Condorcet Markets

Stéphane Airiau, Nicholas Kees Dupuis, Davide Grossi

Abstract

The paper studies information markets about single events from an epistemic social choice perspective. Within the classical Condorcet error model for collective binary decisions, we establish equivalence results between elections and markets, showing that the alternative that would be selected by weighed majority voting (under specific weighting schemes) corresponds to the alternative with highest price in the equilibrium of the market (under specific assumptions on the market type). This makes it possible to implement specific weighted majority elections, which are known to have superior truth-tracking performance, through information markets and, crucially, without needing to elicit voters' competences.

1 Introduction

Information markets (also known as prediction markets) [1, 13, 5, 14] are markets of all-or-nothing contracts (so-called Arrow securities) that pay one unit of currency if a designated event occurs and nothing otherwise. Under the view, inspired by [12], that markets are good aggregators of the information dispersed among traders, proponents of information markets have argued that equilibrium prices are accurate estimates of the probability of the designated event. Much research—theoretical and empirical—has probed this interpretation of prices in information markets, finding that equilibrium prices successfully track the traders' average belief about the event, under several models of trader's utilities [16, 18].

In this paper we address a closely related, but different question: *if we are to take a decision based on the information we extract from the equilibrium price, how accurate would such a decision be?* In other words, rather than relating equilibrium prices to belief aggregation, we relate them directly to the quality of the decision they would support. We frame the above question within the standard binary choice framework of epistemic social choice, stemming from the Condorcet jury theorem tradition [6, 11, 19] and the maximum-likelihood estimation approach to voting [11, 7, 17, 10].

Contribution To the best of our knowledge, the above one is a novel perspective on information markets. In particular, counter to the common assumption on traders' beliefs being subjective, we study information markets when traders' beliefs are obtained by Bayesian update from a private independent signal with known (to the trader) accuracy, just like in the classic jury theorems setting. In other words, we study 'jurors' as if they were 'traders' who, instead of relaying their vote to a central mechanism, trade in an information market. In taking this perspective, we ask the above question by comparing the decisions that would be taken based on the equilibrium price of an information market, with the decisions that would be taken by specific weighted majority elections, whose truth-tracking behavior is already well-understood [11]. Specifically, we aim at identifying correspondences between classes of markets and of weighted majority elections which are equivalent from a decision-making

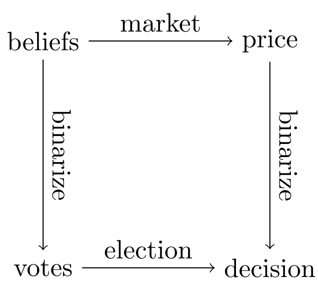

Figure 1: Elections and information markets commute.

point of view. That is, they are such that the weighted majority when agents vote according to the event they believe more likely, always coincides with the event whose price is highest in equilibrium when agents trade in Arrow securities according to their beliefs. Figure 1 depicts such relationship via a commutative diagram. Such results open up the possibility of implementing weighted majority voting with proven truth-tracking performance without needing to know jurors' competences, which may be hard to truthfully elicit or estimate [3].

Paper outline Section 2 introduces the standard binary truth-tracking framework and presents our model of information markets. Section 3 presents results on equilibrium prices in two of the three types of markets we consider (Naive and Kelly markets) and Section 4 proves 'Figure 1-type' results for those markets. Section 5 then shows how such results could be lifted even to the case of majority voting where jurors are weighted perfectly according to their competence. Section 6 outlines future research directions. Two examples illustrating our framework and analysis are provided in Appendix A. All proofs not directly included in the main text are available in Appendix B and C.

2 Preliminaries

2.1 Collective truth-tracking

We are concerned with a finite set of agents $N = \{1, \ldots, n\}$ who have to decide collectively on the correct state of the world $x \in \{A, B\}$. A prior probability $P(x = A) = \pi = 0.5$ is given, that the correct state is A. Each agent i observes a private independent signal $y_i \in \{A, B\}$ that has quality $q_i \in (0.5, 1)$, that is $q_i = P(y = A \mid x = A) = P(y = B \mid x = B)$. Each q_i represents the competence or accuracy of i. We call each vector $\mathbf{q} = (q_1, \ldots, q_n)$ of individual accuracies an *accuracy* or *competence profile* of the group. Having observed her private signal, each agent then forms a posterior belief $b_i = P(x = A \mid y = A)$ about state $x = A$ by Bayes rule. Observe that, by the conditions we impose on competences and prior, if $b_i > 0.5$ then, by Bayes rule, $b_i = q_i$ (the belief in A equals q_i), and $b_i = 1 - q_i$ (the belief in A equals $1 - q_i$) otherwise. This gives us, for all $i \in N$:

$$b_i = \mathbb{1}(b_i > 0.5) \cdot (2q_i - 1) + (1 - q_i) \tag{1}$$

where $\mathbb{1}$ denotes the indicator function. Individual beliefs are then collected in a *belief profile* $\mathbf{b} = (b_1, \ldots, b_n) \in [0, 1]^n$. Given an accuracy profile $\mathbf{q}$, the set of possible belief profiles is denoted $\mathcal{B}_\mathbf{q} = \{\mathbf{b} \in [0, 1]^n \mid P(\mathbf{b} \mid \mathbf{q}) > 0\}$. Observe that the size of this set equals 2^n: the number of all signal realizations.

Based on a profile $\mathbf{b}$ of individual beliefs, the group then takes a decision by mapping the profile to A or to B. In this process of aggregation, agents may have different weights. These weights are collected in a *weight profile* $\mathbf{w} = (w_1, \ldots, w_n) \in [0, 1]^n$. We refer to $\mathbf{1} = (1, \ldots, 1)$ as the *egalitarian weight profile* in which all agents have equal weight. So, assuming a given weight profile $\mathbf{w}$, we call *aggregator* any function

$$\mathcal{A}^\mathbf{w} : [0, 1]^n \to 2^{\{1,0\}} \backslash \emptyset \tag{2}$$

mapping belief profiles to alternatives, where $\{1\}$ denotes A; $\{0\}$ denotes B; and $\{1, 0\}$ denotes a tie.

Types of aggregators We will study two classes of mechanisms to implement aggregators. In the first class, agents cast binary ballots based on their beliefs and these ballots are submitted to a voting mechanism. The winning alternative is the outcome of the aggregation process. In the second class, agents' trade in special types of securities, based on their beliefs. The equilibrium price of this securities market is then used as a proxy for the group's belief in the probability of state A. In this case, it is the alternative favored by this collective belief to be the outcome of the aggregation process.

Let us make the above notions more precise. First of all, a belief $b \in [0, 1]$ is translated into binary opinions, or *votes*, for A or B via the following binarization function:

$$\widehat{b} = \begin{cases} \{1\} & \text{if } b > 0.5 \\ \{0\} & \text{if } b < 0.5 \\ \{0, 1\} & \text{otherwise} \end{cases} \tag{3}$$

That is, agents are assumed to vote in accordance to their posterior belief (this is sometimes referred to as sincere voting [2]). A binarized belief profile $\widehat{\mathbf{b}} = (\widehat{b}_1, \ldots, \widehat{b}_n)$ is therefore a binary vector and we will referred to such vectors also as *voting profiles* and denote them by $\mathbf{v} = (v_1, \ldots, v_n)$.[1]

Given a weight profile $\mathbf{w}$, a (belief) merger is a function $F^{\mathbf{w}} : [0, 1]^n \to [0, 1]$ taking as input a belief profile and outputting a group belief. A choice function is a function $f^{\mathbf{w}} : \{1, 0\}^n \to 2^{\{0,1\}} \backslash \emptyset$ taking as input a voting profile and outputting a possibly tied choice between 1, i.e., A, and 0, i.e., B.

We will study aggregators of the type $f^{\mathbf{w}} \circ \widehat{\ }$ (voting) and $\widehat{\ } \circ F^{\mathbf{w}}$ (trading). A voting mechanism is a choice function f^{w} which, applied to a binarized belief profile $\widehat{\mathbf{b}}$, yields a collective choice $f^{\mathbf{w}}(\widehat{\mathbf{b}})$ (under the weight profile $\mathbf{w}$). A market mechanism is a belief aggregation function F that, once applied to a belief profile $\mathbf{b}$, yields a collective belief $F^{\mathbf{w}}(\mathbf{b})$ whose binarization $\widehat{F^{\mathbf{w}}(\mathbf{b})}$ yields a collective choice (under the weight profile $\mathbf{w}$).

Group accuracy We will study aggregators from a truth-tracking perspective. The accuracy $Q(\mathcal{A}^{\mathbf{w}}, \mathbf{q})$ of an aggregator $\mathcal{A}^{\mathbf{w}}$ under the accuracy profile $\mathbf{q}$, is the conditional probability that the outcome of the aggregator is x if the state of the world is x. The above describes an epistemic social choice setting where the group is confronted with a maximum-likelihood estimation task in a dichotomous choice situation (see [10]).

2.2 Voting and market mechanisms

We turn now to the description of the mechanisms we are concerned with.

2.2.1 Voting mechanisms

After observing their private signal, agents decide whether to vote for A or B according to Equation (3). A weighted majority rule is then applied to these votes to determine the group's choice:

$$M^{\mathbf{w}}(\mathbf{v}) = \begin{cases} \{1\} & \text{if } \sum_{i \in N} w_i v_i > \frac{1}{2} \sum_{i \in N} w_i \\ \{0\} & \text{if } \sum_{i \in N} w_i v_i < \frac{1}{2} \sum_{i \in N} w_i \\ \{0, 1\} & \text{otherwise} \end{cases} \tag{4}$$

We will be working in particular with three variants of Equation (4) defined by three different weight profiles: the egalitarian weight profile $\mathbf{1}$; the weight profile allocating to each agent

[1] As individual beliefs cannot equal 0.5, the reduction function always outputs a singleton $\{0\}$ or $\{1\}$ on individual beliefs. This will not be the case, however, for collective beliefs, which may be undecided ($\{0, 1\}$).

i a weight proportional to $q_i - 0.5$; the weight profile allocating to each agent i a weight proportional to $\log \frac{q_i}{1-q_i}$. The rule induced by the egalitarian weight profile is the *simple majority* rule. We will see that the second weight profile simulates decision-making according to the average belief. The latter weight profile can be inferred from Bayes theorem and induces the weighted majority rule which we refer to as *perfect majority*, and which has been proven to optimize the truth-tracking ability of the group:

Theorem 1 ([11]). *For any accuracy profile* $\mathbf{q}$, $Q(M^{\mathbf{w}}, \mathbf{q})$ *is maximal if* $\mathbf{w}$ *is such that* $w_i \propto \ln\left(\frac{q_i}{1-q_i}\right)$ *for all* $i \in N$.

2.2.2 Markets

The market model we use is borrowed from [14, 5]. Two symmetric Arrow securities are traded: securities of type A, which cost $p_A \in [0, 1]$ and pay 1 unit of currency if $x = A$, and 0 otherwise; securities of type B, which cost $p_B \in [0, 1]$ and pay 1 unit if $x = B$ and 0 otherwise. After observing their private signal, agents decide what fraction of their endowment to invest in which securities. We assume that all agents have the *same endowment* consisting of 1 unit of currency. We also assume that agents *invest in at most one* of these securities, so if $s^A > 0$ then $s^B = 0$ and vice versa. We call agents investing in A, *A-traders* and agents investing in B, *B-traders*. In our setting, this assumption is without loss of generality (see Proposition 1 in the appendix[2]). When the true state of the world is revealed, the market resolves and payouts based on the agents' investments are distributed. We refer to tuples $\mathbf{s}^A = (s^A_1, \ldots, s^A_n)$ (respectively, $\mathbf{s}^B = (s^B_1, \ldots, s^B_n)$) as investment profiles in A-securities (respectively, B-securities). We refer to a pair $\mathbf{s} = (\mathbf{s}^A, \mathbf{s}^B)$ as an *investment profile*. We proceed now to define the notions of price, utility and equilibrium.

Market mechanism We assume that when the market opens all purchasing orders for each security are executed by the the market operator, who therefore sells all requested securities to agents when the market opens and pays the winning securities out immediately when the market resolves, that is, when it is determined whether A or B is the case. We further assume that the operator makes no profits and incurs no losses. So, for every A-security sold at price p^A a B-security is sold at price $p^B = 1 - p^A$ and vice versa. In other words, the price of the risk-less asset consisting of one of each security is $p_A + p_B = 1$. In this way the operator finances the payout of any bet by the pay-in of the opposite bet.

Under the above assumptions, the market clears[3] when the total amount of individual wealth invested in A-securities, divided by the price of A-securities (demand of A-securities) matches the amount of individual wealth invested in B-securities, divided by the price of B-securities (demand of B-securities), that is:[4]

$$\frac{1}{p^A} \sum_{i \in N} s^A_i = \frac{1}{1 - p^A} \sum_{i \in N} s^B_i. \tag{5}$$

It follows that, given an investment profile $\mathbf{s}$, solving Equation (5) for p^A, yields the clearing price $\frac{\sum_{i \in N} s^A_i}{\sum_{i \in N} s^A_i + \sum_{i \in N} s^B_i}$ to which we refer as $p(\mathbf{s})$. Note that the price is undefined if either $p_A = 0$ or $p_A = 1$. We come back to this issue in Remark 2.

[2] We are indebted to Marcus Pivato for bringing this issue to our attention.

[3] A market is normally said to clear when supply and demand match. In our model, supply and demand are implicit in the following way: purchasing one A-security at price p^A (i.e., reducing one's endowment to $1 - p^A$) is equivalent to selling one B-security thereby $p^B = 1 - p^A$.

[4] It may be worth observing that by the above design we are effectively treating the operator as an extra trader in the market, who holds a risk-less asset consisting of $\frac{1}{p^A} \sum_{i \in N} s^A_i$ A-securities and $\frac{1}{1-p^A} \sum_{i \in N} s^B_i$ B-securities. We are indebted to Marcus Pivato for this observation.

When the market resolves, each agent receives a different payout depending on how much of each security she owns, how the market resolves, and how much of her endowment is not invested. The payout, that is, the amount of wealth obtained by an agent with a given strategy s_i^A investing in A under a price p^A, is defined as follows:

$$z(p^A, s_i^A) = \begin{cases} \frac{s_i^A}{p^A} & A \text{ is correct} \\ 1 - s_i^A & \text{otherwise} \end{cases} \tag{6}$$

The payout for an investment in B-securities is defined in the same manner.

Remark 1. *In what follows, to simplify notation, we will refer to the price of A-securities as p instead of p^A and to the price of B-securities as $1 - p$ instead of p^B.*

Utility We study price p by making assumptions on how much utility agents extract from their payout at that price. We consider two types of utility functions:

Naive Given a price $p \in [0, 1]$, the naive utility function of a A-trader i is $u(p, s_i^A) = z(p, s_i^A)$ Similarly, for a B-trader, it is $u(1 - p, s_i^B) = z(1 - p, s_i^B)$. The expected utility for investment in A-securities is then:

$$U_i^A(p, s_i) = \mathbb{E}[u(p, s_i^A)] = b_i \left(\frac{s_i^A}{p} - s_i + 1 \right) + (1 - b_i)(1 - s_i^A). \tag{7}$$

The expected utility for investment in B-securities is, correspondingly, $b_i(1 - s_i^B) + (1 - b_i) \left(\frac{s_i^B}{1-p} - s_i^B \right)$. We will refer to markets under a naive utility assumption as *Naive markets.*

Kelly Given a price $p \in [0, 1]$ the Kelly [13] utility function of an A-trader i is $u(p, s_i^A) = \ln(z(p, s_i^A))$, and mutatis mutandis for B-traders. The expected Kelly utility for an A-trader is therefore:

$$U_i^A(p, s_i^A) = \mathbb{E}[u(p, s_i^A)] = b_i \ln \left(s_i^A \frac{1 - p}{p} + 1 \right) + (1 - b_i) \ln(1 - s_i^A). \tag{8}$$

Correspondingly, the expected utility of investment s_i^B for a B-traders is $b_i \ln(1 - s_i^B) + (1 - b_i) \ln \left(\frac{s_i^B}{1-p} - s_i^B \right)$. We will refer to markets under such logarithmic utility assumption as *Kelly markets.* Investing with a logarithmic utility function is known as Kelly betting and is known to maximize bettor's wealth over time [13]. Information market traders with Kelly utilities have been studied, for instance, in [5].

Equilibria For each of the above models of utility we will work with the notion of equilibrium known as competitive equilibrium [16]. This equilibrium assumes that agents optimize the choice of their investment strategy s_i under the balancing assumption of Equation (5), while not considering the effect of their choice on the price (they behave as 'price takers').

Definition 1 (Competitive equilibrium)**.** *Given a belief profile* $\mathbf{b}$*, an investment profile* $\mathbf{s}$ *is in competitive equlibrium with respect to price p if and only if:*

1. *Equation* (5) *holds, that is, $p = p(\mathbf{s})$;*
2. *for all $i \in N$, if i is a t-trader in* $\mathbf{s}$ *then $s_i^t \in \arg\max_{x \in [0,1]} U_i^t(p^t, x)$, for $t \in \{A, B\}$.*

So, when the investment profile $\mathbf{s}$ is in equilibrium with respect to price p, no agent would like to purchase more securities of any type when the price of A-securities is p. If $\mathbf{s}$ is in equilibrium with respect to $p(\mathbf{s})$ we say that $\mathbf{s}$ is an equilibrium. If equilibria always exist and are all such with respect to one unique price, then such price can be interpreted as the market's belief that the state of the world is A, given the agents' underlying beliefs $\mathbf{b}$. We can therefore view a market as a belief merger $F^{\mathbf{w}} : [0,1]^n \to [0,1]$ mapping belief profiles to the equilibrium price.

Remark 2 (Null price). *Under equation* (5) *a price* $p = 0$ *(respectively,* $p = 1$*) implies that there are no A-traders (respectively, no B-traders). In such cases Equations* (6), (7) *(Naive utility) and* (8) *(Kelly utility) would be formally undefined. Such situations, however, cannot occur in equilibrium because as p approaches* 0 *(respectively,* 1*), the utility for* $s_i^A > 0$ *(respectively,* $s_i^B > 0$*) approaches* ∞ *under both utility models. No investment profile can therefore be in equilibrium with respect to prices* $p = 0$ *or* $p = 1$*.*

3 Equilibrium price in Naive and Kelly markets

In order to see markets as belief aggregators we need to show the above market types always admit equilibria and, ideally, that equilibrium prices are unique, thereby making the aggregator resolute. The present section is concerned with these issues.

3.1 Equilibrium p in Naive markets is the $(1-p)$-quantile belief

Let us start by observing that under naive utility agents maximize their utility by investing all their wealth, unless their belief equals the price, in which case any level of investment would yield the same utility to them.

Lemma 1. *For any* $\mathbf{q}$ *and* $\mathbf{b} \in \mathcal{B}_{\mathbf{q}}$*, and* $p \in [0,1]$ *we have that, for any* $i \in N$*:*

$$\arg\max_{x\in[0,1]} U_i^A(p,x) = \begin{cases} \{1\} & \textit{if } p < b_i \\ \{0\} & \textit{if } p > b_i \\ [0,1] & \textit{otherwise} \end{cases} \; ; \quad \arg\max_{x\in[0,1]} U_i^B(p,x) = \begin{cases} \{1\} & \textit{if } (1-p) < (1-b_i) \\ \{0\} & \textit{if } (1-p) > (1-b_i) \\ [0,1] & \textit{otherwise} \end{cases}$$

Proof. We reason for A. The argument for B is symmetric. Observe first of all that Equation (7) can be rewritten as $U_i^A(p, s_i^A) = \frac{b_i}{p}(s_i^A(1-p)+p) + (1-b_i)(1-s_i^A)$. So, the utility for strategy $s_i^A = 1$ is $\frac{b_i}{p}$ and for $s_i^A = 0$ is 1. If $\frac{b_i}{p} > 1$, $U_i^A(p, s_i) \in [1, \frac{b_i}{p}]$ and so $s_1^A = 1$ maximizes Equation (7). By our assumptions, we therefore also have $s_i^B = 0$. If $\frac{b_i}{p} < 1$, instead, $U_i^A(p, s_i) \in [\frac{b_i}{p}, 1]$ and $s_1^A = 0$ maximizes Equation (7). The agent then takes the opposite side of the bet and maximizes $U_i^B(p, s_i^B)$ by setting $s_i^B = 1$. Finally, if $\frac{b_i}{p} = 1$, all investment strategies yield utility of 1. □

So, if $\mathbf{s}$ is in competitive equilbrium with respect to price $p(\mathbf{s})$ in a Naive market, then for each agent i: $s_i^A = 1$ if $b_i > p(\mathbf{s})$, $s_i^A = 0$ if $b_i < p(\mathbf{s})$, and $s_i \in [0,1]$ if $b_i = p(\mathbf{s})$, and correspondingly for s_i^B.

Let us denote by $NC(\mathbf{b})$ the set of investment profiles $\mathbf{s}$ in competitive equilibrium (under naive utilities). We show now that such equilibria always exist and are unique.

Lemma 2. *For any* $\mathbf{q}$ *and* $\mathbf{b} \in \mathcal{B}_{\mathbf{q}}$*,* $|NC(\mathbf{b})| \geq 1$*.*

Proof. We prove the claim by construction via Algorithm 1, showing that the algorithm outputs an investment profile which is in competitive equilibrium. The algorithm consists

Algorithm 1: Competitive equilibria in Naive markets

input : A belief profile $\mathbf{b} = (b_1, \ldots, b_n)$ ordered from highest to lowest beliefs
output: An investment profile $\mathbf{s} = (\mathbf{s}^A, \mathbf{s}^B)$

1 $\mathbf{s}^A \leftarrow (0, \ldots, 0)$; `/* We start by assuming no agent invests in` A `*/`
2 **for** $1 \le i < n$ **do**
3   **if** $b_i \ge \frac{i}{n} \ge b_{i+1}$ **then**
4     $\mathbf{s}^A \leftarrow (\underbrace{1, \ldots, 1}_{i \text{ times}}, 0, \ldots, 0)$ and $\mathbf{s}^B \leftarrow (\underbrace{0, \ldots, 0}_{i \text{ times}}, 1, \ldots, 1)$;
5   **return** $(\mathbf{s}^A, \mathbf{s}^B)$ and exit ; `/* profile with price` $\frac{i}{n}$ `*/`
6   **end**
7 **end**
8 **for** $1 \le i \le n$ **do**
9   **if** $\frac{i-1}{n} < b_i < \frac{i}{n}$ **then**
10     $x \leftarrow$ solve $\frac{1}{b_i}((i-1)+x) = \frac{1}{1-b_i}(n-i)$; `/* partial` A `investment */`
11   **if** $x \ge 0$ **then**
12       $s_i^A \leftarrow x$;
13       $\mathbf{s}^A \leftarrow (\underbrace{1, \ldots, 1}_{i-1 \text{ times}}, s_i^A, 0, \ldots, 0)$ and $\mathbf{s}^B \leftarrow (\underbrace{0, \ldots, 0}_{i-1 \text{ times}}, 0, 1, \ldots, 1)$;
14   **return** $(\mathbf{s}^A, \mathbf{s}^B)$ and exit
15   **else**
16       $x \leftarrow$ solve $\frac{1}{b_i}(i-1) = \frac{1}{1-b_i}((n-i)+x)$; `/* partial` B `investment */`
17       $s_i^B \leftarrow x$;
18       $\mathbf{s}^B \leftarrow (\underbrace{0, \ldots, 0}_{i-1 \text{ times}}, s_i^B, 1, \ldots, 1)$ and $\mathbf{s}^A \leftarrow (\underbrace{1, \ldots, 1}_{i-1 \text{ times}}, 0, 0, \ldots, 0)$;
19   **return** $(\mathbf{s}^A, \mathbf{s}^B)$ and exit ; `/* profile with price` b_i `*/`
20   **end**
21   **end**
22 **end**

of two routines: lines 1-7, lines 8-21. We first show that the conditions of the loops of the two routines are such that an output is always obtained. The two routines compare entries in two vectors: the n-long vector of beliefs $(b_1, \ldots, b_n)$, assumed to be ordered by decreasing values (thus, stronger beliefs first); the $n+1$-long vector $(0, \frac{1}{n}, \frac{2}{n}, \ldots, \frac{n}{n})$, ordered therefore by increasing values. The two vectors define two functions from $\{0, \ldots, n\}$ to $[0, 1]$ (we postulate $b_0 = 1$). Because the first function is non-increasing, and the second one is increasing and its image contains both 0 and 1, there exists $i \in \{0, \ldots n\}$ such that the two segments $[b_{i+1}, b_i]$ and $[\frac{i}{n}, \frac{i+1}{n}]$ have a non-empty intersection. There are four ways in which the two segments can overlap giving rise to two exhaustive cases: $\frac{i}{n}$ lies in $[b_{i+1}, b_i]$, in which case the condition of the first routine applies; or b_{i+1} lies in $[\frac{i}{n}, \frac{i+1}{n}]$, in which case the condition of the second loop applies.

It remains to be shown that the outputs of the two routines are equilibria. The output of the first routine is an investment profile $\mathbf{s} = (\mathbf{s}^A, \mathbf{s}^B)$ where i agents fully invest in A and $n - i$ agents fully invest in B, yielding a price $p(\mathbf{s}) = \frac{i}{n} \in [b_i, b_{i+1}]$. By Lemma 1 such a profile is an equilibrium. The output of the second routine is an investment profile $\mathbf{s}$ where $i - 1$ agents fully invest in A, $n - i$ agents fully invest in B and agent i, whose belief equals the price, invests partially in either A or B in order for the market to clear (Equation (5)). Again by Lemma 1 we conclude that the profile is in equilibrium with respect to b_i. □

Observe that the price constructed by Algorithm 1 is either equal to the belief of the i-th agent (ordered from stronger to weaker beliefs) or falls in the interval between the belief of the i-th and the $i+1$-th agents.

Lemma 3. *For any* $\mathbf{q}$ *and* $\mathbf{b} \in \mathcal{B}_{\mathbf{q}}$*,* $|NC(\mathbf{b})| \leq 1$*.*

Proof. Assume towards a contradiction there exist $\mathbf{s} \neq \mathbf{t} \in NC(\mathbf{b})$. It follows that $p(\mathbf{s}) \neq p(\mathbf{t})$. Assume w.l.o.g. that $p(\mathbf{s}) < p(\mathbf{t})$. By Equation 5 and the definition of competitive equilibrium, it follows that $\sum_{i \in N} s_i^A \leq \sum_{i \in N} t_i^A$ (larger A-investment in $\mathbf{t}$). By Lemma 1 it follows that there are more agents i such that $b_i > p(\mathbf{t})$ rather than $b_i > p(\mathbf{s})$, and therefore that $p(\mathbf{t}) < p(\mathbf{s})$. Contradiction. □

We can thus conclude that the equilibrium price in Naive markets is unique.

Theorem 2. *For any* $\mathbf{q}$ *and* $\mathbf{b} \in \mathcal{B}_{\mathbf{q}}$*,* $NC(\mathbf{b})$ *is a singleton.*

We will refer to such equilibrium profile as $\mathbf{s}_{NC}(\mathbf{b})$ and to its equilibrium price as $p_{NC}(\mathbf{b})$. An interesting consequence of the above results is that such equilibrium price behaves like a quantile of $\mathbf{b}$, splitting the belief profile into segments roughly proportional to the price.

Corollary 1. *Fix* $\mathbf{q}$*. For any belief profile* $\mathbf{b} \in \mathcal{B}_{\mathbf{q}}$ *there are* $n \cdot p(\mathbf{s})$ *agents* i *such that* $b_i \geq p_{NC}(\mathbf{b})$ *and there are* $n \cdot (1 - p(\mathbf{s}))$ *agents* i *such that* $b_i \leq p_{NC}(\mathbf{b})$*.*

The equilibrium price $p_{NC}(\mathbf{b})$ corresponds to the $(1 - p_{NC}(\mathbf{b}))$-quantile of $\mathbf{b}$.[5]

3.2 The average belief is the equilibrium price in Kelly markets

The two following lemmas are known results, which we restate here for completeness.

Lemma 4 ([13])**.** *For any* $b_i \in [0, 1]$ *and* $p \in [0, 1]$*:*

$$\arg\max_{x \in [0,1]} U_i^A(p, x) = \begin{cases} \frac{b_i - p}{1 - p} & \textit{if } p < b_i \\ 0 & \textit{otherwise} \end{cases} ; \quad \arg\max_{x \in [0,1]} U_i^B(p, x) = \begin{cases} \frac{p - b_i}{p} & \textit{if } (1 - p) < (1 - b_i) \\ 0 & \textit{otherwise} \end{cases}$$

So, a strategy profile $\mathbf{s}$ is in Kelly competitive equilibrium with respect to price $p(\mathbf{s})$ whenever Equation (5) is satisfied together with the 'Kelly conditions' of Lemma 4. Unlike in the case of Naive markets it is easy to see that such equilibrium is unique. So, for a given belief profile $\mathbf{b}$ let us denote by $\mathbf{s}_{KC}(\mathbf{b})$ such competitive equilibrium and by $p_{KC}(\mathbf{b})$ the price at such equilibrium. We then have:

Lemma 5 ([5])**.** *For any* $\mathbf{q}$ *and* $\mathbf{b} \in \mathcal{B}_{\mathbf{q}}$*,* $p_{KC}(\mathbf{b}) = \frac{1}{|N|} \sum_{i \in N} b_i$*.*

4 Truth-Tracking via Equilibrium Prices

In this section we show how competitive equilibria in Naive and Kelly markets correspond to election by simple majority and, respectively, by a majority in which agents carry weight proportional to their competence minus 0.5.

[5] A similar observation, but for a continuum of players ($N = [0, 1]$) and for subjective beliefs is reported in [15].

4.1 Simple majority and Naive markets

Simple majority is implemented in competitive equilibrium by a Naive market:

Theorem 3. *For any* $\mathbf{q}$ *and* $\mathbf{b} \in \mathcal{B}_{\mathbf{q}}$*:* $M^{\mathbf{1}}(\widehat{\mathbf{b}}) = \widehat{p_{NC}(\mathbf{b})}$.

Proof. The claim follows from the observation that, by Corollary 1, $p_{NC}(\mathbf{b}) > 0.5$ if and only if there exists a majority of traders whose beliefs are higher than the price. From which we conclude that $\hat{\mathbf{b}}$ determines a voting profile with a majority of votes for A. □

Put otherwise, the theorem tells us that the outcome of simple majority always consists of the security that the $(1-p)$-quantile belief (where p is the equilibrium price) would invest in equilibrium when the market is naive. So we can treat NC as a belief aggregator $[0,1]^n \to [0,1]$ mapping belief profiles to prices induced by competitive equilibria. In other words, for any belief profile $\mathbf{b}$ induced by independent individual competences in $(0.5, 1]$, the diagram on the right commutes.

$$\begin{array}{ccc} \mathbf{b} & \xrightarrow{NC} & p_{NC}(\mathbf{b}) \\ \downarrow\widehat{} & & \downarrow\widehat{} \\ \widehat{\mathbf{b}} & \xrightarrow{M^{\mathbf{1}}} & M^{\mathbf{1}}(\widehat{\mathbf{b}}) \end{array}$$

Remark 3. *It is worth observing that, by Theorem 3, known extensions of the Condorcet Jury Theorem with heterogeneous competences [11] directly apply to Naive markets in competitive equilibrium. In particular with* $N \to \infty$ *the probability that* $p_{NC}(\mathbf{b})$ *is correct approaches* 1 *for any* $\mathbf{b}$ *induced by a given competence profile.*

4.2 Weighted majority and Kelly markets

A weighted majority rule with weights for each i proportional to $q_i - 0.5$ is implemented in competitive equilibrium by Kelly markets. Intuitively, such markets then implement a majority election where individuals' weights are proportional to how much better the individual is compared to an unbiased coin.

Theorem 4. *For any* $\mathbf{q}$*, and* $\mathbf{b} \in \mathcal{B}_{\mathbf{q}}$*:* $M^{\mathbf{w}}(\widehat{\mathbf{b}}) = \widehat{p_{KC}(\mathbf{b})}$*, where* $\mathbf{w}$ *is s.t. for all* $i \in N$*,* $w_i \propto 2q_i - 1$.

Intuitively, by implementing a weighted average of the beliefs, the competitive equilibrium price for Kelly utilities behaves like a weighted majority where agents' weights are a linear function of their individual competence $(2q_i - 1)$. So, for any belief profile $\mathbf{b}$ induced by a competence profile $\mathbf{q}$ and weights $w_i = 2q_i - 1$, the diagram on the right commutes.

$$\begin{array}{ccc} \mathbf{b} & \xrightarrow{KC} & p_{KC}(\mathbf{b}) \\ \downarrow\widehat{} & & \downarrow\widehat{} \\ \widehat{\mathbf{b}} & \xrightarrow{M^{\mathbf{w}}} & M^{\mathbf{w}}(\widehat{\mathbf{b}}) \end{array}$$

5 Markets for Perfect Elections

In this section we show how, by introducing a specific tax scheme, we can modify Kelly markets in such a way as to make their equilibrium price implement a perfect weighted majority, that is, a majority in which the weight of each individual is proportional to the natural logarithm of their competence ratio. The intuition of our approach is the following: Theorem 4 has shown that Kelly markets correspond to elections where individuals are weighted proportionally to their competence in excess of 0.5; in order to bring such weights closer to the ideal values of Theorem 1 we need therefore to allow more competent agents to exert substantially more influence on the equilibrium price; we do so by designing a tax scheme which achieves such effect asymptotically in one parameter of the scheme.

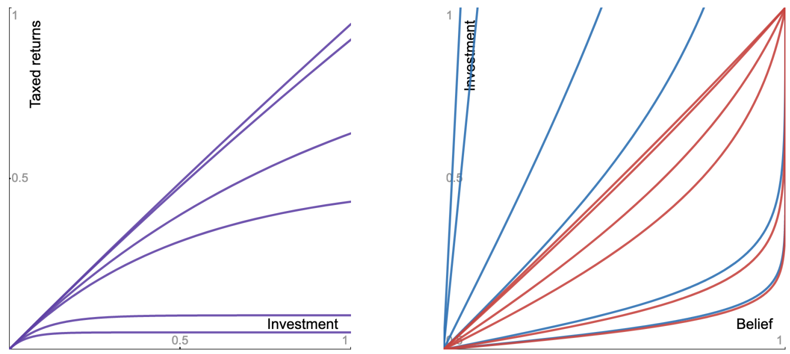

Figure 2: Left: returns after taxation by T as a function of investment (Equation (9)). Right: investment strategy (red) approximating $\ln\left(\frac{b_i}{1-b_i}\right)\frac{1}{k}$ (blue) as k grows when price equals 0.5. Functions plotted for $k \in \{0.1, 0.2, 1, 2, 10, 20\}$.

5.1 Taxing payouts

We modify Equation (8) by building in the effects of a tax scheme T on utility as follows:

$$U_i^A(p, s_i) = b_i \ln T\left(s_i \frac{1-p}{p} + 1\right) + (1 - b_i)\ln(1 - s_i) \tag{9}$$

where

$$T(x) = \frac{1 - e^{-kx\frac{p}{1-p}}}{k\frac{p}{1-p}} \tag{10}$$

with $k \in \mathbb{R}^{>0}$. Observe that as parameter k approaches 0, $T(x)$ approaches x, that is, null taxation is approached.

The best way to gain an intuition of the working of function T is to observe its effects on the agent's optimal investment strategy supposing the price 0.5. For $p = 0.5$ the optimal strategy of a Kelly trader is $2b_i - 1$ (Lemma 4). Function T makes that strategy asymptotically proportional to $\ln\left(\frac{b_i}{1-b_i}\right)$ (Figure 2) as k grows.

We call markets under the utility in Equation (9) *taxed markets* and denote their equilibrium prize by $p_{TC}(\mathbf{b})$, for any belief profile $\mathbf{b}$.

5.2 Equilibria in taxed Kelly markets

Like for Naive and Kelly markets, we first determine the optimal strategy of the traders. We do that for A-traders, as the lemma for B-traders is symmetric.

Lemma 6. *For any $i \in N$, if $b_i > p$, then as $k \to \infty$,*

$$\arg\max_{x\in[0,1]} U_i^A(p, x) \propto \ln\left(\frac{1-p}{p} \cdot \frac{b_i}{1-b_i}\right).$$

Proof. We start from i's utility, given by Equation (9). By setting $\frac{dU_i^A}{ds_i} = 0$ (first order condition) we obtain:

$$\frac{bT'(s\frac{1-p}{p})\frac{1-p}{p}}{1 + T(s\frac{1-p}{p})} = \frac{1-b_i}{1-s_i} \tag{11}$$

If we plug Equation (10) into Equation (11), we obtain:

$$\frac{be^{-ks_i}\frac{1-p}{p}}{1+\frac{1-e^{-ks_i}}{k\frac{p}{1-p}}}=\frac{1-b_i}{1-s_i} \tag{12}$$

and therefore

$$\frac{kbe^{-ks_i}}{k\frac{p}{1-p}+1-e^{-ks_i}}=\frac{1-b}{1-s_i}. \tag{13}$$

As k approaches infinity, s_i approaches zero. For this reason we rescale strategies by k and consider a value $y = sk$. This allows us to understand the form to which strategies tend to as they approach zero. We thus obtain:

$$\frac{kbe^{-y}}{k\frac{p}{1-p}+1-e^{-y}}=\frac{1-b}{1-\frac{y}{k}}. \tag{14}$$

As k approaches infinity this approaches:

$$\frac{be^{-y}}{\frac{p}{1-p}}=(1-b) \tag{15}$$

which can be rewritten in turn as

$$y=\ln\left(\frac{1-p}{p}\frac{b}{1-b}\right) \tag{16}$$

from which we conclude $s_i = \frac{1}{k}\log(\frac{1-p}{p}\frac{b}{1-b})$ as desired. □

As k tends to infinity, the optimal investment strategy will tend to 0 for all agents. However, it will do so in such a way that as k grows, the optimal investment strategy tends to be proportional to $\ln(\frac{1-p}{p}\cdot\frac{b_i}{1-b_i})$ as desired.

So, as k grows large, a strategy profile $\mathbf{s}$ is in competitive equilibrium in a taxed market with respect to price $p(\mathbf{s})$ whenever Equation (5) is satisfied together with the condition identified by Lemma 6. We denote by $\mathbf{s}_{TC}(\mathbf{b})$ such competitive equilibrium and by $p_{TC}(\mathbf{b})$ the price at such equilibrium. We then have:

Lemma 7. *For any* $\mathbf{q}$ *and* $\mathbf{b} \in \mathcal{B}_{\mathbf{q}}$*, as* $k \to \infty$*,*

$$\ln\left(\frac{p_{TC}(\mathbf{b})}{1-p_{TC}(\mathbf{b})}\right) \propto \sum_i^n \ln\left(\frac{b_i}{1-b_i}\right).$$

Proof. To lighten notation we write p for $p_{TC}(\mathbf{b})$. From the equilibrium condition (Equation (5)) and Lemma 6 we have that

$$\frac{1}{p}\sum_{i\in N^A}\ln\frac{b_i}{1-b_i}=\frac{1}{1-p}\sum_{i\in N^B}\frac{1-b_i}{b_i} \tag{17}$$

where $N^A = \{i \in N \mid b_i > p\}$ and $N^B = \{i \in N \mid b_i < p\}$. From the above we obtain:

$$0=\sum_i^N \ln\left(\frac{1-p}{p}\frac{b_i}{1-b_i}\right) \tag{18}$$

which rewrites to

$$\ln\left(\frac{p}{1-p}\right)=\frac{1}{N}\sum_i^N \ln\left(\frac{b_i}{1-b_i}\right) \tag{19}$$

as desired. □

That is, the equilibrium price ratio between A and B securities in a taxed market tends to be proportional, in logarithmic scale, to the average belief ratio.

Theorem 5. *For any* $\mathbf{q}$, $\mathbf{b} \in \mathcal{B}_{\mathbf{q}}$ *and* $k \to \infty$*:*

$$M^{\mathbf{w}}(\widehat{\mathbf{b}}) = \widehat{p_{TC}(\mathbf{b})}$$

where that $\mathbf{w}$ *is s.t. for all* $i \in N$ $w_i \propto \ln \frac{q_i}{1-q_1}$.

This last result shows that elections that are perfect from a truth-tracking perspective (Theorem 1) can be implemented increasingly faithfully by Kelly markets once the taxation scheme T is applied and the taxation parameter k in Equation (10) grows larger. So, for any belief profile $\mathbf{b}$ induced by a competence profile $\mathbf{q}$ and weights $w_i = \frac{q_i}{1-q_i}$, the diagram on the right commutes as k tends to infinity and, therefore, taxation grows.

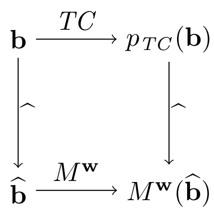

6 Conclusions and Outlook

Our paper is, to the best of our knowledge, the first one establishing a formal link between voting and information markets from an epistemic social choice perspective. The link consists specifically of correspondence results between weighted majority voting on the one hand, and information markets under three types of utility on the other. Such results open up the possibility to implement weighted majority voting with strong epistemic guarantees even without having access to individual competences because such information becomes indirectly available in the market via the equilibrium price. Notice that, in particular, while it may be difficult to elicit truthful weights from agents, investment strategies are subject to the natural incentive of maximizing investment returns. Whether this can prove advantageous also in practice, for instance in the setting of classification markets [4] or voting-based ensembles [8], should definitely be object of future research.

The study we presented is subject to at least three main limitations. First, our analysis inherits all assumptions built into standard jury theorems, in particular: jurors' independence; homogeneous priors; equivalence of type-1 and type-2 errors in jurors' competences. Future research should try to lift our correspondence to more general settings relaxing the above assumptions (see [9] for a recent overview). Secondly, our study limited itself to one-shot interactions. However, markets and specifically Kelly betting make most sense in a context of iterated decisions. Extending our results to the iterated setting, along the lines followed for instance in [5], is also a natural avenue for future research. Thirdly, our market model makes use of the notion of competitive equilibrium. Although such notion of equilibrium is standard in information markets, it responds to the intuition individuals operate in a large group and, therefore, behave as price takers. We consider it interesting to study, at least experimentally, how different notions of equilibrium that do not make such assumption (e.g., Nash equilibrium) would behave within our framework.

Acknowledgments This research was (partially) funded by the Hybrid Intelligence Center, a 10-year programme funded by the Dutch Ministry of Education, Culture and Science through the Netherlands Organisation for Scientific Research, `https://hybrid-intelligence-centre.nl`, grant number 024.004.022. Davide Grossi wishes to also thank Université Paris Dauphine and the Netherlands Institute for Advanced Studies (NIAS), where parts of this research were completed. The authors also wish to thank the anonymous reviewers of COMSOC'23 for several helpful suggestions.

References

[1] C. Alós-Ferrer and A. Ania. The asset market game. *Journal of Mathematical Economics*, 41:67–90, 2005.

[2] D. Austen-Smith and J. Banks. Information aggregation, rationality, and the Condorcet jury theorem. *American Political Science Review*, 90:34–45, 1996.

[3] Eyal Baharad, Jacob Goldberger, Moshe Koppel, and Shmuel Nitzan. Beyond Condorcet: Optimal aggregation rules using voting records. *Bar-Ilan University Department of Economics Research Paper*, (2010-20), 2010.

[4] Adrian Barbu and Nathan Lay. An introduction to artificial prediction markets for classification. *J. Mach. Learn. Res.*, 13:2177–2204, 2012.

[5] Alina Beygelzimer, John Langford, and David M. Pennock. Learning performance of prediction markets with kelly bettors. In Wiebe van der Hoek, Lin Padgham, Vincent Conitzer, and Michael Winikoff, editors, *International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2012, Valencia, Spain, June 4-8, 2012 (3 Volumes)*, pages 1317–1318. IFAAMAS, 2012.

[6] Marquis de Condorcet, M.J.A.N. de C. *Essai sur l'Application de l'Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix.* Imprimerie Royale, Paris, 1785.

[7] Vincent Conitzer and Tuomas Sandholm. Common voting rules as maximum likelihood estimators. In *UAI '05, Proceedings of the 21st Conference in Uncertainty in Artificial Intelligence, Edinburgh, Scotland, July 26-29, 2005*, pages 145–152. AUAI Press, 2005.

[8] Cristina Cornelio, Michele Donini, Andrea Loreggia, Maria Silvia Pini, and Francesca Rossi. Voting with random classifiers (VORACE): theoretical and experimental analysis. In Piotr Faliszewski, Viviana Mascardi, Catherine Pelachaud, and Matthew E. Taylor, editors, *21st International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2022, Auckland, New Zealand, May 9-13, 2022*, pages 1929–1931. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), 2022.

[9] F. Dietrich and K. Spiekermann. Jury theorems. *Stanford Encyclopedia of Philosophy*, 2021.

[10] Edith Elkind and Arkadii Slinko. Rationalizations of voting rules. In *Handbook of Computational Social Choice*, pages 169–196. 2016.

[11] B. Grofman, G. Owen, and S. Feld. Thirteen theorems in search of the truth. *Theory and Decision*, 15(3):261–278, 1983.

[12] F. Hayek. The use of knowledge in society. *American Economic Review*, 35:519–530, 1945.

[13] J. L. Kelly. A new interpretation of information rate. *The Bell System Technical Journal*, 35(4):917–926, 1956.

[14] Willemien Kets, David M. Pennock, Rajiv Sethi, and Nisarg Shah. Betting strategies, market selection, and the wisdom of crowds. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27 -31, 2014, Québec City, Québec, Canada*, pages 735–741, 2014.

[15] Charles F Manski. Interpreting the predictions of prediction markets. *economics letters*, 91(3):425–429, 2006.

[16] David Pennock. *Aggregating Probabilistic Beliefs: Market Mechanisms and Graphical Representations*. 1999.

[17] Marcus Pivato. Voting rules as statistical estimators. *Social Choice and Welfare*, 40(2):581–630, 2013.

[18] Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of economic perspectives*, 18(2):107–126, 2004.

[19] H. P. Young. Condorcet's theory of voting. *American Political Science Review*, 82(4):1231–1244, 1988.

A Examples

Assume $N = \{1, \dots, 5\}$ with competence profile $\mathbf{q} = (0.9, 0.7, 0.6, 0.6, 0.6)$. Assume further that only the first and last agent receive signal A while the rest receives signal B. This gives us the following belief profile by Bayesian update: $\mathbf{b} = (0.9, 0.3, 0.4, 0.4, 0.6)$.

The above beliefs result in the voting profile $\mathbf{v} = (1, 0, 0, 0, 1)$, from which we obtain:

- $M^{\mathbf{1}}(\mathbf{v}) = \{0\}$, that is, standard majority selects B
- $M^{\mathbf{w}}(\mathbf{v}) = \{1\}$ where $\mathbf{w} = (0.8, 0.4, 0.2, 0.2, 0.2)$ (weight profile given by $2q_i - 1$) as $0.8 + 0.2 - (0.4 + 0.2 + 0.2) > 0$, that is, the sum of weights of the first and last agents are larger then the sum of weights of the remaining agents.
- $M^{\mathbf{w}}(\mathbf{v}) = \{1\}$ where $\mathbf{w} = \left(\ln \frac{0.9}{0.1}, \ln \frac{0.7}{0.3}, \ln \frac{0.6}{0.4}, \ln \frac{0.6}{0.4}, \ln \frac{0.6}{0.4}\right)$ (optimal weights) as
$$\ln \frac{0.9}{0.1} + \ln \frac{0.6}{0.4} - \left(\ln \frac{0.7}{0.3} + 2 \cdot \ln \frac{0.6}{0.4}\right) \tag{20}$$
is positive.

We move now to the choices made by the markets based on equilibrium prices. We have that: by Algorithm 1, $p_{NC}(\mathbf{b}) = \frac{2}{5}$ (Naive market equilibrium) where the two agents who received the A signal invest all their endowment in A-securities, and the remaining agents invest all their endowment in B-securities; $p_{KC}(\mathbf{b}) = \frac{2.6}{5}$ (Kelly market equilibrium) corresponding to the mean belief in $\mathbf{b}$. So, a Naive market given the above beliefs selects B while the Kelly market selects A by a very small margin. As to the taxed markets for perfect elections, our results do not give us a closed expression for $p_{TC}(\mathbf{b})$ but rather determine whether the price favors A- or B-securities based on the logarithm of the ratio between the two prices, which is proportional to the logarithm of the weighed support for A and for B when the taxation parameter k tends to infinity (Theorem 5). In this example, we thus have that $\ln \frac{p_{TC}(\mathbf{b})}{1-p_{TC}(\mathbf{b})}$ is proportional to Equation (21) and therefore points to security A.

Assume $N = \{1, \dots, 4\}$ with competence profile $\mathbf{q} = (0.8, 0.6, 0.6, 0.6)$. Assume further that only the first agent receives signal A while the rest receives signal B. This gives us the following belief profile by Bayesian update: $\mathbf{b} = (0.8, 0.4, 0.4, 0.4)$.

The above beliefs result in the voting profile $\mathbf{v} = (1, 0, 0, 0)$, from which we obtain:

- $M^{\mathbf{1}}(\mathbf{v}) = \{0\}$, that is, standard majority selects B

- $M^{\mathbf{w}}(\mathbf{v}) = \{0, 1\}$ where $\mathbf{w} = (0.6, 0.2, 0.2, 0.2)$ (weight profile given by $2q_i - 1$) as $0.6 - (0.2 + 0.2 + 0.2) = 0$. That is, we have a split weighted majority.
- $M^{\mathbf{w}}(\mathbf{v}) = \{1\}$ where $\mathbf{w} = \left(\ln \frac{0.8}{0.2}, \ln \frac{0.6}{0.4}, \ln \frac{0.6}{0.4}, \ln \frac{0.6}{0.4}\right)$ (optimal weights) as

$$\ln \frac{0.8}{0.2} - 3 \cdot \ln \frac{0.6}{0.4} \tag{21}$$

is positive.

As to equilibrium prices, by applying Algorithm 1, we have that also in this case $p_{NC}(\mathbf{b}) = \frac{2}{5}$ (Naive market equilibrium). This price equals the posterior beliefs of the three agents that receive signal B. By the algorithm, the agent receiving signal A invests all its wealth in A, one of the agents receiving signal B invests $\frac{1}{3}$ of their wealth in A (to guarantee market clearing at that price, line 10 of Algorithm 1), and the remaining agents invest all their endowment in B-securities. The equilibrium price in Kelly markets is in this case 0.5 (mean belief). So, a Naive market given the above beliefs selects B while the Kelly market remains undecided. In taxed markets, we have that $\ln \frac{p_{TC}(\mathbf{b})}{1-p_{TC}(\mathbf{b})}$ (ratio of prices in logarithmic scale) is proportional to Equation (21) and therefore points to security A.

B Omitted proofs

Proof of Lemma 4 [13]. We reason for A. The argument for B is symmetric. First we compute the first derivative of $U_i^A(p, x)$:

$$\begin{aligned}
\frac{d}{dx} U_i^A(p, x) &= \frac{d}{dx} \left(b_i \ln(x \frac{1-p}{p} + 1) + (1 - b_i) \ln(1 - x) \right) \\
&= b_i \cdot \frac{\frac{d}{dx}(x \frac{1-p}{p} + 1)}{x \frac{1-p}{p} + 1} + (1 - b_i) \cdot \frac{\frac{d}{dx}(1 - x)}{1 - x} \\
&= b_i \cdot \frac{\frac{1-p}{p}}{x \frac{1-p}{p} + 1} + (1 - b_i) \cdot \frac{-1}{1 - x} \\
&= \frac{b_i(1-p)}{x + p(1-x)} - \frac{1 - b_i}{1 - x}
\end{aligned}$$

The first-order condition is $\frac{b_i(1-p)}{x+p(1-x)} - \frac{1-b_i}{1-x} = 0$, that is, $\frac{b_i - p}{1-b} = \frac{x}{1-x}$. Solving for x we obtain $\frac{b_i - p}{1 - b_i}$. We need then to check the second-order condition. The second derivative is $-\frac{1-b_i}{(1-x)^2} - \frac{b_i(1-p)^2}{(p(1-x)+x)^2}$, which is negative at every point and therefore, as desired, also at $\frac{b_i - p}{1 - b_i}$. Now observe that $\frac{b_i - p}{1 - b_i}$ is positive whenever $b_i > p$ and equals 0 when $b_i = p$. When $b_i < p$ then the expression is negative, and the trader takes the opposite side of the bet, investing in B. □

Proof of Lemma 5 [5]. Let us denote $p_{KC}(\mathbf{b})$ by p. By Lemma 4, Equation (5) becomes:

$$\frac{1}{p} \sum_{i \in N^A} \frac{b_i - p}{1 - p} = \frac{1}{1 - p} \sum_{i \in N^B} \frac{p - b_i}{p} \tag{22}$$

where $N^A = \{i \in N \mid b_i > p\}$ and $N^B = \{i \in N \mid b_i < p\}$. Notice we can assume $N = N^A \cup N^B$ as any agent with belief equal to the price does not invest in either A- or B-securities in Kelly markets (by Lemma 4). From Equation (22), by some basic algebra, we obtain the following series of equations

$$
\begin{aligned}
\frac{1-p}{p}\sum_{i\in N^A}\frac{b_i-p}{1-p} &= \sum_{i\in N^B}\frac{p-b_i}{p}\\
\sum_{i\in N^A} b_i - p &= \sum_{i\in N^B} p - b_i\\
np &= \sum_{i\in N} b_i\\
p &= \frac{1}{n}\sum_{i\in N} b_i
\end{aligned}
$$

thereby proving the claim. ☐

Proof of Theorem 4. By the assumed weight profile $\mathbf{v}$ and Equation (4) we can rewrite $M^{\mathbf{v}}(\widehat{\mathbf{b}})$ as $\widehat{R(\mathbf{b})}$ where:

$$
\begin{aligned}
R(\mathbf{b}) &= \frac{\sum_{i\in N}\mathbb{1}(b_i>0.5)(2q_i-1)}{\sum_{i\in N}(2q_i-1)}\\
&= \frac{\sum_{i\in N}(2q_i-1)}{n}\cdot\left(\frac{\sum_{i\in N}\mathbb{1}(b_i>0.5}{(}2q_i-1))\sum_{i\in N}(2q_i-1)-\frac{1}{2}\right)+\frac{1}{2}\\
&= \frac{\sum_{i\in N}(2q_i-1)}{n}\cdot\left(\frac{2\sum_{i\in N}\mathbb{1}(b_i>0.5)(2q_i-1)-\sum_{i\in N}(2q_i-1)}{2\sum_{i\in N}(2q_i-1)}\right)+\frac{1}{2}\\
&= \frac{\sum_{i\in N}\mathbb{1}(b_i>0.5)(2q_i-1)-\sum_{i\in N}(q_i-\frac{1}{2})}{n}+\frac{1}{2}\\
&= \frac{\sum_{i\in N}\mathbb{1}(b_i>0.5)(2q_i-1)-\sum_{i\in N}(q_i-\frac{1}{2})+\frac{n}{2}}{n}\\
&= \frac{\sum_{i\in N}\mathbb{1}(b_i>0.5)(2q_i-1)-\sum_{i\in N}q_i+\frac{n}{2}+\frac{n}{2}}{n}\\
&= \frac{\sum_{i\in N}\mathbb{1}(b_i>0.5)(2q_i-1)+\sum_{i\in N}(1-q_i)}{n}\\
&= \frac{\sum_{i\in N}\mathbb{1}(b_i>0.5)(2q_i-1)+(1-q_i)}{n} \qquad \text{(recall Equation (1))}\\
&= \frac{\sum_{i\in N} b_i}{n}.
\end{aligned}
$$

From this and Lemma 5 we obtain $M^{\mathbf{v}}(\widehat{\mathbf{b}}) = \widehat{\frac{\sum_{i\in N} b_i}{n}} = p_{KC}(\mathbf{b})$ as desired. ☐

Proof of Theorem 5. First of all, observe that: $\widehat{p_{TC}(\mathbf{b})} = \{1\}$ iff $\ln\left(\frac{p_{TC}(\mathbf{b})}{1-p_{TC}(\mathbf{b})}\right) > 0$; $\widehat{p_{TC}(\mathbf{b})} = \{0,1\}$ iff $\ln\left(\frac{p_{TC}(\mathbf{b})}{1-p_{TC}(\mathbf{b})}\right) = 0$; and $\widehat{p_{TC}(\mathbf{b})} = \{0\}$ iff $\ln\left(\frac{p_{TC}(\mathbf{b})}{1-p_{TC}(\mathbf{b})}\right) < 0$. Then, by Lemma 7, Equation (1) and some algebra we obtain the following relations:

$$\begin{aligned}
\ln\left(\frac{p_{TC}(\mathbf{b})}{1-p_{TC}(\mathbf{b})}\right) &\propto \sum_{i}^{n} \ln\left(\frac{b_i}{1-b_i}\right) \\
&= \sum_{i}^{n} \mathbb{1}(b_i > 0.5) \cdot \ln\left(\frac{q_i}{1-q_i}\right) \\
&= \sum_{i:b_i>0.5} \ln\left(\frac{q_i}{1-q_i}\right) + \sum_{i:b_i<0.5} \ln\left(\frac{1-q_i}{q_i}\right) \\
&= \sum_{i:b_i>0.5} \ln\left(\frac{q_i}{1-q_i}\right) - \sum_{i:b_i<0.5} \ln\left(\frac{q_i}{1-q_i}\right)
\end{aligned}$$

The last expression is: positive whenever weighted voting with optimal weights returns $\{1\}$; negative whenever it returns $\{0\}$; and 0 whenever it returns $\{0, 1\}$ (Equation (4)). □

C Supplementary results

Let us define the investment strategy of a trader i by the pair $s_i = (s_i^A, s_i^B)$ where $s_i^A \in [0, 1]$ is the amount of endowment invested in A-securities (similarly for s_i^B). Let then $c_i = w_i - (s_i^A + s_i^B)$, that is, the unspent endowment of i given strategy (s_i^A, s_i^B). The assumption that traders invest their full endowment amounts to $s_i^A + s_i^B = 1$. We can therefore simply refer to s_i^A as s and to s_i^B as $1 - s$.

In Naive markets, it is a corollary of Lemma 1 that, in equilibrium, any strategy in which traders invest all their endowment is equivalent to a strategy investing in at most one security.

For Kelly markets we can prove that the same observation holds. Under a full investment assumption Equation (8) needs to be modified to:

$$U_i(p, s) = b_i \ln\left(\frac{s}{p}\right) + (1 - b_i) \ln\left(\frac{1-s}{1-p}\right). \tag{23}$$

which can be further generalized to

$$U_i(p, s) = b_i \ln\left(\frac{s}{p} + c\right) + (1 - b_i) \ln\left(\frac{1-s}{1-p} + c\right). \tag{24}$$

when i can invest any amount of endowment in A- or B-securities and still keep unspent endowment c.

Proposition 1. *In Kelly markets, for all $i \in N$, $p \in [0, 1]$ and s_i such that $s_i^A + s_i^B = 1$ (whole endowment is spent) there exists $t_i = (t_i^A, 0)$ or $t_i = (0, t_i^B)$ which yield the same utility as s_i, and vice versa.*

Proof. First observe that the first-order condition for Equation (23) gives us the optimal investment strategy, which is b_i. Such optimal strategy will give i $\frac{b_i}{p}$ A-securities and $\frac{1-b_i}{1-p}$ B-securities. Left to right Suppose that $b_i > p$. It follows that $\frac{b_i}{p} > \frac{1-b_i}{1-p}$, that is, i purchases more A-securities than B-securities. Were i not be compelled to invest all her endowment, she would invest according to Equation (24). Now it is easy to see that if i invests b_i in A-securities, 0 in B-securities, and keeps $1 - b_i$ as cash, she obtains the same payoff by Equation (24) as she would obtain by Equation (23) when she invests $1 - b_i$ in B-securities. There exists therefore $t_i = (t_i^A, 0)$ yielding the same utility. The same reasoning applies, symmetrically, for $b_i < p$. Right to left. Given an investment strategy t_i investing in only one security, a utility-equivalent full-investment strategy can be constructed using the same reasoning used in the previous case via the utility for unrestricted investments of Equation (24). □

Stéphane Airiau
University Paris-Dauphine
France
Email: stephane.airiau@dauphine.fr

Nicholas Kees Dupuis
Email: nkdupuis@protonmail.com

Davide Grossi
University of Groningen and University of Amsterdam
The Netherlands
Email: d.grossi@rug.nl