Computing*.
- M. Kearns, M. Pai, A. Roth, and J. Ullman. 2014. Mechanism design in large games: Incentives and privacy. In *5th Conf. on Innovations in Theoretical Computer Science*.
- N. S. Lambert, J. Langford, J. W. Vaughan, Y. Chen, D. Reeves, Y. Shoham, and D. M. Pennock. 2015. An axiomatic characterization of wagering mechanisms. *Journal of Economic Theory* 156 (2015), 389–416.
- N. S. Lambert, J. Langford, J. Wortman, Y. Chen, D. Reeves, Y. Shoham, and D. M. Pennock. 2008. Self-financed wagering mechanisms for forecasting. In *9th ACM Conference on Electronic Commerce*.
- F. McSherry. 2009. Privacy Integrated Queries: An extensible platform for privacy-preserving data analysis. In *ACM SIGMOD International Conference on Management of Data*.
- D. M. Pennock, S. Lawrence, C. L. Giles, and F. A. Nielsen. 2002. The real power of artificial markets. *Science* 291 (2002), 987–988.
- C. Plott and K.-Y. Chen. 2002. Information aggregation mechanisms: Concept, design and field implementation. (2002). Cal Tech Social Science Working Paper 1131.
- P. M. Polgreen, F. D. Nelson, and G. R. Neumann. 2007. Using prediction markets to forecast trends in infectious diseases. *Clin. Infect. Dis.* 44, 2 (2007), 272–279.
- R. Roll. 1984. Orange juice and weather. *Am. Econ. Rev.* 74, 5 (1984), 861–880.
- L. J. Savage. 1971. Elicitation of personal probabilities and expectations. *J. Amer. Statist. Assoc.* 66, 336 (1971), 783–801.
- R. H. Thaler and W. T. Ziemba. 1988. Anomalies: Parimutuel betting markets: Race-tracks and lotteries. *J. of Economic Perspectives* 2, 2 (1988), 161–174.
- B. Waggoner, R. Frongillo, and J. Abernethy. 2015. A market framework for eliciting private data. In *Advances in Neural Information Processing Systems* 28.

APPENDIX

A.1. Omitted Proofs from Section 3

PROOF OF LEMMA 3.2. For each $i \in \mathcal{N}$,

$$\mathbb{E}[x_i(p_i, \omega)] = \frac{\alpha s(p_i, \omega) + \beta}{1 + \beta} - \beta \frac{1 - \alpha s(p_i, \omega)}{1 + \beta} = \alpha s(p_i, \omega) \quad (4)$$

and so

$$\mathbb{E}[\Pi_i(\mathbf{p}, \mathbf{m}, \omega)] = m_i \left(\alpha s(p_i, \omega) - \frac{\sum_{j \in \mathcal{N}} m_j \alpha s_j(p_j, \omega)}{\sum_{j \in \mathcal{N}} m_j} \right).$$

This is precisely the profit to bettor i in a WSWM with scoring rule αs . \square

PROOF OF THEOREM 3.4. For any $j \in \mathcal{N}$, consider the quantity $m_j x_j(p_j, \omega)$. From Equation 4, $\mathbb{E}[m_j x_j(p_j, \omega)] = m_j \alpha s(p_j, \omega)$. Additionally we can bound $m_j x_j(p_j, \omega) \in [-m_j \beta, m_j]$. Hoeffding's inequality then implies that with probability at least $1 - \delta$,

$$\left| \sum_{j \in \mathcal{N}} m_j \alpha s(p_j, \omega) - \sum_{j \in \mathcal{N}} m_j x_j(p_j, \omega) \right| \leq \|\mathbf{m}\|_2 (1 + \beta) \sqrt{\frac{\ln(2/\delta)}{2}}.$$

From the definition of the private wagering mechanism and Lemma 3.2, we then have that with probability at least $1 - \delta$, for any $i \in \mathcal{N}$,

$$\begin{aligned} |\Pi_i(\mathbf{p}, \mathbf{m}, \omega) - \mathbb{E}[\Pi_i(\mathbf{p}, \mathbf{m}, \omega)]| &= \frac{m_i}{\sum_{j \in \mathcal{N}} m_j} \left| \sum_{j \in \mathcal{N}} m_j \alpha s(p_j, \omega) - \sum_{j \in \mathcal{N}} m_j x_j(p_j, \omega) \right| \\ &\leq m_i \frac{\|\mathbf{m}\|_2}{\|\mathbf{m}\|_1} (1 + \beta) \sqrt{\frac{\ln(2/\delta)}{2}} \end{aligned}$$

as desired. \square

A.2. Omitted Proofs from Section 4

The proof of Lemma 4.2 makes use of the following technical lemma, which says that it is profitable in expectation to sell shares as long as the price remains above p^* or to purchase shares as long as the price remains below p^* . In this statement, q can be interpreted as the current market state and x as a new purchase (or sale); $C'(q^*)x - C(q+x) + C(q) \geq 0$ would then be the expected profit of a trader making this purchase or sale if $\omega \sim p^* = C'(q^*)$.

LEMMA A.1. *Fix any convex function C and any q^* , q , and x such that $q+x \geq q^*$ if $x \leq 0$ and $q+x \leq q^*$ if $x \geq 0$. Then $C'(q^*)x - C(q+x) + C(q) \geq 0$.*

PROOF. Since C is convex, the assumptions in the lemma statement imply that if $x \leq 0$ then $C'(q+x) \geq C'(q^*)$, while if $x \geq 0$ then $C'(q+x) \leq C'(q^*)$. Therefore, in either case $C'(q+x)x \leq C'(q^*)x$, and

$$C'(q^*)x - C(q+x) + C(q) \geq C'(q+x)x - C(q+x) + C(q) = D_C(q, q+x) \geq 0.$$

\square

PROOF OF LEMMA 4.2. From the definition of the market maker's loss, we can rewrite $\mathbb{E}[L_T(x_1, \dots, x_T, \eta_1, \dots, \eta_T, \omega)] = \sum_{t=1}^T \mathbb{E}[\pi_t]$ where π_t is the expected (over just the randomness in ω) loss of the market maker from the t th trade, i.e.,

$$\pi_t = C'(q^*)x_t - C(q'_t + x_t) + C(q'_t).$$

By definition of the target strategy \mathbf{s}^* and Lemma A.1, $\pi_t \geq 0$ for all t .

Consider a round t in which $|q'_t - q^*| \geq \gamma$. Suppose first that $q'_t \geq q^* + \gamma$, so a trader playing the target strategy would sell. By definition of \mathbf{s}^* , $x_t = -\min\{q'_t - q^*, k\} \leq -\gamma$. We can write

$$\begin{aligned}\pi_t &= C'(q^*)(x_t + \gamma) - C'(q^*)\gamma - C(q'_t + x_t) + C(q'_t - \gamma) - C(q'_t - \gamma) + C(q'_t) \\ &\geq -C'(q^*)\gamma - C(q'_t - \gamma) + C(q'_t) \\ &\geq -C'(q^*)\gamma - C(q^*) + C(q^* + \gamma) \\ &= D_C(q^* + \gamma, q^*) \geq \chi,\end{aligned}$$

where χ is defined as in the lemma statement. The first inequality follows from an application of Lemma A.1 with $q = q'_t - \gamma$ and $x = x_t + \gamma$. The second follows from the convexity of C and the assumption that $q'_t \geq q^* + \gamma$.

If instead $q'_t \leq q^* - \gamma$ (so a trader playing the target strategy would buy), a similar argument can be made to show that $\pi_t \geq D_C(q^* - \gamma, q^*) \geq \chi$.

Putting this all together, we have

$$\mathbb{E}[L_T(x_1, \dots, x_T, \eta_1, \dots, \eta_T, \omega)] = \sum_{t=1}^T \mathbb{E}[\pi_t] \geq \sum_{t=1}^T \chi \Pr(|q'_t - q^*| \geq \gamma)$$

as desired. The fact that $\chi > 0$ follows from the fact that it is the minimum of two Bregman divergences, each of which is strictly positive since C is nonlinear (and thus strictly convex) in the region around q^* and the arguments are not equal. \square