Submitted for publication.

# A Market Framework for Pooling Opinions

**David M. Pennock and Michael P. Wellman**
University of Michigan AI Laboratory
1101 Beal Avenue
Ann Arbor, MI 48109-2110 USA
{dpennock,wellman}@umich.edu

July 8, 1999

**Abstract**

Several individuals disagree on the probability of rain tomorrow. For decades, statisticians have grappled with how best to combine their beliefs together to form a consensus assessment. Many aggregation functions or *opinion pools* have been proposed, each sensible under particular assumptions or measures. We examine a *market-based* approach, where the equilibrium prices of *securities*—lottery tickets that pay off $1 contingent on, for example, whether it rains—are interpreted as the consensus beliefs. We develop a simple model and identify correspondences between the consensus probabilities derived by the opinion pools and by the market, where expert weights in the former context are parameters of the agents' utilities in the latter. In an extended model, in which agents learn from prices, the weights are shown to additionally depend on self-assessed confidence. Even when agents themselves do not learn, we show that the market as a whole can adapt over time in the manner of a Bayesian-rational individual, by transferring wealth from inaccurate to accurate agents. We also propose tests whereby an outside observer might distinguish a group's aggregate behavior from that of an individual.

# 1 Introduction

The notion that a rational decision maker should behave according to a subjective probability distribution is widely accepted. Defining an analogous concept for a *group* of decision makers has met with more controversy. Many *belief aggregation* methods, or opinion pools, have been proposed to derive consensus probabilities for groups. Forming consensus beliefs is a classic problem in statistics [11, 17, 23, 36, 68, 69, 70], and has found application in many areas, from group decision making [43, 67, 76] to artificial intelligence [30, 40, 50, 53, 54]. In some group situations there is a natural leader ultimately responsible for decisions, and standard techniques may apply [37]. Other groups are more egalitarian; in these situations, understanding what beliefs can be the object of a consensus, and how collective decisions can be made, may be of considerable interest. Despite decades of research on opinion pools, including a history of contributions reported in this journal [4, 5, 6, 12, 44, 45, 46, 72, 73, 74], definitive answers have been elusive, and various impossibility results [29, 9, 19, 22, 34, 66] and definitional controversies [17] in the literature cast doubt on the prospects for an entirely satisfactory solution. Much debate revolves around which properties of the aggregate distribution are "reasonable", and which pooling functions exhibit them. We take instead a bottom-up approach, beginning with characterizations of individual behavior, and deriving the implications for aggregate properties from there.

One of the most convincing justifications for *individual* subjective probabilities is de Finetti's famous *no Dutch book* argument [10]. Imagine that a bettor $B$ must decide at what price to buy and sell each of two lottery tickets: one pays \$1 if an event $A_1$ occurs, and nothing otherwise; the other pays \$1 if $A_2$. If $A_1$ and $A_2$ are disjoint events, then $B$'s total price for both tickets should never exceed \$1—otherwise, $B$ is exposed to a guaranteed loss: a shrewd opponent $O$ will sell both tickets to $B$ for more than \$1, paying out at most \$1 in all outcomes. More generally, if $B$'s internal prices for a set of such lottery tickets are not legal probabilities, then $B$ is subject to certain losses—also called *Dutch books* or *arbitrage* [64]. Now imagine that a *group* of bettors, $B_1, B_2, \ldots$ are locked in a room. $O$ is outside, and slides offers to buy and sell lottery tickets under the door. $O$ does not care who accepts the offers: if opportunities for risk-free profit exist, $O$ will take advantage of them. From $O$'s perspective, such opportunities imply that the group or "room" is irrational. To be concrete, imagine that $B_1$ will sell "\$1 if $A_1$"

at \$0.5, and $B_2$ will buy the same ticket at \$0.7. $O$ could potentially buy a ticket from $B_1$ and sell it to $B_2$ for a riskless profit of \$0.2. But if the group had first exhausted all opportunities for within-group trades, then $B_1$ would sell instead to $B_2$. Each bettor's transactions are exactly the same, and the group as a whole is better off by \$0.2. Only if there is *excess demand* within the group for some lottery ticket at some price (unmatched offers to buy or sell), should they accept a countering outside offer. If the group is in an *equilibrium* state—that is, all possible within-group trades have been exhausted—then arbitrage opportunities cannot exist [48]. From $O$'s perspective, the group appears rational. In other words, by de Finetti's same argument, the group behaves as if guided by a single *representative* probability distribution.

The hypothetical group of bettors in the above example essentially created a market in lottery tickets. Each bettor acts rationally, buying and selling according to his own beliefs and utility for money. The equilibrium prices in the market can be interpreted as the group's consensus probability distribution. In this paper, we examine the reasonableness of this interpretation under various market modeling assumptions.

In Section 2, we cover necessary background material on opinion pools and markets, and introduce our basic market framework. In Section 3, we show that under each of two homogeneity assumptions regarding agents' utility types, equilibrium prices have the same functional form as a standard opinion pool. The so-called *expert weights*, notoriously hard to define in the context of opinion pools [17], correspond to parameters of the agents' utility functions. In Section 4, we present an extended model where agents update their beliefs based on market signals, by treating prices as if they were the beliefs of another individual. In this case, expert weights depend not only on utilities, but on assessments of confidence as well. In Section 5, we describe a situation where the market itself adapts over time; the redistribution of wealth among agents leads to price dynamics that actually mimic rational Bayesian updates. In Section 6, we show how an outside observer can distinguish group behavior from that of an individual, by determining whether equilibrium prices conform to specific properties developed within the opinion pool literature. In Section 7, we describe previous proposals for market-based belief aggregation and disucss related work in experimental economics. Some final remarks appear in Section 8.

## 2 Background, Notation, and Framework

We presume that $N$ agents are uncertain about which of many possible world *states* will eventually obtain. Denote the set of possibilities as $\Omega = \{\omega_1, \omega_2, \ldots\}$, where the $\omega$ are disjoint and exhaustive. There are $S$ *events* of interest (i.e., events for which a consensus is to be reached), $A_1, A_2, \ldots, A_S$, where each $A_j$ is a subset of $\Omega$. Agent $i$ has a subjective probability distribution $\Pr_i(\cdot)$ over $\Omega$; its probability for an event $A_j$ is then just $\Pr_i(A_j) = \sum_{\omega \in A_j} \Pr_i(\omega)$. In general events can overlap, though we consider in detail the special case of *mutually exclusive* events. Agent $i$'s utility for $y$ dollars is $u_i(y)$. When only one agent or one event is under consideration, we simplify by dropping the appropriate subscript(s).

An *opinion pool* combines the probabilistic beliefs of a group into a single, consensus belief, denoted $\Pr_0$.

$$\Pr_0 \equiv f(\Pr_1, \Pr_2, \ldots, \Pr_n)$$

### 2.1 Opinion Pools

A variety of authors have proposed or advocated a corresponding variety of aggregation functions—Genest and Zidek [23] and French [17] provide comprehensive surveys. Three approaches are generally distinguished. The first assumes a single Bayesian decision maker $h$ (real or fictitious, within or outside the group), called the *supra Bayesian*, with a joint distribution over all events and all participants' beliefs. The supra Bayesian updates its beliefs via Bayes's rule, given the "evidence" of everyone else's beliefs. The resulting posterior is taken to be the consensus belief,

$$\Pr_0(\omega | \Pr_1, \ldots, \Pr_N) \propto \Pr_h(\Pr_1, \ldots, \Pr_N | \omega) \Pr_h(\omega). \tag{1}$$

The second approach is to apply a prespecified function, usually some form of weighted average, that maps any set of probability distributions to a singleton. Note that some such functions can be interpreted as the the updating procedure of a supra Bayesian. Each pooling function is usually justified axiomatically, by assuming a "reasonable" set of properties of the aggregate distribution. The two most common and well-studied aggregation functions are the linear and logarithmic opinion pools (LinOP, LogOP). The LinOP is a weighted arithmetic mean of the members' probabilities,

$$\mathrm{Pr}_0(\omega) = \sum_{i=1}^{N} \alpha_i \mathrm{Pr}_i(\omega), \tag{2}$$

and the LogOP is a normalized, weighted geometric mean,

$$\mathrm{Pr}_0(\omega) \propto \prod_{i=1}^{N} [\mathrm{Pr}_i(\omega)]^{\alpha_i}, \tag{3}$$

where the $\alpha_i$ are called *expert weights*, usually nonnegative numbers that sum to one. The LinOP and LogOP can actually be characterized as two instances of a parameterized family of weighted aggregation functions [7]. A third, relatively new, approach to pooling opinions is based on maximum entropy inference. The consensus is the unique probability distribution that maximizes Shannon entropy, chosen from among the distributions that are consistent with all available information, including the experts' beliefs, their past performance, and/or dependencies among experts [35, 46].

The debate over which aggregation method is best continues to rage [2, 7, 30, 49, 74]. Several authors (most emphatically Lindley [36, 37]) argue that the supra Bayesian approach is superior, as it is grounded in standard normative Bayesian theory [6, 43, 44, 70, 73]. However, several conceptual questions arise in characterizing the *single* supra Bayesian responsible for the consensus of a (possibly leaderless) group [23]. If the supra Bayesian is real, is it a member of the group or an outsider? If an outsider, will the group accept its authority to generate the consensus? If a member, which one? If it is fictitious, where do its prior probabilities come from? These must somehow be agreed upon in advance, but then how is *that* consensus reached? A second problem is analytic tractability—the supra Bayesian must have a joint distribution over all members' beliefs and the events in question. In almost every treatment of this method, strong assumptions are made for the supra Bayesian's priors and likelihood function—for example, by assuming that the members' log-odds are normally distributed and independent—in order to yield tractable results [6, 30, 37, 73]. Weighted aggregation functions are not without their own drawbacks. Each method implies different properties for the combined distribution. Many of these properties seem reasonable, but disagreement persists on which are essential.[1] Each of many opinion pools,

[1]For example, Lindley [36] regards the so-called *marginalization property* as an "ad-hockery" while Cooke [7] characterizes any consensus function that does *not* respect it as "downright queer".

including the LinOP and LogOP, can be shown to be the *only* form possible, if one assumes an appropriate, corresponding set of axioms [4, 45, 20, 18, 21, 68]. It follows that no pool can satisfy a sufficiently large set of potentially desirable properties, resulting in several so-called *impossibility* theorems [8, 9, 19, 22, 34, 66]. Another point of contention is how best to determine the expert weights. In most cases they are chosen in an ad hoc manner to encode some measure of confidence, reliability, or importance [2, 17, 72]. Some more formal methods to derive weights have been proposed, by making assumptions concerning the form of, or interdependence among, participants' beliefs [7, 12, 30, 44], or through iterative self-weighting procedures [11]. Little work, in any of the traditional categories of opinion pools, explicitly addresses truth incentives for reporting either self-assessed weights or the probabilities themselves.

## 2.2 Securities Markets and Rational Expectations

In certain real markets, a common interpretation of prices is as an amalgamation of opinions. For example, the price of a stock represents the "market evaluation" of the expected present value of future dividends, and odds in a horse race summarize the bettors' beliefs about the winning horse's identity. Under certain regularity conditions, the exact dependence of prices on beliefs can be characterized [28, 61, 63, 71], and compared with the opinion pools of Section 2.1. In contrast to the Bayesian or axiomatic approaches, economic models are usually built upon behavioral assumptions (e.g., that agents are rational utility maximizers) and descriptions of the price-formation process.

In a classical, pure-exchange economy *without* uncertainty, agents have *endowments* of commodities or *goods* and exchange them according to their utility for consumption. Under uncertainty, endowments and utilities are in general state-dependent, and thus so are the spot prices of goods. Agents will generally exhibit varying marginal value of wealth in the different states. Consequently, if they are risk-averse, they will desire to *hedge* or *insure* against risk by distributing wealth across states.[2] The *Arrow-Debreu securities market* is the fundamental theoretical framework in economics and finance for resource allocation under uncertainty [1, 13, 39]. A *security*, denominated in money or other exchangeable good, pays off variously con-

[2]For example, insuring a package delivery transfers wealth from the *package-delivered* state to the *package-lost* state.

tingent upon the realization of an uncertain state. In the standard model, agents trade securities prior to revelation of the world state, and then trade real goods once the state is known. A market is termed *complete* if it contains at least $|\Omega| - 1$ *linearly independent* securities,[3] plus spot markets in the original goods. Such a market guarantees, under classical assumptions, that equilibrium entails a Pareto optimal allocation of both goods and risk.

When one unit of each security pays out one dollar, it follows, from de Finetti's argument, that the equilibrium prices in a securities market form a legal probability distribution [13]. Moreover, all agents' *risk-neutral* probabilities are in agreement, and coincide with these prices [48]. Risk-neutral probabilities are defined as $\Pr^{\mathrm{RN}}(\omega) \propto \Pr(\omega)u'(Y_\omega)$, where $Y_\omega$ is the agent's payoff in state $\omega$ and $u'$ is the derivative of utility [47]. A hypothetical agent with transformed probabilities $\Pr^{\mathrm{RN}}(\omega)$ and transformed (state-dependent) *risk-neutral* utility $u^{\mathrm{RN}}(y) \equiv u(y)/u'(Y_\omega)$ has precisely the same betting behavior as an agent with $\Pr(\omega)$ and $u(y)$. Since any two risk-neutral agents with differing probabilities will desire to trade, the market cannot be in equilibrium unless all agents' risk-neutral probabilities agree, and equal the prices [48].

To derive explicit equations for price, early work centered on the *prior information* (PI) equilibrium concept, where agents' beliefs are presumed to remain fixed during price formation [28, 61, 63]. As discussed in Section 3, prices can in this case be characterized a combination of the agents' beliefs, but cannot reasonably be regarded as an accumulation of diverse *information*. The latter, perhaps more satisfying, interpretation does become admissible under the *rational expectations* (RE) equilibrium assumption [24, 38]. This theory asserts that prices actually reflect all of the information available to all agents. Even when some agents have exclusive access to inside information, prices equilibrate exactly as if everyone had access to all information. The procedural explanation is that prices *reveal* to the ignorant agents any initially private information; that is, agents *learn* from prices. Several authors show that, if agents begin with identical priors and disparate evidence, repeated observation of some aggregate statistic (e.g., price) will converge to a consensus on posteriors, for various sufficient statistics [25, 41, 42, 51]. We describe a supra Bayesian model of RE in Section 4.

Both Rubinstein [61] and Wilson [71] develop formal conditions, examined

[3] This is often realized with one security paying out in each state of nature, though any set whose payoff-by-state matrix has rank $|\Omega| - 1$ will do.

further in Section 3.1.3, under which agents in a securities market behave collectively in a manner indistinguishable from a single, *composite* agent whose beliefs equal the equilibrium prices [28]. Raiffa [59] (pp. 211–216) recounts an example, due to Pratt, where the optimal group action requires randomization, and thus a composite agent *cannot* exist. Nevertheless, Kasanen and Trigeorgis [32] demonstrate empirically that the analogy of a market as a representative investor can sometimes apply, even outside the theoretical bounds. The composite-agent construction constitutes perhaps the most compelling interpretation for prices as consensus beliefs.

## 2.3 A Securities Market Model for Belief Aggregation

Several economic constructs central to the study of securities markets (e.g., spot markets in real goods, consumption, production, and endowments) become somewhat extraneous in a model focusing on belief aggregation. In what follows, we develop a stripped-down model sufficient for our purposes.

The market consists of securities "\$1 if $A_j$", abbreviated $\langle A_j \rangle$, for every event of interest $A_j$. One unit of the security pays off one dollar if the event occurs, and nothing otherwise. If an agent must pay a price $p$ per unit of security $\langle A \rangle$, then purchasing $x$ units is equivalent to accepting a lottery with payoff $(1-p)x$ if $A$ occurs, and $-px$ otherwise. Positive $x$ indicates a quantity to buy, and negative $x$ a quantity to sell. *Either* action incurs risk: buying $\langle A \rangle$ at price $p$ yields exactly the same lottery as selling $\langle \bar{A} \rangle$ at price $1-p$.

Agents are *competitive*; that is, they maximize expected utility at a given set of prices, or equivalently, ignore any effect their own demand may have on prices. We assume that utility for money is independent of state and monotonically increasing. To avoid unbounded demand, we further assume that agents are *risk averse* [33, 58], that is, have concave $u$, always preferring the expected payoff of a lottery to the lottery itself. We also restrict attention to linearly independent securities (i.e., none can be replicated as a combination of others). Without this, agents could buy one security and sell the equivalent portfolio without risk, and demand may be nonunique or unbounded.[4] We consider the PI and RE equilibrium conditions in Sections 3 and 4, respectively.

In a market of $S$ securities, let $x^{(j)}$ represent a quantity of security

[4]This assumption implies that each event is strictly uncertain ($0 < \Pr(A) < 1$).

$\langle A_j \rangle$, and $p^{(j)}$ the security's price. The agent's utility for a bundle $\mathbf{x} = \langle x^{(1)}, x^{(2)}, \ldots, x^{(S)} \rangle$ of securities is its expected $u$:

$$U(\mathbf{x}) = \sum_{\omega \in \Omega} \Pr(\omega) u(Y_\omega), \tag{4}$$

where

$$Y_\omega = \sum_{j=1}^{S} (1_{\omega \in A_j} - p^{(j)}) x^{(j)}$$

is the payoff in state $\omega$, with $1_{\omega \in A_j}$ the indicator function that equals one if $\omega \in A_j$, zero otherwise.

The agent's *demand function*, $\mathbf{x}(\mathbf{p}) = \langle x^{(1)}(\mathbf{p}), x^{(2)}(\mathbf{p}), \ldots, x^{(S)}(\mathbf{p}) \rangle$, represents the quantities of securities that maximize utility (4) at prices $\mathbf{p} = \langle p^{(1)}, p^{(2)}, \ldots, p^{(S)} \rangle$. Concavity of $u$ (risk aversion) entails concavity of $U$, and hence any critical point of (4) is a global maximum.

Equation 4 captures the optimization problem of a single agent faced with the decision to invest in securities. In an economy of $N$ agents, each continually maximizing (4), prices adjust until all buy orders (positive $x^{(j)}$) match with sell orders (negative $x^{(j)}$) for each security $\langle A_j \rangle$. A market is in *competitive equilibrium*[5] at prices $\mathbf{p}$ if and only if

$$\sum_{i=1}^{N} \mathbf{x}_i(\mathbf{p}) = \mathbf{0}. \tag{5}$$

Most opinion pools require that agents directly, and truthfully, reveal their probabilities to a central location. Scoring rules [72], promissory notes [10], or other incentive devices could be leveraged to elicit honest assessments, though these may require outside subsidies and, as French [17] points out, a group may collectively have greater incentive to lie. Kadane and Winkler [31] demonstrate that the probabilities revealed by such elicitation aids are inextricably confounded with utility, whenever the agent is risk averse and may

[5]This formulation assumes a Walrasian environment, where exchange occurs only after equilibrium is reached. In an alternative formulation, agents make incremental trades for securities, at mutually beneficial prices, until no such opportunities exist. Under our assumptions, such a process would also reach an equilibrium, where agents' risk-neutral probabilities are equal (and thus represent clearing prices). Solutions for equilibrium prices for the case of negative exponential utility (derived below) would remain the same, though our other analytic results are specific to the Walrasian formulation. The dynamic model of Section 5 reflects some aspects of the incremental approach.

own stakes in the events. The securities market framework provides built-in rewards for participation, honesty,[6] and the collection of cost-effective evidence. Agent beliefs are revealed only implicitly through their trades. In fact, if measures are taken to ensure the privacy of bids, the *only* personal information directly revealed is an agent's risk-neutral probability distribution at equilibrium, which equals the consensus anyway. Since supply and demand are balanced in equilibrium, the mechanism requires no subsidy, and so the only cost involved is that of organizing and running the market.

An agent whose beliefs differ from the going prices will desire to trade, and its choices will always yield, from its perspective, nonnegative expected payoff. The quantity demanded will be bounded as long as utility is sufficiently convex[7] and securities are linearly independent. If all agents maintain identical beliefs across all events, then equilibrium prices will equal these beliefs. This property, called *unanimity*, is often considered a minimal standard of acceptability for an opinion pool.

# 3 Market Pooling Functions

In this section, we characterize the dependence of prices on beliefs within our model, concentrating first on the case where all events are mutually exclusive. We consider, in turn, two specific assumptions regarding agents' utility forms.

## 3.1 Disjoint Securities

Events are disjoint but, to retain linear independence, are not exhaustive. For notational purposes, define the event $A_0 \equiv \bar{A}_1\bar{A}_2\cdots\bar{A}_S$, though it does not have a corresponding security.[8] Mutually exclusive events are certainly a special case, however any event space can partitioned into securities representing disjoint events.

[6] Incentives exist, though the market is not fully *incentive compatible* [39].

[7] Risk aversion alone does not guarantee bounded demand. One sufficient condition is that $\lim_{y\to\infty}\frac{u(y)}{u(-y)} = 0$.

[8] $\langle A_0\rangle$ can be replicated by selling equal amounts of all securities.

### 3.1.1 Negative Exponential Utility

First, suppose that all agents have negative exponential utility for money, or $u(y) = -e^{-cy}$. This utility is synonymous with *constant absolute risk aversion* (CARA), where $c$ is the coefficient of risk aversion, or $1/c$ the risk tolerance.

Each agent's unique optimal demand is:

$$x^{(j)}(\mathbf{p}) = \frac{1}{c} \ln \left( \frac{\Pr(A_j)}{p^{(j)}} \cdot \frac{1 - p^{(1)} - p^{(2)} - \cdots - p^{(S)}}{\Pr(A_0)} \right). \tag{6}$$

The condition for competitive equilibrium (5) is $\sum_{i=1}^{N} x_i^{(j)}(\mathbf{p}) = 0$, for all $j = 1 \ldots S$. Solving for $\mathbf{p}$, we arrive at the following observation.

**Observation 1 (CARA $\Rightarrow$ LogOP)** *In a market of disjoint securities, each agent having CARA $c_i$, the prices*

$$p^{(j)} = \frac{\prod_{i=1}^{N} [\Pr_i(A_j)]^{\alpha_i}}{\prod_{i=1}^{N} [\Pr_i(A_0)]^{\alpha_i} + \prod_{i=1}^{N} [\Pr_i(A_1)]^{\alpha_i} + \cdots + \prod_{i=1}^{N} [\Pr_i(A_S)]^{\alpha_i}} \tag{7}$$

*form a competitive equilibrium, where $\alpha_i = (1/c_i)/\sum_l (1/c_l)$ is the normalized risk tolerance of agent $i$.*

These prices are the same probabilities as derived by the LogOP (3), where the expert weights are normalized risk tolerances. The market model, therefore, provides one way to ground a well-known centralized pooling mechanism in terms of individual behavior. Observation 1 can be thought of as a special case of Rubinstein's result [61] within our simplified context, the major difference being that the normalization of prices falls out naturally in our derivation, yielding a legal probability distribution. Wilson [71] also arrives at the weighted geometric average relationship when agents (with CARA) share the profits of a security that pays off contingent on a single, continuous random variable. Nau and McCardle [48] derive a similar result for more than one continuous variable when all beliefs are multivariate normal, and quadratic options are available in addition to the standard securities.

Observation 1 provides a decision-theoretic interpretation for the notoriously slippery concept of expert weights. In the usual interpretation, the exponents in (7) encode some sort of degree of expertise, confidence, or reliability, and are almost always chosen in an ad hoc manner [2, 17, 23, 50, 72].

In the market model, weights are *derived*, and are directly related to a well-understood quantity, namely risk aversion. It may seem that weights *should* relate to a notion of confidence, however loosely defined, rather than to risk tolerance. In Section 4, we propose an extension of the market model in which agents learn from prices, and show how a confidence-related interpretation arises naturally.

#### 3.1.2 Generalized Logarithmic Utility

Next consider an assumption that all agents exhibit *generalized logarithmic utility* (GLU) for money: $u(y) = \ln(y + b)$. With the further constraint that $b > 0$, GLU is a standard *decreasingly* risk-averse form. Note that $u$ is imaginary on the interval $(-\infty, -b)$. Since utility approaches $-\infty$ as $y \to -b$ from the right, and $u$ is increasing, we can safely extrapolate that it is $-\infty$ for all $y \leq -b$. The upshot is that the agent would *never* put $b$ or more dollars at risk, regardless of probability assessments. Note that any wealth accumulated in dollars is additive with $b$, so a natural interpretation of $b$ is as the agent's initial wealth.

Unique optimal demand is

$$x^{(j)}(\mathbf{p}) = b\left(\frac{\Pr(A_j)}{p^{(j)}} - \frac{\Pr(A_0)}{1 - p^{(1)} - p^{(2)} - \cdots - p^{(S)}}\right), \tag{8}$$

and equilibrium prices can again be characterized in closed form.

**Observation 2 (GLU $\Rightarrow$ LinOP)** *In a market of disjoint securities, each agent having GLU with initial wealth $b_i$, the prices*

$$p^{(j)} = \sum_{i=1}^{N} \beta_i \Pr_i(A_j) \tag{9}$$

*form a competitive equilibrium, where $\beta_i = b_i / \sum_l b_l$ is agent $i$'s normalized initial wealth.*

This is exactly the LinOP formula (2), where the expert weights are again directly related to risk tolerance. This result can again be seen as a simplified derivation of that by Rubinstein [61, 62, 63].

### 3.1.3 Composite Agents

With disjoint securities, and under each of the utility assumptions (CARA, GLU) the market as a whole behaves as a single, *composite* agent with probabilities equal to prices. An outsider, offering securities to the group at various prices, cannot differentiate their demand response from that of an individual.

**Observation 3 (Composite CARA Agent)** *Suppose that, in a market of disjoint securities, all agents exhibit CARA with risk aversion $c_i$, and effect equilibrium prices $\mathbf{q}$. Then their total demand for security $\langle A_j \rangle$ offered at price $p^{(j)}$ is:*

$$x_0^{(j)} = \frac{1}{c_0} \ln \left( \frac{\Pr_0(A_j)}{p^{(j)}} \cdot \frac{1 - p^{(1)} - \cdots - p^{(S)}}{\Pr_0(A_0)} \right),$$

*or equal to that of a composite agent with beliefs $\Pr_0(A_j) = q^{(j)}$ and CARA with risk tolerance $1/c_0 = \sum_{i=1}^{N} 1/c_i$.*

**Observation 4 (Composite GLU Agent)** *Suppose that, in a market of disjoint securities, all agents exhibit GLU with initial wealth $b_i$, and effect equilibrium prices $\mathbf{q}$. Then their total demand for security $\langle A_j \rangle$ offered at price $p^{(j)}$ is:*

$$x_0^{(j)} = b_0 \left( \frac{\Pr_0(A_j)}{p^{(j)}} - \frac{\Pr_0(A_0)}{1 - p^{(1)} - p^{(2)} - \cdots - p^{(S)}} \right),$$

*or equal to that of a composite agent with beliefs $\Pr_0(A_j) = q^{(j)}$ and GLU with initial wealth $b_0 = \sum_{i=1}^{N} b_i$.*

Furthermore, any subset of the agents can be aggregated together, interacting with the rest of the system as if it were an individual. Wilson [71] first described conditions for a composite agent, which he called a *surrogate*, extended later by Rubinstein [61] and covered in Huang and Litzenberger's textbook [28]. Notice that group rationality emerges from the individual decisions of its members, unlike the classic supra Bayesian, whose existence is imposed at the outset. Assuming positive $c_i$, the composite agent's risk aversion is strictly less than that of any individual; in other words, the group together is willing to take on more risk than any proper subset of its members. Raiffa [59] elucidates the connection between risk sharing and group decision making, discussing the case of CARA in depth.

A mixed population, consisting of some agents with CARA and some with GLU, can be summarized as *two* composite agents, each representing its respective subpopulation. In this case, numeric solutions, or approximate formulas, for price are possible.

### 3.2 Arbitrary Securities

Deriving closed-form pricing formulae for arbitrary securities is difficult, even under homogeneous utility constraints, though a few things can be said. First, if $S = |\Omega| - 1$, and securities are linearly independent, then the market is *complete* and prices define a full probability distribution over $\Omega$. Moreover, the distribution is irrespective of which $|\Omega| - 1$ states are chosen [39]. In particular, assuming either CARA or GLU, the consensus is the same as implied by (7) or (9), respectively, in a market with securities corresponding to all but one state.

More can be said as well when agents have CARA, and certain independence conditions among events hold.

**Proposition 5 (Separation of independent events)** *Suppose that an agent with CARA considers $A_1$ independent of all other events. Then its optimization problem for $x^{(1)}$ is separable, and reduces to the single-security case.*

**Proof.** Let $X = \{A_2, \ldots, A_S\}$ be the set of all events except $A_1$. Utility (4) is

$$\begin{aligned} & -\sum_X \left(\Pr(A_1)\Pr(X)e^{-c[(1-p^{(1)})x^{(1)}+Y_X]} + \Pr(\bar{A}_1)\Pr(X)e^{-c[-p^{(1)}x^{(1)}+Y_X]}\right) \\ = & \; -\left(\Pr(A_1)e^{-c(1-p^{(1)})x^{(1)}} + \Pr(\bar{A}_1)e^{-c(-p^{(1)}x^{(1)})}\right)\sum_X \Pr(X)e^{-cY_X} \end{aligned}$$

where $\sum_X$ is a sum over all possible joint outcomes of events in $X$, and $Y_X$ is the payoff in the selected outcome. The optimal choice for $x^{(1)}$ is independent of demand for other securities, and has the same solution as if $\langle A_1 \rangle$ were offered in isolation. $\Box$

It follows that when *all* agents have CARA, and everyone agrees that $A_1$ is independent of all other events, the equilibrium price is given by (7) as if $\langle A_1 \rangle$ were the only security present ($S = 1$). Moreover, because its price does not depend on the outcome of any other event, $A_1$ remains independent of all other events in the distribution implied by equilibrium prices. Elsewhere [54], we derive more general independence preservation properties of the LogOP; implications for securities markets will be addressed in a future report.

In general, including the case of GLU, independence of events does *not* lead to separable demand. Even if everyone agrees that an event is independent of all others, a dependence will generally be introduced among the equilibrium prices.

Any prices that form incoherent probabilities (that allow arbitrage) cannot be part of an equilibrium. If any agent, with *any* utility function that is monotonically increasing for money, believes that some outcome is certain ($\Pr(A) \in \{0, 1\}$), or that one subset of events is logically equivalent to another, then it will effectively dictate that relationship in the consensus, by demanding infinite quantities if prices suggest otherwise.[9]

# 4 Learning from Prices: A Supra Bayesian Model of Rational Expectations

We have thus far assumed a classical *prior information* (PI) equilibrium in which agents' beliefs are fixed. Yet our advocacy of a market-based pooling device presupposes that exogenous prices do reveal meaningful aggregate information. If the agents themselves agree, then we would expect them to update their beliefs based on the observed prices. Anecdotally, an agent whose probability is poorly aligned with the going price may begin to suspect that everyone else knows something that it does not, and adjust its belief accordingly. This is a motivating premise of the theory of *rational expectations* (RE) [24, 38]. The equilibrium conditions remain that all agents are optimizing and that total demand is zero. Implicit in the former requirement however is the dependence of each agent's *posterior* probability on the equilibrium prices. In terms of the model outlined in Section 2.3, we replace each fixed prior probability $\Pr(\cdot)$ in all equations with the conditional probability $\Pr(\cdot|\mathbf{p})$. This substitution in general transforms closed-form solutions for price into implicit equations. In some cases, explored below, we are able to resolve $\mathbf{p}$.

What is a reasonable model for $\Pr(\cdot|\mathbf{p})$? Recall that $\mathbf{p}$ has all the proper-

[9] Recall that in this treatment, we impose no restrictions on demand magnitude, either by budget constraints, or through market rules.

ties of a probability distribution and that, at least under certain homogeneity conditions, a formal equivalency between the market and a composite individual can be established. In a strong sense, the agent faces the same problem as one decision maker updating beliefs based on the probability assessments of one expert, where the "expert" in this case is the market itself. This is precisely the scenario where the supra Bayesian method (1) is both justified and well-developed. The philosophical concerns reviewed in Section 2.1 surrounding the identification of the supra Bayesian are not relevant here; moreover, some of the computational burden is reduced, since the agent updates based only on one composite distribution rather than all $N-1$ distributions. Note that this interpretation makes sense only when the agent does not consider the effect of its own decisions on prices.

Depending on how the situation is modeled, several supra Bayesian formulations are plausible [3, 6, 37, 44, 70]. We adopt one similar to that proposed by Morris [45], as it ultimately yields closed-form results for equilibrium prices. Consider first a market with a single security $\langle A\rangle$. The agent envisions the corresponding event $A$ as one trial in a series of identical, independent Bernoulli trials with probability of success $\Pi$. The parameter $\Pi$ is itself uncertain; assume that, after witnessing $s$ positive outcomes or *successes* among $n$ repetitions of the event, the agent's prior probability is $\Pr(\Pi=\pi) \propto \pi^{s-1}(1-\pi)^{n-s-1}$, or a Beta distribution with parameters $s$ and $n-s$. The agent is myopic, or acts as if the market is a one-shot game, so uses $\Pr(A) = E[\Pi] = s/n$ to guide its decisions.

As judged by the agent, the "market" (i.e., everyone else) has observed $s'$ successful outcomes of the event $A$ in a total of $n'$ trials. The likelihood function $\Pr(p|\Pi=\pi) \propto \pi^{s'}(1-\pi)^{n'-s'}$ is the corresponding Binomial distribution. For consistency, it must be that $p = s'/n'$. The desired posterior probability

$$\Pr(\Pi=\pi|p) \propto \Pr(p|\Pi=\pi)\Pr(\Pi=\pi)$$

is straightforward to compute, as the Beta and Binomial distributions form a natural conjugate class [60]. The result is another Beta distribution with parameters $s+s'$ and $n+n'-(s+s')$.

Reformulating in terms of $A$, we find that

$$\begin{aligned}
\Pr(A|p) &= E[\Pi|p] \\
&= \frac{s+s'}{n+n'} = \frac{n}{n+n'}\cdot\frac{s}{n} + \frac{n'}{n+n'}\cdot\frac{s'}{n'} \\
&= w\Pr(A) + (1-w)p,
\end{aligned}$$

where $w \equiv n/(n+n')$. The agent's posterior probability given the price is just a weighted average of its prior and the price. The weighting term captures the agent's perception of its own confidence, expressed in terms of number of observations, as compared to the market's.

The independent Bernoulli trials model need not be taken literally; one may prefer to interpret $s/n$ and $s'/n'$ as subjective measures of probability, and $n$ and $n'$ as measures of relative precision. In particular, the four parameters need not be constrained to the integers.

Generalize this to a market of $S$ disjoint securities, and assume that the agents exhibit GLU. The demand function (8) becomes

$$x^{(j)}(\mathbf{p}) = wb\left(\frac{\Pr(A_j)}{p^{(j)}} - \frac{\Pr(A_0)}{1 - p^{(1)} - \cdots - p^{(S)}}\right).$$

In words, the relative stock that an agent places on its own prior probability versus public prices serves to attenuate demand by a multiplicative factor.

In a market of likewise agents, the equilibrium prices still obey (9), though these are now implicit equations:

$$\begin{aligned} p^{(j)} &= \sum_{i=1}^{N} \beta_i \Pr_i(A_j | p^{(j)}) \\ &= \sum_{i=1}^{N} \frac{w_i b_i}{\sum_l b_l} \Pr_i(A_j) + \sum_{i=1}^{N} \frac{(1-w_i) b_i}{\sum_l b_l} p^{(j)}. \end{aligned}$$

Solving for $p^{(j)}$, we find that

$$p^{(j)} = \sum_{i=1}^{N} \frac{w_i b_i}{\sum_l w_l b_l} \Pr_i(A_j). \tag{10}$$

Prices retain the form of the LinOP, but with weights proportional to the product of risk tolerance and self-assessed confidence. The introduction of a weighting factor that depends on the precision of probability assessments rather than solely on parameters of utility may be more palatable to advocates of the standard opinion pools. Why should consensus probabilities depend on utilities at all? Hylland and Zeckhauser [29] prove that any composite agent whose $\Pr_0$ and $u_0$ are aggregated separately must *fail* to either (1) uphold unanimous decisions, (2) maintain unanimous agreement on probabilities, or (3) allow $\Pr_0$ to depend on more than one agent.

Under some conditions we can show that prices do encode the "correct" aggregate probabilities. That is, the market reaches a RE equilibrium where prices are exactly as if all agents had witnessed all observations for all events. Each agent has seen $s_i^{(j)}$ positive occurrences of $A_j$ out of $n_i$ total observations. Suppose that agents' observations are independent, that each agent knows its relative number of observations (i.e., $w_i = n_i / \sum_h n_h$), and that all agents have the same risk tolerance. Then (10) becomes

$$p^{(j)} = \frac{\sum_{i=1}^{N} s_i^{(j)}}{\sum_{i=1}^{N} n_i}.$$

Even though the agents' posterior probabilities are *not* in complete agreement, this is precisely the same equilibrium that would have arisen had all agents (or a composite agent) seen all observations.

Finally, analogous derivations apply when agents have CARA and the posterior update rule is a weighted *geometric* average of prior and price:

$$\Pr(A_j | p^{(j)}) \propto [\Pr(A_j)]^w [p^{(j)}]^{(1-w)}.$$

# 5 Wealth Dynamics and Market Adaptation

We have seen that markets can mimic both the LinOP and LogOP formulas, and that simple learning from prices can lead to confidence-based weights. These results highlight the market's flexibility, showing that popular methods for deriving consensus are in fact special cases of the market mechanism. More heterogeneous mixtures of agents entail still other pooling functions, most of which are unlikely to admit closed-form expression. Yet this begs the question of which aggregation rule is the most appropriate and/or accurate. In contrived scenarios, one or the other is clearly warranted. For example, if the agents' beliefs are distributed such that their mean is the "true" probability, then the LinOP with equal weights will recover it.[10] Similarly, if the agents' log-odds are distributed such that their mean is the true log-odds, then the LogOP is justified. Of course, it is unreasonable to assume knowledge of the distribution of agent beliefs in relation to the true probability. In this section, we argue that, over time, the market may implicitly *learn*

[10]This is precisely the construction of the comparative simulation experiments of Ng and Abramson [49]; in this light, their finding that the LinOP is best is perhaps not surprising.

this distribution. Markets are not only flexible but *adaptive*, rewarding accuracy and punishing inaccuracy. Agents with better probabilistic models gain wealth and, consequently, greater influence on prices.

We report simulation results in support of this claim. The market consists of a single security traded repeatedly over the course of $T$ *periods* or rounds. With only one security per period, we can let the subscript on events $A_t$ and the superscripts on price $p^{(t)}$ and demand $x^{(t)}$ index time. Unbeknownst to the agents, each event unfolds as an independent Bernoulli trial with probability of success $\pi$. At the beginning of time period $t$, the realization of $A_t$ is unknown and agents trade in $\langle A_t \rangle$ until equilibrium. Then the outcome—either $A_t$ or $\bar{A}_t$—is revealed, and the agents' holdings pay off accordingly. As time period $t+1$ begins, the outcome of $A_{t+1}$ is uncertain. Agents trade in the $t+1$ period security until equilibrium, the event outcome is revealed, payoffs are collected, and the process repeats.

Agents are *myopic*, or make decisions as if the current time period were the final one, maximizing expected utility (4) for each period in isolation. They do not update their beliefs based on either observations or prices.

Figure 1 plots the price of the security over fifty time periods, in a market composed of one hundred agents, each with GLU, initial wealth $b_i^{(0)} = 1$, and $\Pr_i(A)$ generated randomly and uniformly on $(0, 1)$. Let $b_i^{(t)}$ indicate agent $i$'s wealth after period $t$ ends. For this simulation, we set $\pi = 0.5$. For comparison, the figure also shows the *observed frequency*, or the number of times that $A$ has occurred divided by the number of periods. The market price tracks the observed frequency extremely closely. Note that price changes are due only to a transfer of wealth from inaccurate agents to accurate agents, who then wield more power in the market; individual beliefs remain fixed.

Figure 2 provides further insight into this transfer of wealth. The graph provides a snapshot of agents' wealth $b_i^{(15)}$ versus their belief $\Pr_i(A)$ after period 15. In this run, $A$ has occurred in three out of the fifteen trials. The maximum in wealth obtains at a belief near 3/15. The solid line in the figure is a Beta distribution with parameters $3 + 1$ and $12 + 1$, scaled only by a constant. This distribution is precisely the posterior probability of success that results from the observation of 3 successes out of 15 independent Bernoulli trials, when the prior probability of success is uniform on (0,1). The fit is essentially perfect, and can be proved in the limit.

**Proposition 6 (Market Beta)** *Suppose that agents with GLU, identical initial wealth* $b_i^{(0)} = b$, *and* $\Pr_i(A)$ *generated randomly and uniformly on*

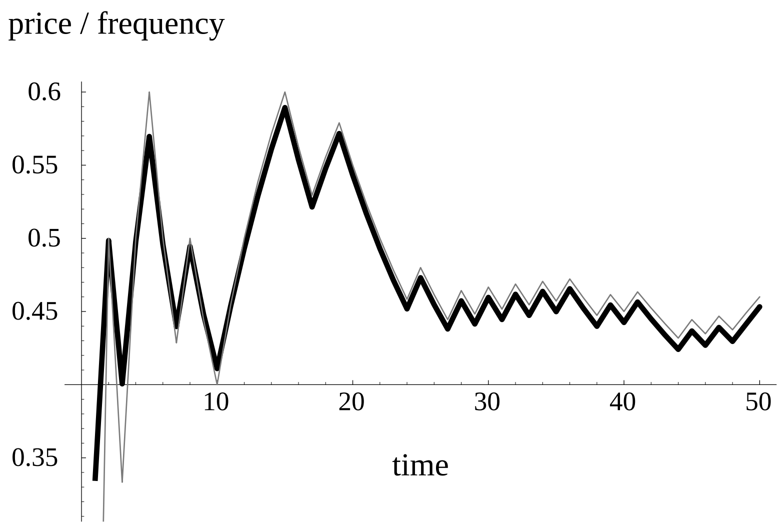


Figure 1: A single security's price (dark line) and the observed frequency (gray line) of the event over fifty time periods. The market consists of one hundred agents with GLU and $b_i^{(0)} = 1$.

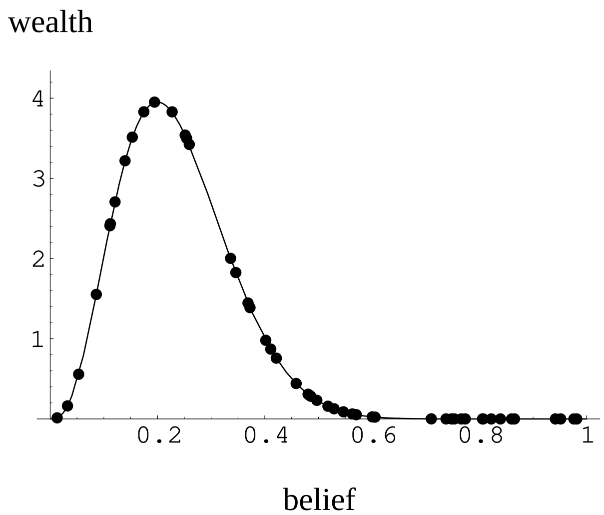


Figure 2: Wealth versus belief for fifty agents with GLU and $b_i^{(0)} = 1$, after fifteen time periods. The event has occurred in three of the fifteen trials. The solid line is the posterior Beta distribution consistent with the observation of three successes in fifteen independent Bernoulli trials.

*$(0,1)$ trade in a single-security market over $n$ time periods. For notational brevity, let $\pi_i \equiv \Pr_i(A)$. Then for sufficiently large $N$, the ratio of any two agents' wealth is*

$$\frac{b_h^{(n)}}{b_i^{(n)}} \approx \frac{(\pi_h)^s(1-\pi_h)^{n-s}}{(\pi_i)^s(1-\pi_i)^{n-s}}, \tag{11}$$

*where $s$ is the number of positive occurrences of $A$ observed in the $n$ periods.*

**Proof.** We proceed by induction. The base case $n = 0$ holds by construction. Assume that (11) holds for $n = t$. Therefore, agent $h$'s wealth after period $t$ is $b_h^{(t)} = \kappa(\pi_h)^s(1-\pi_h)^{t-s}$, where $\kappa$ is independent of $h$. The agent's demand for $\langle A_{t+1}\rangle$ is according to (6) with $S = 1$. The price $p^{(t+1)}$ is a weighted sum of beliefs $(\sum_{i=1}^N \pi_i b_i^{(t)})/\sum_{l=1}^N b_l^{(t)}$ as prescribed by (9). Because beliefs $\pi_i$ are generated uniformly on (0,1), as $N \to \infty$, this summation approaches

$$p^{(t+1)} \to \int_0^1 \pi \cdot \frac{\kappa(\pi)^s(1-\pi)^{t-s}}{\int_0^1 \kappa(u)^s(1-u)^{t-s}du} d\pi = E[\mathrm{Beta}(s+1, t-s+1)] = \frac{s+1}{t+2}.$$

We substitute this price into the demand equation and compute the agent's revised wealth $b_h^{(t+1)}$ after appropriate payoffs under both possible outcomes:

**Case $A_{t+1}$.** After a positive outcome, the agent's new wealth is $b_h^{(t+1)} = \kappa(\pi_h)^{s+1}(1-\pi_h)^{t-s}(1/p^{(t+1)})$, which is proportional to a Beta distribution reflecting one additional successful trial than at time $t$. The rightmost factor depends only on price, and will be identical for all agents.

**Case $\bar{A}_{t+1}$.** Revised wealth is $b_h^{(t+1)} = \kappa(\pi_h)^s(1-\pi_h)^{t-s+1}(1/(1-p^{(t+1)}))$, or proportional to a Beta distribution reflecting an additional trial without success.

The period $t+1$ wealth of agent $i$ is analogous, and so (11) holds for $n = t+1$. □

Although individual agents are not adaptive, the market's composite agent appears to compute a proper Bayesian update. Specifically, wealth is reallocated proportionally to a Beta distribution corresponding to the observed number of successes and trials, and price is approximately the expected value of this Beta distribution.[11] Moreover, this correspondence holds

[11] As $t$ grows, this expected value rapidly approaches the observed frequency plotted in Figure 1.

regardless of the number of successes or failures, or the temporal order of their occurrence. A kind of collective Bayesianity *emerges* from the interactions of a group of myopic, fixed-belief agents subject to the reward structure of a securities market.

In our simulations, we observe that wealth among a subpopulation of GLU agents, mixed together also with CARA agents, continues to coincide with the appropriate Beta distribution.

# 6 Exposing a Group

Under what conditions can the collective behavior of a group be differentiated from that of an individual? Can an outsider $O$, by querying the group's demand for securities at various prices, determine whether $N > 1$? The composite agent results in Section 3.1.3 identify conditions under which no distinction can be made within the context of a *fixed* market configuration. In this section, we discuss how $O$ might deduce the makeup of the group by examining measurements collected under different configurations.

Assume that agents are myopic, optimizing (4) across *currently available* securities. The agents do not attempt to predict what $O$ will offer next, but simply react according to whatever markets exist at the moment. $O$ has the power to open and cancel markets at will; the agents are essentially memoryless, and blindly recompute their optimization problem in the new environment, with negligible animosity toward $O$'s probing. Agents may have state-dependent utilities or prior stakes in events, though the latter change only via trades in available securities. That is, the agents are isolated from outside markets, or any other betting opportunities, beyond those offered by $O$. Beliefs are assumed constant throughout the observation procedure.

Suppose that $O$ first opens a market in a single security and observes the equilibrium price. $O$ then withdraws this market and opens two more in disjoint securities that form a partition of the first. That is, $O$ determines both of the following:

- the equilibrium price $p^{\langle A_1 \rangle}$ for a single security $\langle A_1 \rangle$
- the equilibrium prices $p^{\langle A_1 A_2 \rangle}$ and $p^{\langle A_1 \bar{A}_2 \rangle}$ for securities $\langle A_1 A_2 \rangle$ and $\langle A_1 \bar{A}_2 \rangle$

If $N = 1$ then, by de Finetti's argument,

$$p^{\langle A_1 \rangle} = p^{\langle A_1 A_2 \rangle} + p^{\langle A_1 \bar{A}_2 \rangle}. \tag{12}$$

Thus, if (12) does *not* hold, we can be certain that $N > 1$. Consensus beliefs that always adhere to (12) satisfy what is called the *marginalization property* (MP). Under exceptional circumstances—for example, all agents have GLU and thus prices form a LinOP, or all beliefs are identical—Equation 12 is upheld by a group as well [20]. Generically, however, prices revealed by a group will *not* conform to (12).

Now imagine that $O$ conducts, in turn, three experiments to elicit the following information:

- the equilibrium price $p^{\langle A_1 A_2 \rangle}$ for security $\langle A_1 A_2 \rangle$
- the equilibrium price $p^{\langle A_1 | A_2 \rangle}$ for *conditional* security[12] $\langle A_1 | A_2 \rangle$
- the equilibrium price $p^{\langle A_2 \rangle}$ for security $\langle A_2 \rangle$

An individual necessarily reveals

$$p^{\langle A_1 A_2 \rangle} = p^{\langle A_1 | A_2 \rangle} p^{\langle A_2 \rangle}. \tag{13}$$

An opinion pool that always respects (13) satisfies what we call the *family aggregation* (FA) property, named in the context of graphical models. We have shown elsewhere [54] that *no* combination of beliefs can simultaneously satisfy FA, unanimity, and nondictatorship.[13] The latter two properties are implicit in the market design. It follows that, even when the conditions for a standard composite agent are met (all have CARA or all have GLU), Equation 13 is generically true if and only if $N = 1$.

# 7 Related Work

Eisenberg and Gale [14] propose a betting-based aggregation mechanism built not upon the securities market framework, but rather on the common method for deriving odds in horse races. They consider a pari-mutuel scheme where agents place bets across a partition of events, yielding a consensus probability equal to the proportion of the total bet on each event. If event $A$ obtains,

[12] A conditional security $\langle A_1 | A_2 \rangle$ pays off \$1 if $A_1 A_2$ occurs and nothing if $\bar{A}_1 A_2$ occurs. If $\bar{A}_2$ occurs, then the bet is called off, and any price paid for the security is returned.

[13] FA essentially requires *both* MP and *external Bayesianity* (EB) [18] to hold. Genest [19], and indirectly Dalkey [8, 9], prove that MP, EB, unanimity, and nondictatorship are mutually incompatible.

agents share the total amount bet in proportion to their bets on $A$. Agents bet to maximize expected payoffs, where expectation is with respect to their probability distribution over the events, subject to a budget constraint limiting their total bets. Eisenberg and Gale show that this mechanism yields a unique set of equilibrium probabilities, and Norvig [52] presents a dynamic process for reaching this equilibrium through iterated bids.

As the authors point out, however, this scheme can yield pathological (their word) results. For example, if there are two bettors with equal budgets, then whichever has more uniform probabilities will dictate the results. According to Genest and Zidek [23], the pari-mutuel approach to belief aggregation "has never enjoyed much popularity" for this reason. The pathological behavior, we believe, can be attributed to the role of arbitrary budgets. In the approach developed here, we impose no budgets, but rather rely on risk aversion to limit bets to the finite range.

Plott et al. [55] investigate whether pari-mutuel markets are able to aggregate information, as postulated by RE theory, in a laboratory setting. In one set of experiments, each subject was given inside knowledge that a subset of horses would definitely *not* win. Although all agents were uncertain as to the outcome, their *collective* information was enough to identify the winning horse with certainty. Information aggregation did occur, and RE-based predictions fit the data well. In a second set of experiments, subjects were given probabilistic evidence in the form of a sample of random draws from the true distribution. Some weak aggregation did occur, but not to the extent predicted by the RE model; instead, a prior information model proved more accurate.

In earlier work, Plott and Sunder [56, 57] conducted laboratory experiments to test the reasonableness of the RE assumption in the context of a securities market. Agents were initially unaware which of three states would occur. In one study [56], privileged insiders were given categorical knowledge of the underlying state. The RE model's predictions—that equilibrium prices and other economic variables converge as if everyone were aware of the true state—were significantly more accurate than those of other models, including the classical Walrasian or PI hypothesis that agents do not revise their beliefs based on prices. When insiders were given less than certain information, the results were not definitive. In a second study [57], insiders were told only that one of the three states would *not* occur. The combined knowledge of all agents was sufficient to logically infer the true state, though no single insider could directly do so. It was found that, in a complete market of three securities, the RE predictions were again the most accurate. In a single security market with agents of different preference types, the RE equilibrium was *not* realized. On the other hand, even in this last condition, Forsythe and Lundholm [15], with a similar experimental design, found that RE *was* verified, as long as subjects were experienced and preferences were common knowledge.

Beyond the controlled setting of the laboratory, researchers have investigated the aggregative properties of public markets. Hanson [27, 26] proposes an *Idea Futures* market, where participants trade in securities that pay off contingent on future developments in science, technology, or other arenas of public interest. He argues that the reward structure of such a market encourages honest revelation of opinions among scientists, and provides more accurate probability assessments for use by funding agencies, public policy leaders, the media, and other interested parties. The concept is operational as a game (i.e., played with fake money), called the "Foresight Exchange", on the world-wide web at `http://www.ideosphere.com/`.

In a similar vein, though run with real money, the Iowa Electronic Market (IEM) (`http://www.biz.uiowa.edu/iem/`) supports trading in securities tied to the outcome of political and financial events. Their 1988 market, open only to University of Iowa students and employees, offered securities that paid off proportionally to the number of votes received by various candidates in that year's US Presidential election. The final prices predicted Bush's final margin of victory more closely than any of the six major polls [16]. Since opening to the public, subsequent US Presidential election markets have attracted wide participation and following.

West [68] presents one axiomatic justification for the (unnormalized) LogOP, based on the betting behavior of a group. He defines $S_i$ as the value for which agent $i$ is indifferent between $D_i$ units of security $\langle A \rangle$, and $S_i$ dollars for sure. Consider an assumption that the group as a whole will be indifferent between the allocation of $D = \langle D_1, \ldots, D_N \rangle$ units of the security, and the certain payoff $S = \langle S_1, \ldots, S_N \rangle$. From the premise that this is true for any $D$ (along with some other conditions), West proves that the group's $\Pr_0(A)$ must obey an *unnormalized* geometric mean or unnormalized LogOP. In contrast to the normalized version (3), this pooling function does not in general yield a probability (e.g., if any agents disagree, $\Pr(A) + \Pr(\bar{A}) < 1$). West goes on to show that the degree by which the aggregate diverges from a probability can be used as a measure of disagreement among the agents. However, we would make the case that the group should not be indifferent between $S$ dollars and security distribution $D$, unless the individual beliefs happen to be the same. In fact, the group should prefer the latter. The reason is that if the group has the security distribution, then there will be trading opportunities that can make everyone better off. We suspect that this preference can account for the normalization factor in the pooling procedure, a consequence of the market model presented here.

Varian [65] analyzes the effects of divergence of opinion among agents on asset valuation and trading volume in a securities market, under various preference models.

# 8 Conclusion

We have outlined a methodology for pooling opinions based on individually rational agents (utility maximizers) trading in a competitive securities market, where the resultant equilibrium prices represent consensus beliefs. Our framework is motivated from the bottom up: we begin with several common assumptions about agents (mainly that each has a subjective probability distribution, has risk averse utility for money, and behaves competitively), and seek to analyze properties of the implied price equilibrium. We establish a few properties for general markets, and provide closed-form characterizations of demand and equilibrium prices for the two cases of negative exponential and generalized logarithmic utility for money. The expert weights in the standard pools coincide with normalized measures of risk tolerance in our model, providing a decision-theoretic interpretation for an often ungrounded concept.

We further explored the possibility that agents update their beliefs based on price feedback. In a simple model of such behavior, where agents can express their own as well as the market's information in terms of equivalent samples, their posterior probabilities are a weighted average of their prior beliefs and the observed market price. Given our standard utility models, the resulting equilibrium prices retain the same functional form, but the expert weights depend on the agents' own assessments of their relative qualities of information.

We investigated a complementary situation, where the agents do not learn, but the market as a whole adapts over time. Price fluctuations are due solely to redistribution of wealth from historically inaccurate agents to those enjoying relative success. In a one-security model, under particular assumptions, these price changes can be rationalized as the Bayesian updates of a representative agent witnessing the outcomes of repeated Bernoulli trials. While this extends Rubinstein's notion of a composite agent [61] to the multiperiod case, our examination of how to "expose" a group points out that, even under conditions in which a standard composite agent exists, the group may appear irrational if market configurations can change.

Future plans include continued theoretical investigations concurrent with computer simulations in more complex markets, to test robustness of convergence, properties of aggregate prices, and further learning and adaptation effects. Our computational testbed is the Michigan Internet AuctionBot (`http://auction.eecs.umich.edu/`), which provides fairly general facilities to set up and run distributed computational markets, and provides an interface for both human and computer agents [75].

**Acknowledgments**

Thanks to Didier Dubois, Robin Hanson, Jeffrey Mackie-Mason, Robert Nau, Charles Plott, Steven Pollock, Ross Shachter, James Smith, Mike West, Fredrik Ygge, and the anonymous reviewers for providing helpful insights, feedback, encouragement, and relevant pointers.

# References

[1] Kenneth J. Arrow. The role of securities in the optimal allocation of risk-bearing. *Review of Economic Studies*, 31(2):91–96, 1964. Translated from *E*conométrie 1953.

[2] Jon Atli Benediktsson and Philip H. Swain. Consensus theoretic classification methods. *IEEE Transactions on Systems, Man, and Cybernetics*, 22(4):688–704, 1992.

[3] James O. Berger. *Statistical Decision Theory and Bayesian Analysis.* Springer-Verlag, New York, second edition, 1985.

[4] Robert F. Bordley. A multiplicative formula for aggregating probability assessments. *Management Science*, 28(10):1137–1148, 1982.

[5] Robert F. Bordley and R. W. Wolfe. On the aggregation of individual probability estimates. *Management Science*, 27(8):959–964, 1981.

[6] Robert T. Clemen and Robert L. Winkler. Aggregation of point estimates: A flexible modeling approach. *Management Science*, 39(4):501–515, 1993.

[7] Roger M. Cooke. *Experts in Uncertainty: Opinion and Subjective Probability in Science*. Environmental Ethics and Science Policy Series. Oxford University Press, New York, 1991.

[8] Norman C. Dalkey. An impossibility theorem for group probability functions. P-4862, The Rand Corporation, Santa Monica, CA, 1972.

[9] Norman C. Dalkey. Toward a theory of group estimation. In H. A. Linstone and M. Turoff, editors, *The Delphi Method: Techniques and Applications*, pages 236–261. Addison-Wesley, Reading, MA, 1975.

[10] Bruno de Finetti. *Theory of Probability: A Critical Introductory Treatment*, volume 1. J. Wiley, New York, 1974.

[11] Morris H. Degroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.

[12] Morris H. Degroot and Julia Mortera. Optimal linear opinion pools. *Management Science*, 37(5):546–558, 1991.

[13] Jacques H. Drèze. Market allocation under uncertainty. In Jacques H. Drèze, editor, *Essays on Economic Decisions under Uncertainty*, pages 119–143. Cambridge University Press, Cambridge, 1987.

[14] Edmund Eisenberg and David Gale. Consensus of subjective probabilities: The pari-mutuel method. *Annals of Mathematical Statistics*, 30:165–168, 1959.

[15] Robert Forsythe and Russell Lundholm. Information aggregation in an experimental market. *Econometrica*, 58(2):309–347, 1990.

[16] Robert Forsythe, Forrest Nelson, George R. Neumann, and Jack Wright. Anatomy of an experimental political stock market. *American Economic Review*, 82(5):1142–1161, 1992.

[17] Simon French. Group consensus probability distributions: A critical survey. *Bayesian Statistics*, 2:183–202, 1985.

[18] Christian Genest. A characterization theorem for externally Bayesian groups. *Annals of Statistics*, 12(3):1100–1105, 1984.

[19] Christian Genest. A conflict between two axioms for combining subjective distributions. *Journal of the Royal Statistical Society*, 46(3):403–405, 1984.

[20] Christian Genest. Pooling operators with the marginalization property. *Canadian Journal of Statistics*, 12(2):153–163, 1984.

[21] Christian Genest, Kevin J. McConway, and Mark J. Schervish. Characterization of externally Bayesian pooling operators. *Annals of Statistics*, 14(2):487–501, 1986.

[22] Christian Genest and Carl G. Wagner. Further evidence against independence preservation in expert judgement synthesis. *Aequationes Mathematicae*, 32(1):74–86, 1987.

[23] Christian Genest and James V. Zidek. Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*, 1(1):114–148, 1986.

[24] Sanford J. Grossman. An introduction to the theory of rational expectations under asymmetric information. *Review of Economic Studies*, 48(4):541–59, 1981.

[25] Robin Hanson. Consensus by identifying extremists. *Theory and Decision*, 44(3):293–301, 1998.

[26] Robin Hanson. Decision markets. *IEEE Intelligent Systems*, 14(3):16–19, 1998.

[27] Robin D. Hanson. Could gambling save science? Encouraging an honest consensus. *Social Epistemology*, 9(1):3–33, 1995.

[28] Chi-Fu Huang and Robert H. Litzenberger. *Foundations for Financial Economics*. Elsevier Science, Amsterdam, 1988.

[29] Aanund Hylland and Richard Zeckhauser. The impossibility of Bayesian group decision making with separate aggregation of beliefs and values. *Econometrica*, 47(6):1321–1336, 1979.

[30] Robert A. Jacobs. Methods for combining experts' probability assessments. *Neural Computation*, 7(5):867–888, September 1995.

[31] Joseph B. Kadane and Robert L. Winkler. Separating probability elicitation from utilities. *Journal of the American Statistical Association*, 83(402):357–363, 1988.

[32] Eero Kasanen and Lenos Trigeorgis. Market utility approach to investment valuation. *European Journal of Operational Research*, 74(2):294–309, 1994.

[33] Ralph L. Keeney and Howard Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. John Wiley and Sons, New York, 1976.

[34] Keith Lehrer and Carl G. Wagner. Probability amalgamation and the independence issue: A reply to Laddaga. *Synthese*, 55(3):339–346, 1983.

[35] William B. Levy and Hakan Delic. Maximum entropy aggregation of individual opinions. *IEEE Transactions on Systems, Man, and Cybernetics*, 24(4):606–613, 1994.

[36] Dennis V. Lindley. Reconciliation of discrete probability distributions. In *Bayesian Statistics 2*, pages 375–390, Amsterdam, 1985. North-Holland.

[37] Dennis V. Lindley. The use of probability statements. In *Accelerated Life Testing and Experts' Opinions in Reliability*, pages 25–57, Amsterdam, 1988. Elsevier.

[38] Robert E. Lucas. Expectations and the neutrality of money. *Journal of Economic Theory*, 4(2):103–24, 1972.

[39] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.

[40] Izhar Matzkevich and Bruce Abramson. Some complexity considerations in the combination of belief networks. In *Ninth Conference on Uncertainty in Artificial Intelligence*, pages 152–158, 1993.

[41] Richard D. McKelvey and Talbot Page. Common knowledge, consensus, and aggregate information. *Econometrica*, 54(1):109–127, 1986.

[42] Richard D. McKelvey and Talbot Page. Public and private information: An experimental study of information pooling. *Econometrica*, 58(6):1321–1339, 1990.

[43] Peter A. Morris. Decision analysis expert use. *Management Science*, 20(9):1233–1241, 1974.

[44] Peter A. Morris. Combining expert judgments: A Bayesian approach. *Management Science*, 23(7):679–693, March 1977.

[45] Peter A. Morris. An axiomatic approach to expert resolution. *Management Science*, 29(1):24–32, 1983.

[46] In Jae Myung, Sridhar Ramamoorti, and Andrew D. Bailey. Maximum entropy aggregation of expert predictions. *Management Science*, 42(10):1420–1436, 1996.

[47] Robert F. Nau. Coherent decision analysis with inseparable probabilities and utilities. *Journal of Risk and Uncertainty*, 10(1):71–91, 1995.

[48] Robert F. Nau and Kevin F. McCardle. Arbitrage, rationality, and equilibrium. *Theory and Decision*, 31:199–240, 1991.

[49] Keung-Chi Ng and Bruce Abramson. Consensus diagnosis—A simulation study. *IEEE Transactions on Systems, Man, and Cybernetics*, 22(5):916–928, 1992.

[50] Keung-Chi Ng and Bruce Abramson. Probabilistic multi-knowledge-base systems. *Applied Intelligence*, 4(2):219–236, 1994.

[51] Lars Tyge Nielsen, Adam Brandenburger, John Geanakoplos, Richard McKelvey, and Talbot Page. Common knowledge of an aggregate of expectations. *Econometrica*, 58(5):1235–1239, 1990.

[52] T. Norvig. Consensus of subjective probabilities: A convergence theorem. *Annals of Mathematical Statistics*, 38:221–225, 1967.

[53] David M. Pennock and Michael P. Wellman. Representing aggregate belief through the competitive equilibrium of a securities market. In *Thirteenth Conference on Uncertainty in Artificial Intelligence*, pages 392–400, Providence, RI, USA, July 1997.

[54] David M. Pennock and Michael P. Wellman. Graphical representations of consensus belief. In *Fifteenth Conference on Uncertainty in Artificial Intelligence (UAI–99)*, July 1999.

[55] C. R. Plott, J. Wit, and W. C. Yang. Parimutuel betting markets as information aggregation devices: experimental results. Technical Report Social Science Working Paper 986, California Institute of Technology, April 1997.

[56] Charles R. Plott and Shyam Sunder. Efficiency of experimental security markets with insider information: An application of rational-expectations models. *Journal of Political Economy*, 90(4):663–98, 1982.

[57] Charles R. Plott and Shyam Sunder. Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica*, 56(5):1085–1118, 1988.

[58] John W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32:122–136, 1964.

[59] Howard Raiffa. *Decision Analysis: Introductory Lectures on Choices under Uncertainty.* Addison-Wesley, Reading, MA, 1968.

[60] Howard Raiffa and R. Schlaifer. *Applied Statistical Decision Theory.* Harvard University, Boston, 1961.

[61] Mark Rubinstein. An aggregation theorem for securities markets. *Journal of Financial Economics*, 1(3):225–244, 1974.

[62] Mark Rubinstein. Securities market efficiency in an Arrow-Debreu economy. *Americian Economic Review*, 65(5):812–824, 1975.

[63] Mark Rubinstein. The strong case for the generalized logarithmic utility model as the premier model of financial markets. *Journal of Finance*, 31(2):551–571, 1976.

[64] Hal R. Varian. The arbitrage principle in financial economics. *Journal of Economic Perspectives*, 1(2):55–72, Fall 1987.

[65] Hal R. Varian. Differences of opinion in financial markets. In Courtenay C. Stone, editor, *Financial Risk: Theory, Evidence and Implications*, pages 3–37. Kluwer Academic, Norwell, MA, 1989.

[66] Carl Wagner. Aggregating subjective probabilities: Some limitative theorems. *Notre Dame Journal of Formal Logic*, 25(3):233–240, 1984.

[67] S. Weerahandi and J. V. Zidek. Multi-Bayesian statistical decision theory. *Journal of the Royal Statistical Society. Series A. General*, 144(1):85–93, 1981.

[68] Mike West. Bayesian aggregation. *Journal of the Royal Statistical Society. Series A. General*, 147(4):600–607, 1984.

[69] Mike West and Jo Crosse. Modeling agent forecast distributions. *Journal of the Royal Statistical Society. Series B. Methodological*, 54(2):553–567, 1992.

[70] Mike West and Jo Crosse. Modeling probabilistic agent opinion. *Journal of the Royal Statistical Society. Series B. Methodological*, 54(1):285–299, 1992.

[71] Robert Wilson. The theory of syndicates. *Econometrica*, 36(1):119–132, 1968.

[72] Robert L. Winkler. The consensus of subjective probability distributions. *Management Science*, 15(2):B61–B75, 1968.

[73] Robert L. Winkler. Combining probability distributions from dependent sources. *Management Science*, 27(4):479–488, 1981.

[74] Robert L. Winkler. Expert resolution. *Management Science*, 32(3):298–303, 1986. Introduces commentary by Lindley, Schervish, Clemen, French, and Morris.

[75] Peter R. Wurman, Michael P. Wellman, and William E. Walsh. The Michigan Internet AuctionBot: A configurable auction server for human and software agents. In *Second International Conference on Autonomous Agents*, pages 301–308, Minneapolis, May 1998.

[76] James V. Zidek. Group decision analysis and its application to combining opinions. *Journal of Statistical Planning and Inference*, 20(3):307–325, 1988.