JMLR: Workshop and Conference Proceedings 29:373–387, 2013 ACML 2013

Aggregating Predictions via Sequential Mini-Trading

Mindika Premachandra MINDIKA.PREMACHANDRA@ANU.EDU.AU
Mark D. Reid MARK.REID@ANU.EDU.AU
The Australian National University & NICTA, Canberra, Australia

Editor: Cheng Soon Ong and Tu Bao Ho

Abstract

Prediction markets which trade on contracts representing unknown future outcomes are designed specifically to aggregate expert predictions via the market price. While there are some existing machine learning interpretations for the market price and connections to Bayesian updating under the equilibrium analysis of such markets, there is less of an understanding of what the instantaneous price in sequentially traded markets means. In this paper we show that the prices generated in sequentially traded prediction markets are stochastic approximations to the price given by an equilibrium analysis. We do so by showing the equilibrium price is a solution to a stochastic optimisation problem which is solved by stochastic mirror descent (SMD) by a class of sequential pricing mechanisms. This connection leads us to propose a scheme called "mini-trading" which introduces a parameter related to the learning rate in SMD. We prove several properties of this scheme and show that it can improve the stability of prices in sequentially traded prediction markets.

Keywords: Prediction Markets, Stochastic Optimisation, Belief Aggregation

1. Introduction

The main purposes of prediction markets are eliciting and aggregating beliefs over an unknown future outcome. Traders with different beliefs trade on contracts whose payoff's are related to the unknown future outcome and the market prices of the contracts are considered as the aggregated beliefs. While there are some existing machine learning interpretations for the market price under classical market analysis techniques (*i.e.*, via the Walrasian equilibrium price), there is less of an understanding of what the instantaneous price in sequentially traded markets means. The main purpose of this paper is to interpret the instantaneous price of sequential markets as stochastically minimizing the same objective as its Walrasian equilibrium, thus unifying the meaning of instantaneous and equilibrium prices via an optimisation point of view for markets satisfying certain conditions.

While machine learning aggregation models (such as boosting) perform aggregations over expert beliefs or machine learning models, prediction markets can be used to appeal to the wisdom of the crowds (Surowiecki, 2004) and may also be used to aggregate expert/model beliefs with human judgement. But unlike in machine learning algorithms where expert beliefs are readily available for aggregation, traders in a prediction market may have to be incentivized to participate in the market. Recently many connections between prediction markets and machine learning have been shown based on the Walrasian equilibrium analysis. Also sequentially traded markets have been shown to be closely related to follow-the-regularized-leader algorithms.

Our work relates the prices of sequentially traded markets with SMD and interprets its equilibrium as solving the same optimisation problem, thus establishing a connection between sequential and equilibrium market prices.

1.1. Overview

The closest work to ours is by Frongillo et al. (2012) who first interpreted sequential markets as performing a SMD and established a connection between sequential markets and equilibrium analysis using limiting conditions on sequential markets. Our interpretation of sequential markets as SMD is a variation of their result for a more broader class of trades and a more restrictive class of markets. We further show that the solution of the stochastic optimisation problem is the equilibrium price of the market (Theorem 1 in Section 3), thus establishing a connection between sequential markets and equilibrium without using limiting conditions.

Section 4 introduces "mini-trading" as a mechanism for implementing the learning rate parameter that is used for convergence in SMD. The idea behind "mini-trading" is to allow repeated but small-scale trader interactions as opposed to single but large-scale interactions in a sequential market. Then we show that "mini-trading" has desirable properties like more stable prices (Theorem 2) and bounded worst-case loss.

2. Definitions, Notation, and Background

We will use the following notation and conventions throughout the paper. Given vectors $x, y \in \mathbb{R}^N$, their inner product will be denoted $\langle x, y\rangle := \sum_{n=1}^N x_n y_n$ and we use $\mathbf{1} := (1, \dots, 1) \in \mathbb{R}^N$ to denote the vector of all ones. The set $\{1, \dots, N\}$ will be written as $[N]$ and $\Delta_N := \{\pi \in [0,1]^N : \langle \pi, \mathbf{1}\rangle = 1\}$ will denote the set of probability distributions over $[N]$. For $X \subseteq \mathbb{R}^N$ and functions $f : X \to \mathbb{R}$, the convex conjugate of f will be denoted $f^*(y) := \sup_{x\in X}\{\langle x, y\rangle - f(x)\}$. If f is also differentiable on the interior of X then its derivative at $x \in \mathbb{R}^N$ will be written $\nabla f(x) := \left(\frac{\partial f}{\partial x_1}(x), \dots, \frac{\partial f}{\partial x_N}(x)\right)$.

2.1. Contracts and Prices

We are interested in mechanisms for aggregating trader beliefs about a single future event with N of possible outcomes (e.g., who will win an election or horse race) and will label them $1, \dots, N$. We will assume that the outcomes are mutually exclusive and complete (i.e., exactly one of $1, \dots, N$ can occur). Prediction markets are markets where the "goods" that are traded are N types of contract – one for each of the N outcomes – that pay \$1 if outcome n occurs and nothing otherwise.[1]

The main aim of the mechanisms we consider is to produce a probability $\pi \in \Delta_N$ representing some "consensus belief" of the market about the future outcome. Such a mechanism is loosely analogous to an ensemble for a multiclass class probability estimation problem in the machine learning literature, where traders are playing the role of base predictors.

A common form of trading in markets is based on the assignment of prices to goods. Since the goods we consider are the contracts on future events, we will use $\Pi : \mathbb{R}^N \to \mathbb{R}$ to

1. For a more flexible framework with arbitrary payoff functions see (Abernethy et al., 2011).

denote a *pricing function* that assigns a price $\Pi(s)$ to each *contract bundle* $s \in \mathbb{R}^N$. Each component s_n represents the (possibly fractional) number of contracts bought for outcome $n \in [N]$ with $s_n < 0$ representing a sale of $-s_n$ contracts. We will use $\mathcal{P}$ to denote the set of pricing functions.

2.1.1. Examples

A very simple form of pricing function is one that assigns a fixed price $\pi_n \in [0,1]$ per unit contract for outcome n. The price for a bundle $s \in \mathbb{R}^N$ is then given by $\Pi^\pi(s) := \langle \pi, s \rangle$, which we will call a *fixed-price* pricing function.

Another type of pricing function we will use when we consider sequential market making in Secton 2.3.2 is a *cost function-based* pricing function[2]. We say a convex, differentiable function $C : \mathbb{R}^N \to \mathbb{R}$ is a *cost function* if $\nabla C(x) \in \Delta_N$ for all x. This condition ensures the prices set by a cost function are probabilities and cannot be arbitraged (Abernethy et al., 2011). We define a cost function-based pricing function by

$$\Pi^{C,x}(s) := C(x+s) - C(x).$$

The function C can be interpreted as assigning a dollar value to every position and the derived pricing function charges the difference between the value at position x (which denotes the total number of contracts sold so far) and the value if the position was moved to $x+s$. The value of $\nabla C(x)$ can be seen as the price for buying an infinitesimally small bundle at position x and is therefore called the *instantaneous price* at x.

One of the most well studied cost function-based pricing functions used in prediction markets is the Logarithmic Market Scoring Rule (LMSR) (Hanson, 2007). This is defined by the cost function $C(x) = b \log\left(\sum_{n=1}^N \exp(x_n/b)\right)$ where $b > 0$ is sometimes called the *liquidity parameter*. The LMSR has instantaneous prices $(\nabla C(x))_n = Z^{-1}\exp(x_n/b)$ where $Z = \sum_{n=1}^N \exp(x_n/b)$.

2.2. Traders and Demand

A trader's purchasing behaviour will be modelled through a *demand operator* d that reacts to pricing functions. Formally, $d : \mathcal{P} \to \mathbb{R}^N$ will return a contract bundle $d(\Pi) \in \mathbb{R}^N$ when given a pricing function $\Pi \in \mathcal{P}$. The returned bundle represents the contracts the trader wishes to buy when bundles are priced according to Π.

A common assumption about traders is that they are risk averse expected utility maximisers. That is, their demands are determined by their *belief* $p \in \Delta^N$, *wealth* $W \in \mathbb{R}$, and a concave[3] *utility function* $U : \mathbb{R} \to \mathbb{R}$ which measures the value $U(m)$ of \$$m$ to the trader. Buying a bundle of $s \in \mathbb{R}^N$ contracts reduces a trader's wealth by \$$\Pi(s)$ but will return \$$s_n$ should outcome n occur. If the trader believes each outcome n will occur with probability p_n then her *expected utility* for owning the bundle s is $\mathbb{E}_{n\sim p}[U(W - \Pi(s) + s_n)]$. We will say a trader is *loss avoiding* if, no matter which outcome obtains, the trader will always have wealth that is lower bounded by some value (V) that usually depends on

2. See Abernethy et al. (2011) for a detailed, axiomatic approach.

3. Traders with concave utility functions U are called *risk averse* since their utility for a guaranteed \$$\frac{1}{2}(x+y)$ is always larger than their expected utility for a coin toss resulting in \$$x$ or \$$y$.

their utility function. The set of allowable bundles for a loss avoiding trader is therefore $S_{W,p}(\Pi) := \{s \in \mathbb{R}^N : W - \Pi(s) + s_n \geq V \ \forall n \in [N]\}$. A demand operator

$$d^{\mathrm{MEU}}(\Pi) := \underset{s \in S_{W,p}(\Pi)}{\arg\max}\ \mathbb{E}_{n\sim p}\left[U(W - \Pi(s) + s_n)\right]$$

that returns the optimal expected loss avoiding contract bundle for a given pricing function is called a *maximum expected utility* (MEU) demand operator. Because of their importance in equilibrium analysis (discussed below), we will denote the set of MEU demand operators for loss avoiding traders by $\mathcal{D}^{\mathrm{MEU}}$.

2.2.1. Examples

One of the most common examples of MEU demand operators is for logarithmic utility $U^{\log}(m) := \log m$, which always requires wealth to be non-negative ($m > 0$). Log utility based traders make (positive) demands $d_n(\Pi^\pi) = W.(p_n/\pi_n)$ (spending entire wealth W) which also makes intuitive sense since demands are proportional to the belief/price ratio.

2.3. Markets

There are two broad classes of market mechanisms we consider: equilibrium markets and sequential markets. In equilibrium markets, traders buy and sell contracts with each other based on a fixed-price pricing function. If the price is such that all traders' demands are satisfied then the market is said to be in equilibrium and the price can be interpreted as consensus belief about the outcome probabilities (Easley and Jarrow, 1983; Wolfers and Zitzewitz, 2006). In sequential markets, traders interact sequentially with a "market maker" that sets the contract prices and updates them after each interaction. Each price generated by the market maker can also be viewed as a summary of the beliefs held by the traders who have interacted with the market maker (Frongillo et al., 2012).

Since demand operators are a general way of describing trader behaviour, we model *market* for both scenarios as a set $\mathcal{M} \subset \mathcal{D}$ of demand operators. We further assume that the demand operators for a market are drawn i.i.d. from some distribution[4] σ over $\mathcal{D}$ and refer to the distribution σ as a *stochastic market*.

As noted in the introduction, our main contributions concern the relationship between these two types of market under the assumption of stochastically drawn demands: Theorem 1 relates equilibrium prices to those obtained from sequential mechanisms acting upon the same stochastic market while the "mini-trading" mechanism introduced in Section 4 can be seen as a implementation of stochastic mirror descent with properties (*e.g.*, stability, bounded loss) that make it desirable for finding equilibrium prices.

2.3.1. Equilibrium Market Mechanisms

Classically, the *(Walrasian) equilibrium price* for a market is a fixed price at which there is no excess demand for any good (Varian, 2009). We define the *equilibrium price* for a stochastic market σ to be the price $\pi^*_\sigma \in \Delta_N$ such that the expected demand of the market

4. We want to avoid delving into measure theoretic details so assume throughout that appropriate algebras are defined and measurability conditions are met.

in response to the fixed-price pricing function $\Pi^{\pi^*_\sigma}$ is equal to the total spendings of the traders. Formally, the equilibrium price satisfies

$$\mathbb{E}_{d\sim\sigma}\left[d\left(\Pi^{\pi^*_\sigma}\right)\right] = \mathbb{E}_{d\sim\sigma}\left[\Pi^{\pi^*_\sigma}\left(d\left(\Pi^{\pi^*_\sigma}\right)\right)\right].\mathbf{1} \tag{1}$$

This is similar to the definition used by Storkey et al. (2012) where the expected demand for any outcome equals the total wealth of the traders, assuming that traders spend their entire wealth for trading. If instead traders spend \$0, then we get Frongillo et al. (2012)'s definition that the expected demand equals zero. Also if the expected demand for any outcome equals any $N \in \mathbb{R}$ such that $\mathbb{E}_{d\sim\sigma}\left[d\left(\Pi^{\pi^*_\sigma}\right)\right] = N.\mathbf{1}$, then we have that the total spendings also equal N, satisfying the equilibrium condition.

$$\mathbb{E}_{d\sim\sigma}\left[\Pi^{\pi^*_\sigma}\left(d\left(\Pi^{\pi^*_\sigma}\right)\right)\right] = \langle \pi^*_\sigma, \mathbb{E}_{d\sim\sigma}\left[d\left(\Pi^{\pi^*_\sigma}\right)\right]\rangle = \sum_{n=1}^{N}(\pi^*_\sigma)_n\left(\mathbb{E}_{d\sim\sigma}\left[d_n\left(\Pi^{\pi^*_\sigma}\right)\right]\right) = N \tag{2}$$

Equilibrium prices for a non-stochastic market with K traders with demands $d^1, \ldots, d^K$ can be obtained via a distribution that puts mass $1/K$ over those K demands. The question of the existence of equilibrium prices is complex but resolved. For our purposes it suffices for trader utilities to be concave (Arrow and Debreu, 1954). As such, we will restrict our attention to demand operators in $\mathcal{D}^{\text{MEU}}$ – *i.e.*, loss avoiding MEU demands – whenever discussing equilibria.

Several recent papers have studied the equilibrium prices of prediction markets consisting of so-called "artificial" traders with known utilities and beliefs derived from predictions of machine learning algorithms. Storkey et al. (2012) and Storkey (2011) show that the equilibrium prices of these artificial markets replicate well-known aggregations techniques from machine learning: weighted-means used in boosting and random forests (derived from log utilities); entropic means or α-mixtures (from iso-elastic utilities), weighted-medians (linear utilities), product-model combinations such as log opinion pools or geometric means (exponential negative decay utilities). They also obtain novel aggregations from markets consisting of traders with differing utility functions that have connections to the minimisation of divergence-based distances from the optimisation literature (Ben-Tal et al., 1989; Amari, 2007). Barbu and Lay (2011) also use artificial prediction markets to obtain the weighted-mean aggregation and implement kernel methods, logistic regressions. Beygelzimer et al. (2012) show connections to Bayesian model averaging and Bayesian weight updates for Kelly bettors (log utility based traders) which Storkey et al. (2012) generalize to the class of iso-elastic utility based traders.

2.3.2. Sequential Market Mechanisms

In a sequential market, demand operators are drawn i.i.d. from the stochastic market. At step t, $d^t \sim \sigma$ is presented to a market maker that offers a pricing function Π^t, accepts the bundle $s^t = d^t(\Pi^t)$, and updates its pricing function to Π^{t+1} in preparation for the next step. Formally, the updates are represented by *sequential pricing mechanism* $\Pi(\cdot|s^1, \ldots, s^T) \in \mathcal{P}$ that is a function from histories of bundle purchases $s^1, \ldots, s^T$ for any number of trades T to pricing functions.

A natural constraint on sequential pricing mechanisms is that it is *path independent*, that is, $\Pi(s+s'|s^1,\dots,s^t) = \Pi(s|s^1,\dots,s^t,s') + \Pi(s'|s^1,\dots,s^t)$ for all bundles s, s' and histories $s^1,\dots,s^t$. Theorem 1 of Abernethy et al. (2011) shows that all path independent sequential pricing mechanisms are necessarily the cost function-based pricing functions described in Section 2.1.1. Specifically, $\Pi^{C,x^0}(s|s^1,\dots,s^t) = C(x^t+s) - C(x^t)$ where $x^t = x^0 + \sum_{i=1}^t s^i$.

Sequential cost function-based market markets are known to have a bounded loss for the market-maker (see Chen and Pennock (2007)) like regret analysis in online learning algorithms. Chen and Vaughan (2010) show how such markets can be used to implement follow-the-regularized-leader and achieve no-regret learning. Thus the role of the market-maker becomes much like an algorithm for learning from expert advice. Abernethy et al. (2011) notes further connections between the price update mechanisms in sequential markets and weight update mechanisms in follow-the-regularized-leader algorithms.

For the purposes of our Theorem 1 in Section 3 and the "mini-trading" mechanism in Section 4 we now introduce a pricing mechanism that is somewhere in between the cost function-based and fixed-price mechanisms which we call a *hybrid pricing mechanism* and denote by $\Pi^{\nabla C,x^0}$. Like the cost function-based mechanisms, this pricing mechanism sets prices based on a cost function C and position x that is initialised to x^0. Unlike the cost function-based mechanism, however, the bundle prices are set by a fixed-price function $\Pi^{\nabla C(x)}$. That is, $\Pi^{\nabla C,x^0}(s|s^1,\dots,s^t) := \langle \nabla C(x^t), s\rangle$ where $x^t = x^0 + \sum_{i=1}^t s^i$. We call this a hybrid mechanism since it prices bundles using fixed prices derived from a cost function but updates the price like the cost function-based mechanism.

3. Relating Equilibrium and Sequential Market Prices

In Theorem 1 below we show that, under certain conditions, the instantaneous prices of sequential markets are approximations of its Walrasian equilibrium, thus unifying the meaning of instantaneous and equilibrium prices. We do this by interpreting the stochastic price update mechanism of sequential cost function based markets as performing a stochastic mirror descent (SMD) and showing that the direct solution to this optimisation problem is the Walrasian equilibrium of the market.

This result can be seen as a refinement or variation of the two results by Frongillo et al. (2012). In their Theorem 1, it was shown that the stationary distribution of the sequential market making process was equal to the equilibrium price in the limit of infinite market liquidity, and in their Theorem 2, it was shown that sequential cost function-based market making was a stochastic mirror descent. Our result differs from their two as, although it establishes the similar correspondence with SMD, it does so for a broader range of demand operators and a more restrictive class of cost functions. Furthermore, our result establishes a direct connection between sequential and equilibrium prices without needing to appeal to infinite liquidity limits.

3.1. Background and Assumptions

Before stating the theorem, we first recall the form of stochastic mirror descent algorithms. We then introduce our *potential-based* assumption on traders' demands and a property of cost function-based markets called *liquidity insensitivity* which makes it possible to interpret the stochastic price update as a stochastic mirror descent.

3.1.1. Stochastic Optimisation and Stochastic Mirror Descent

In stochastic optimisation problems there is some convex *objective* $\Phi : X \to \mathbb{R}$ to be minimised that is only accessible through unbiased samples of its value and gradient. That is, $\Phi(x) = \mathbb{E}_{\omega\sim\mu}[F^\omega(x)]$ for some collection of convex functions $F^\omega : X \to \mathbb{R}$ for which we can compute the gradients $\nabla F^\omega(x)$ and distribution μ.

One very successful method for solving a problem of this type is *stochastic mirror descent* (Nemirovski et al. (2009)). This algorithm starts with some initial point $x^0 \in X$ as a candidate solution and iteratively improves it using samples of the objective function. Given a learning rate $\eta \in (0, 1]$ and a strictly convex regulariser $R : X \to \mathbb{R}$, the update step is

$$x^{t+1} = \arg\min_{x\in X}\{\eta\langle x, \nabla F^\omega(x^t)\rangle + D_R(x; x^t)\} = \nabla R^*(\nabla R(x^t) - \eta.\nabla F^\omega(x^t))$$

where, at each step, ω are drawn i.i.d. from μ. The first form of the update involving the $\arg\min$ has an intuitive interpretation as a trade off between taking a step in the steepest descent direction and staying close to the previous solution, as measured by the Bregman distance $D_R(p; q) = R(p) - R(q) - \langle\nabla R(q), (p - q)\rangle$ generated by R. The other form of the update step is derived using convex duality and is the one we show is equivalent to sequential price updates.

3.1.2. Potential-Based Demands and Liquidity Insensitive Cost Functions

In order to relate the sequential price update mechanism to a stochastic mirror descent algorithm, we assume that the demand operators d are linearly related to the negative of the gradient of some convex function. To this end, we say a trader is *potential-based* if there exist functions $F : \Delta_N \to \mathbb{R}^N$ and $f : \Delta_N \to \mathbb{R}$ such that the trader's demand operator d satisfies

$$d(\Pi^\pi) = -\nabla F(\pi) + f(\pi).\mathbf{1} \tag{3}$$

for any fixed-price pricing function Π^π. We will say a stochastic market $\mathcal{M}$ is *potential-based* if every trader that can be drawn from $D_\mathcal{M}$ is potential-based.

As examples for this assumption, consider the MEU demands of some popular utility functions presented in Table 1, where $D_{KL}(p; q) = \sum_n p_n. \ln(p_n/q_n)$ is the KL divergence and $D^P_\beta(p; q) = \frac{\sum_n q_n.(p_n/q_n)^\beta - 1}{\beta(\beta-1)}$ for $\beta > 0$ are the power divergences (Jose et al., 2008).

We note that the set of demand operators that are potential-based in our sense is strictly larger than those considered by Frongillo et al. (2012). In particular, the exponential negative decay utilities are not in the class they consider.

The other property we require is a constraint on the cost function that generates prices for our hybrid mechanism. This constraint – called *liquidity insensitivity* – requires the following: for all positions $x \in \mathbb{R}^N$ and $\alpha \in \mathbb{R}$ the cost function C satisfies $C(x + \alpha.\mathbf{1}) = C(x) + \alpha$. An immediate and important consequence of this property is that the prices given by the C are invariant to purchases of equal quantities of all contracts. That is, $\nabla C(x + \alpha.\mathbf{1}) = \nabla C(x)$ for all $x \in \mathbb{R}^N$ and $\alpha \in \mathbb{R}$. This property has been studied in the context of prediction markets by Othman et al. (2010) and recently given an axiomatic characterisation by Li and Vaughan (2013).

Utility	MEU Demand	Potential-based Representation
Log Utility: $U(w) = \ln(w)$	$d_n(\Pi^\pi) = W\frac{p_n}{\pi_n}$	$F^d(\pi) = W.D_{KL}(p;\pi)$, $f^d(\pi) = 0$
Exp. Neg. Utility ($r > 0$): $U(w) = -\exp(-r.w)$	$d_n(\Pi^\pi) = \frac{1}{r}\ln(p_n/\pi_n) + W + \frac{1}{r}.D_{KL}(\pi;p)$	$F^d(\pi) = \frac{1}{r}D_{KL}(\pi;p)$, $f^d(\pi) = W + \frac{1}{r} + \frac{1}{r}D_{KL}(\pi;p)$
Iso-elastic Utility ($\beta > 0$): $U(w) = \frac{1}{\beta-1}\left(w^{\frac{\beta-1}{\beta}} - 1\right)$	$d_n(\Pi^\pi) = \frac{W}{Z}\left(p_n/\pi_n\right)^\beta$ for $Z = \sum_n \pi_n.\left(p_n/\pi_n\right)^\beta$	$F^d(\pi) = \frac{W}{Z}\beta.D^P_\beta(p;\pi)$, $f^d(\pi) = 0$

Table 1: Potential-based MEU demands for beliefs $p \in \Delta^N$, wealth W and prices $\pi \in \Delta^N$

3.2. Sequential Prices Approximate Equilibrium Prices

We can now state and prove our theorem. This result shows that under the assumption of potential-based demands and liquidity insensitive cost functions, the sequential prices generated by a hybrid pricing mechanism are approximating the solution of a stochastic minimisation problem defined by the traders' demands. Furthermore, the solution of the problem is necessarily the equilibrium price for the stochastic market generating the demand operators.

Theorem 1 *Suppose σ is a potential-based stochastic market. Traders drawn from σ with demand operators $d^1, \ldots, d^T$ interact with a hybrid sequential pricing mechanism $\Pi^{\nabla C, x^0}$ generating sequences of positions $x^t = x^0 + \sum_{i=1}^t s^i$, prices $\pi^t = \nabla C(x^t)$, and bundles $s^t = d^t(\Pi^{\pi^t})$ for $t = 1, \ldots, T$. Then,*

1. *The generated price sequence $\pi^1, \ldots, \pi^T$ is exactly the update sequence for a stochastic mirror descent of the function $\Phi(\pi) := \mathbb{E}_{d\sim\sigma}\left[F^d(\pi)\right]$ using regulariser $R = C^*$.*

2. *Any price π^* minimising $\Phi(\pi)$ is an equilibrium price for the stochastic market σ.*

Proof Since each d drawn from σ is potential-based by assumption, there exists convex functions F^d and functions f^d such that $d(\Pi^\pi) = -\nabla F^d(\pi) + f^d(\pi).\mathbf{1}$. Therefore,

$$\pi^{t+1} = \nabla C(x^t + s^t) = \nabla C(x^t + d^t(\Pi^{\pi^t})) = \nabla C(x^t - \nabla F^{d^t}(\pi^t) + f^{d^t}(\pi^t).\mathbf{1}).$$

However, since C is liquidity insensitive we have that $\pi^{t+1} = \nabla C(x^t - \nabla F^{d^t}(\pi^t))$. As C is differentiable and convex its dual R^* is also and their derivatives satisfy $\nabla C = \nabla R^*$ and $(\nabla C)^{-1} = \nabla R$ (see, *e.g.*, Boyd and Vandenberghe (2004)) and so $x^t = (\nabla C)^{-1}(\pi^t) = \nabla R(\pi^t)$. Thus,

$$\pi^{t+1} = \nabla R^*(x^t - \nabla F^{d^t}(\pi^t)) = \nabla R^*(\nabla R(\pi^t) - \nabla F^{d^t}(\pi^t))$$

which is precisely the stochastic mirror descent update with regulariser R, objective $\Phi(\pi) = \mathbb{E}_{d\sim\sigma}\left[F^d(\pi)\right]$ and step size $\eta = 1$, establishing the first part of the theorem.

For the second part of the theorem, consider the direct solution to the optimisation problem: $\min_\pi \Phi(\pi) = \mathbb{E}_{d\sim\sigma}[F^d(\pi)]$ subject to $\pi \in \Delta_N$. The Lagrangian for this problem is

$\mathcal{L} = \mathbb{E}_{d\sim\sigma}\left[F^d(\pi)\right] + \lambda(\sum_{n=1}^{N}\pi_n - 1)$. The KKT conditions require a solution π^* to satisfy $\nabla\mathcal{L}(\pi^*) = 0$. Because demands are potential-based, we have that $\nabla F^d(\pi^*) = -d(\Pi^{\pi^*}) + f^d(\pi^*).\mathbf{1}$ and so

$$\mathbb{E}_{d\sim\sigma}\left[-d(\Pi^{\pi^*}) + f^d(\pi^*).\mathbf{1}\right] + \lambda.\mathbf{1} = 0$$

and therefore $\mathbb{E}_{d\sim\sigma}\left[d(\Pi^{\pi^*})\right] = (\lambda + \mathbb{E}_{d\sim\sigma}\left[f^d(\pi^*)\right]).\mathbf{1}$.

This final expression is precisely the condition for the Walrasian equilibrium since the total demands for each contract are equal. Thus the solution of the stochastic optimisation problem is an equilibrium price for the market. ■

The correspondence between sequential hybrid pricing and stochastic mirror descent established by the above theorem is not perfect since the price update mechanism has a learning rate parameter η fixed to 1. In the next section, we introduce a market mechanism designed to incorporate a learning rate parameter in a stochastic price update, thus implementing a stochastic mirror descent model with a built-in learning rate parameter. For small values of the learning rate parameter, we show that the hybrid pricing mechanism closely approximates the traditional cost function-based mechanism.

4. Mini-trading

One difficulty with interpreting the instantaneous prices of a sequential market maker is the high variability of prices that can occur when the market maker has low liquidity relative to the wealth of the traders in the market. We propose a simple idea to combat this problem and thereby help to make prediction prices more interpretable. Instead of allowing traders to purchase large bundles of contracts, constrained only by their wealth, we modify the market maker's pricing function in a way that effectively limits how much a single trade can move the market price. Given a parameter $m \in (0, 1]$, we do so by simultaneously scaling down each bundle purchased by m, scaling up the cost of purchasing it by $\frac{1}{m}$, and allowing traders to trade a factor of $\frac{1}{m}$ more times. That is, we allow repeated but small-scale trader interactions as opposed to single but large-scale interactions. We first introduce this *mini-trading* scheme for general pricing functions and then focus our attention on cost function-based and hybrid pricing mechanisms. We show that this approach does in fact improve price stability and relate its worst-case loss to the loss of the original pricing function in the case of cost function-based markets.

Formally, given a pricing mechanism Π, its *mini-trading* version for $m \in (0, 1]$ is the function

$$\Pi_m(s|s^1, \ldots, s^T) = \frac{1}{m}\Pi(ms|ms^1, \ldots, ms^T). \tag{4}$$

In the case of mechanisms Π^π that use a fixed price π to define the bundle cost $\langle\pi, s\rangle$ we see immediately that the mini-trade transformation does not affect the cost since $\frac{1}{m}\langle\pi, ms\rangle = \langle\pi, s\rangle$ for all s, π and m. In particular, suppose $\Pi = \Pi^{\nabla C,0}$ is a hybrid mechanism with initial position 0. The mini-trade transformed version Π_m will assign the same prices for bundles as Π when the two mechanisms are at the some position x. However, it is important to note that the positions of the two mechanisms after t steps will be $x^t = \sum_{i=1}^{n} s^i$ for the original mechanism and mx^t for the mini-trade version. That is, the position is also scaled by m.

To relate mini-trading to the optimisation perspective developed in the previous section, we now consider the price update mechanism for a mini-trade transformed hybrid mechanism. Letting $x^t = \sum_{i=1}^t s^i$, the price at step $t+1$ is $\pi^{t+1} = \nabla C(mx^{t+1}) = \nabla C(mx^t + ms^t)$. However, since $\pi^t = \nabla C(mx^t)$ we have that $mx^t = (\nabla C)^{-1}(\pi^t) = \nabla R(\pi^t)$. By following the rest of the argument in the proof of Theorem 1 we get

$$\pi^{t+1} = \nabla R^*(\nabla R(\pi^t) - m.\nabla F^{d^t}(\pi^t))$$

which is precisely the stochastic mirror descent update with $\eta = m$ and regulariser $R = C^*$.

4.1. Properties

We now discuss some more properties and advantages of mini-trading. After introducing an appropriate definition of price stability, we show that the price stability in a mini-trade mechanism is better than in the original mechanisms.

4.1.1. Price Stability

Given two pricing mechanisms Π and Π', we will say that Π' has better *price stability* than Π if, whenever the mechanisms are at positions that offer the same price, the effect of updating the mechanism will change the price less for Π' than for Π. Formally, let $S = (s^1, \ldots, s^t)$ and $R = (r^1, \ldots, r^{t'})$ be trade histories and $s, r \in \mathbb{R}^N$ be bundles such that $\Pi(s|S) = \Pi'(r|R)$. Then Π' has better price stability than Π if

$$|\Pi(s|S, s) - \Pi(s|S)| > \left|\Pi'(r|R, r) - \Pi'(r|R)\right|. \tag{5}$$

The following theorem shows that mini-trade transformations of hybrid pricing mechanisms always improve price stability.

Theorem 2 *Let $\Pi^{\nabla C,0}$ be a hybrid pricing mechanism with initial price 0. Then mini-trade transformed version $\Pi_m^{\nabla C,0}$ with parameter has better price stability than $\Pi^{\nabla C,x^0}$ for all $m \in (0,1)$.*

Proof For simplicity we will let $\Pi := \Pi^{\nabla C,x^0}$ and $\Pi_m := \Pi_m^{\nabla C,x^0}$. First observe that to meet the pre-condition for price stability – *i.e.*, the two mechanisms to return identical costs for the same bundle – we need to find histories $S = (s^1, \ldots, s^t)$ and $R = (r^1, \ldots, r^{t'})$ and bundles s and r such that

$$\Pi(s|S) = \Pi_m(r|R) \iff \langle \nabla C(x^S), s\rangle = m^{-1}\langle \nabla C(m.x^R), m.r\rangle$$

where $x^S = \sum_{i=1}^t s^i$ and $x^R = \sum_{i=1}^{t'} r^i$. This is clearly satisfied by $R = m^{-1}.S$ and $r = s$ for any S and s. In this case, the pricing function presented by both mechanisms will be $\langle \pi, \cdot\rangle$ where $\pi = \nabla C(x^S) = \nabla C(m.x^R)$.

Due to the convexity of C we know that ∇C is monotonic – *i.e.*, $\langle \nabla C(x) - \nabla C(y), x - y\rangle \geq 0$ for all x, y. Therefore,

$$\langle \nabla C(x^S + s) - \nabla C(x^S), s\rangle \geq 0 \quad \text{and} \quad \langle \nabla C(x^S + m.s) - \nabla C(x^S), m.s\rangle \geq 0 \tag{6}$$

and thus, $\Pi_m(r|R,r) - \Pi_m(r|R) \geq 0$ and $\Pi(s|S,s) - \Pi(s|S) \geq 0$ which means we can remove the absolute value signs in the definition of price stability in (5).

Now consider

$$\begin{aligned}
&(\Pi(s|S,s) - \Pi(s|S)) - (\Pi_m(r|R,r) - \Pi_m(r|R)) \\
=&(\langle \nabla C(x^S + s) - \nabla C(x^S), s\rangle) - (\langle m^{-1}\nabla C(m.x^R + m.r) - m^{-1}\nabla C(m.x^R), m.r\rangle) \\
=&\langle \nabla C(x^S + s) - \nabla C(x^S + m.r), s\rangle \\
=&(1-m)^{-1}\langle \nabla C(x^S + s) - \nabla C(x^S + m.r), (1-m)s\rangle \geq 0
\end{aligned}$$

where the last inequality is once again due to the convexity of C and the last equality is because $\nabla C(m.x^R) = \nabla C(x^S)$. Since the quantities in the difference in top of that chain are both positive by (6), we have established the result. ■

4.1.2. Worst-case Loss

So far our analysis has been based on price updates in sequential markets. Now we turn to the profit/loss analysis in these markets. Typically, we assume that traders make demands so as to maximise their expected profit or utility. So trader's have an incentive to participate in the market. But the automated market-maker who has the objective of eliciting and aggregating trader's beliefs, may end up bearing a loss in the market and this has been viewed as a "price" that a market-maker has to pay in return for "learning" from the trader's beliefs. But when designing such a market, there might be a maximum price that a market-maker is allowed to lose in return for "learning".

In sequential markets with cost function-based pricing, the worst-case loss of the market-maker is defined as,

$$\max_{n \in 1,\dots,N} (x_n^T - x_n^0) - \Sigma_{t=1}^T C(x^t) - C(x^{t-1})$$

For strictly convex cost functions, this worst case loss has been shown to be upper bounded (Hanson, 2007; Chen and Pennock, 2007). For example, the LMSR cost function has worst-case loss bounded by $b.\ln(N)$ if the market-maker set initial prices π^0 as a uniform distribution.

But in a market with hybrid pricing, the worst-case loss of the market-maker would be,

$$\max_{n \in 1,\dots,N} (y_n^T - y_n^0) - \Sigma_{t=1}^T \langle \nabla C(y^{t-1}), (y^t - y^{t-1})\rangle$$

(denoting the contract position by y since traders may react differently). Also by convexity of the cost function, we have that $C(x+s) - C(x) \geq \langle \nabla C(x), s\rangle$, which implies that the sale of a bundle s using hybrid pricing (with prices fixed at $\nabla C(x)$) yields less income to the market-maker than selling the same bundle using cost function-based pricing. Thus we have,

$$\forall n, (y_n^T - y_n^0) - \Sigma_{t=1}^T C(y^t) - C(y^{t-1}) \leq (y_n^T - y_n^0) - \Sigma_{t=1}^T \langle \nabla C(y^{t-1}), (y^t - y^{t-1})\rangle$$

which suggests that the worst-case loss of a hybrid pricing mechanism is higher than the worst-case loss of a cost function-based pricing mechanism.

Due to the definition of the cost function-based pricing, for infinitesimal demands we also have $\lim_{s\to 0} C(x+s) - C(x) = \langle \nabla C(x), s\rangle$. This means for each demand $s^t = y^t - y^{t-1}$, for the limit $s^t \to 0$, the worst-case loss analysis for a market with hybrid pricing is the same as in a market with cost function-based pricing.

$$\forall n, \lim_{s^t \to 0} (y_n^T - y_n^0) - \Sigma_{t=1}^T \langle \nabla C(y^{t-1}), (y^t - y^{t-1})\rangle \approx (y_n^T - y_n^0) - \Sigma_{t=1}^T C(y^t) - C(y^{t-1})$$

In a mini-trade transformed hybrid mechanism, since the mechanism scales down the original demand s to $m.s$, we have that $\lim_{m\to 0} m.s \to 0$. So we conclude that a mini-trade transformed hybrid mechanism has worst-case loss similar to a cost function-based pricing mechanism for the limit $m \to 0$.

4.2. Experimental results

Although our contribution in this paper is primarily a theoretical one, we briefly present a simple experiment that demonstrates the convergence and price stability properties of mini-trading. For this purpose we use log utility based traders d with wealths W^d and beliefs p^d who make MEU demands for fixed prices as shown in Table 1. These traders participate, 1) in an original market with the LMSR cost function with $b = 1$ and a hybrid pricing scheme, 2) in a mini-trade transformed market. We set initial prices as a uniform distribution by setting initial positions to be 0. Using equilibrium analysis as in Theorem 1 we have that the equilibrium price $\mathbb{E}_{d\sim D}[W^d.p^d]$ (*i.e.*, the wealth-weighted mean of the traders' beliefs) minimises $f(\pi) = \mathbb{E}_{d\sim D}[W^d.D_{KL}(p^d;\pi)]$ (see Storkey et al. (2012) and Amari (2007) for direct derivations).

For the experiment we use a binary outcome market and 100 traders who participate once each in the original market and $1/m$ times in the mini-trade market. For the mini-trade market, we set $m = 0.1$. Traders' beliefs p^d are drawn from a normal distribution with mean 0.75 and standard deviation 0.2. All traders have unit wealth $W^d = 1$ so the equilibrium price is simply the mean belief. The instantaneous prices of the markets, averaged over 30 simulations for the same traders is given in Figure 1. For each simulation, the trading order is determined by a random permutation order. From the price history, we see a better convergence of instantaneous prices to the mean belief (equilibrium price) in mini-trades. Since the price fluctuation is less in the mini-trade market, we also see that mini-trading creates more stable prices.

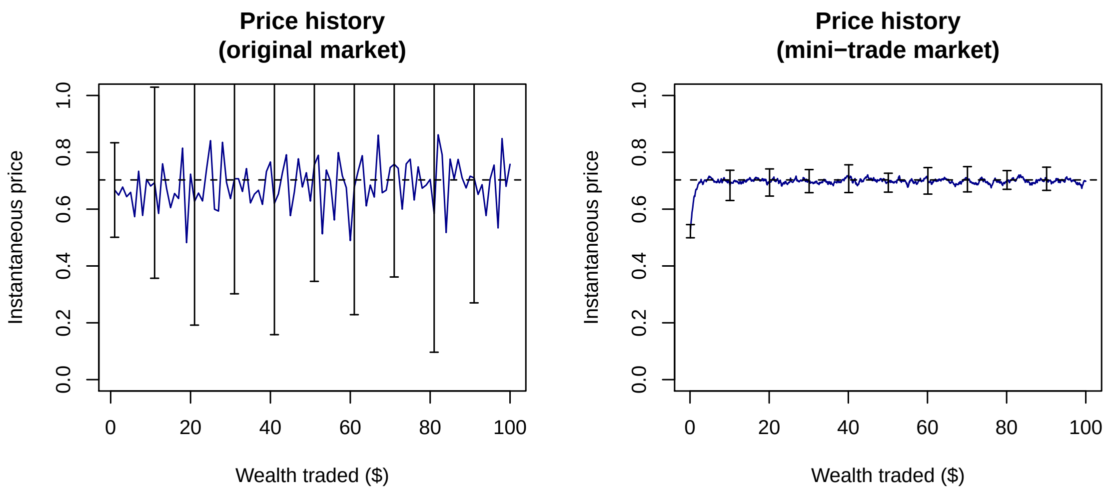

Figure 1: The price history of full-trade and mini-trade markets, averaged over 30 simulations. The dashed line marks the equilibrium price and bars show standard deviation.

5. Conclusion and future work

In this paper we connect the instantaneous price updates in sequential markets and the equilibrium price in equilibrium markets as solving the same optimisation problem. We explore the convergence of instantaneous prices in sequential markets to the equilibrium price via mini-trading, which has desirable properties of price stability and bounded loss for the market-maker, thus mini-trading can used as a framework to produce aggregations that are closer to the equilibrium price using the sequential market approach. As a result of convergence, we suspect that mini-trading also generates profits that are closer to the profits obatined from an equilibrium market. This is desirable because Beygelzimer et al. (2012) and Storkey et al. (2012) interpret equilibrium wealth updates as Bayesian weight updates in an online learning setup. Current experiments support this claim and we leave this analysis as future work.

We have assumed that traders don't adopt strategic behaviour (Chen et al., 2007; Dimitrov and Sami, 2008) given the opportunity for repeated trades in mini-trading, since each trade only results in a "small-scale" impact on the market prices and trades occur in a randomised order. We hope to prove this as future work.

Acknowledgments

This work was supported by the Australian Research Council (ARC). The second author is supported by an ARC Discovery Early Career Research Award (DE130101605). NICTA is funded by the Australian Government and as an ARC ICT Centre of Excellence.

References

J. Abernethy, Y. Chen, and J. W. Vaughan. An Optimization-Based Framework for Automated Market-Making. In *Proceedings of the 12th ACM conference on Electronic Commerce (EC'11)*, 2011.

S. Amari. Integration of stochastic models by minimizing α-divergence. *Neural Computing*, 19(10):2780–2796, 2007.

K.J. Arrow and G. Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22, 1954.

Adrian Barbu and Nathan Lay. An Introduction to Artificial Prediction Markets for Classification. *Journal of Machine Learning Research*, 2011.

Aharon Ben-Tal, Abraham Charnes, and Marc Teboulle. Entropic Means. *Journal of Mathematical Analysis and Applications*, 139(2):537–551, 1989.

Alina Beygelzimer, John Langford, and David Pennock. Learning performance of prediction markets with kelly bettors. 2012.

Stephen P. Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

Yiling Chen and David Pennock. A Utility Framework for Bounded-Loss Market Makers. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*, 2007.

Yiling Chen and Jennifer Wortman Vaughan. A new understanding of prediction markets via no-regret learning. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC'10)*, 2010.

Yiling Chen, Daniel M. Reeves, David M. Pennock, Robin D. Hanson, Lance Fortnow, and Rica Gonen. Bluffing and strategic reticence in prediction markets. In *Proceedings of the 3rd International Conference on Internet and Network Economics (WINE '07)*, pages 70–81, 2007.

Stanko Dimitrov and Rahul Sami. Non-myopic strategies in prediction markets. In *Proceedings of the 9th ACM conference on Electronic Commerce (EC '08)*, 2008.

David Easley and Robert A. Jarrow. Consensus beliefs equilibrium and market efficiency. *The Journal of Finance*, 38(3), 1983.

Rafael Frongillo, Nicolás Della-Penna, and Mark D. Reid. Interpreting prediction markets: A stochastic approach. In *Neural Information Processing Systems*, 2012.

Robin Hanson. Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation . *Journal of Prediction Markets*, 1(1), 2007. (Appears in 2002 as a technical report).

Victor Richmond R. Jose, Robert F. Nau, and Robert L. Winkler. Scoring rules, generalized entropy, and utility maximization. *Operations Research*, 56(5):1146–1157, 2008.

Xiaolong Li and Jennifier Wortman Vaughan. An axiomatic characterization of adaptive-liquidity market makers. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC'13)*, 2013.

Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.

Abraham Othman, Tuomas Sandholm, David M. Pennock, and Daniel M. Reeves. A Practical Liquidity-Sensitive Automated Market Maker. In *Proceedings of the 11th ACM conference on Electronic commerce (EC'10)*, 2010.

Amos J. Storkey. Machine learning markets. In *Proceedings of Artificial Intelligence and Statistics (AISTATS)*, 2011.

Amos J. Storkey, Jonathan J. Millin, and Krzysztof J. Geras. Isoelastic agents and wealth updates in machine learning markets. In *Proceedings 29th International Conference on Machine Learning*, 2012.

James Surowiecki. *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*. Doubleday, 2004. ISBN 978-0385503860.

Hal R. Varian. *Intermediate Microeconomics: A Modern Approach*. W.W. Norton & Company, 8th edition, 2009.

Justin Wolfers and Eric Zitzewitz. Interpreting Prediction Market Prices as Probabilities. NBER Working Paper Series, April 2006.