Technical Report.

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

March 5, 1998

Abstract

Consider a group of Bayesians, each with a subjective probability distribution over a set of uncertain events. An *opinion pool* derives a single consensus distribution over the events, representative of the group as a whole. Several pooling functions have been proposed, each sensible under particular assumptions or measures. Many researchers over many years have failed to form a consensus on which method is best. We propose a *market-based* pooling procedure, and analyze its properties. Participants bet on securities, each paying off contingent on an uncertain event, so as to maximize their own expected utilities. The consensus probability of each event is defined as the corresponding security's equilibrium price. The market framework provides explicit monetary incentives for participation and honesty, and allows agents to maintain individual rationality and limited privacy. "No arbitrage" arguments ensure that the equilibrium prices form legal probabilities. We show that, when events are disjoint and all participants have exponential utility for money, the market derives the same result as the logarithmic opinion pool; similarly, logarithmic utility for money yields the linear opinion pool. In both cases, we prove that the group's behavior is, to an outside observer, indistinguishable from that of a rational individual, whose beliefs equal the equilibrium prices.

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 the group. Forming consensus beliefs is a classic problem in statistics [Degroot, 1974, French, 1985, Genest and Zidek, 1986, Lindley, 1985, West, 1984, West and Crosse, 1992a, West and Crosse, 1992b], and has found application in many areas, from group decision making [Morris, 1974, Weerahandi and Zidek, 1981, Zidek, 1988] to artificial intelligence [Jacobs, 1995, Ng and Abramson, 1994, Pennock and Wellman, 1996, Pennock and Wellman, 1997]. In some group situations there is a natural leader ultimately responsible for decisions, and standard techniques may apply. 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 important results reported in this journal [Bordley and Wolfe, 1981, Bordley, 1982, Clemen and Winkler, 1993, Degroot and Mortera, 1991, Morris, 1977, Myung *et al.*, 1996, Winkler, 1968, Winkler, 1981], definitive answers have been elusive, and various impossibility results [Hylland and Zeckhauser, 1979, Genest and Zidek, 1986, Saari, 1995] and definitional controversies [French, 1985] 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 [de Finetti, 1974]. Imagine that a bettor ("he") must decide at what price to buy and sell each of two lottery tickets: one pays \$1 if A occurs, and nothing otherwise; the other pays \$1 if B. If A and B are disjoint events, then the bettor's total price for both tickets should never exceed \$1—otherwise, he is subject to a guaranteed loss: a shrewd opponent ("she") will sell both tickets to him for more than \$1, paying out at most \$1 in all outcomes. More generally, if the bettor's internal prices for a set of such lottery tickets are not legal probabilities, then he is exposed to sure losses—also called *Dutch books* or *arbitrage* [Varian, 1987]. Now imagine a *group* of bettors locked in a room. The shrewd opponent is outside, and slides offers to buy and sell lottery tickets under the door. She does not care who accepts her offers: if opportunities for risk-free profit exist, she will take advantage of them. From the opponent's perspective, such opportunities imply that the group or "room" is irrational. To be concrete, imagine that bettor #1 will sell "\$1 if A" at \$0.5, and bettor #2 will buy the same ticket at \$0.7. The opponent could potentially buy a ticket from #1 and sell it to #2 for a riskless profit of \$0.2. But if the bettors had first exhausted all opportunities for within-group trades, then bettor #1 would sell instead to bettor #2. Each bettor's transaction is exactly the same, and the group is guaranteed to be 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 the group accept a countering outside offer. If the group is in an *equilibrium* state—that is, all possible within-group trades have been completed—then arbitrage opportunities cannot exist [Nau and McCardle, 1991]. From the outsider's perspective, the group or "room" appears rational. In other words, by de Finetti's same argument, it 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, [1] also called *securities* or *contingent claims*. Each participant still maintains his own rationality; that is, he buys and sells securities according to his own beliefs and utility for money. The equilibrium prices in the market are interpreted as the group's consensus probability distribution. In this paper, we formalize the use of a securities market as a vehicle for pooling opinions, analyze some of its properties, and compare it to other proposed methods. In Section 2 we review some previous proposals for belief aggregation, including the two most common pooling functions: the linear opinion pool (LinOP) and the logarithmic opinion pool (LogOP). In Section 3, we introduce the setup and notation for the market and derive several general characteristics of the system. In Section 4, we show that, for markets in disjoint securities, if all agents have exponential utility (constant risk aversion), then equilibrium prices are exactly the probabilities derived from the LogOP. In Section 5, we demonstrate the same equivalence between logarithmic utility and the LinOP. The "expert weights", notoriously hard to define [French, 1985], are simple parameters of the agents'

[1] Perhaps because they were locked in a room with nothing better to do?

utility functions. In both cases we prove that, to an outside observer, the group's betting behavior is indistinguishable from that of a rational, utility-maximizing individual. In Section 6 we consider the case when participants update their beliefs based on market prices. In Section 7, we place our work within the perspective of other belief aggregation and opinion pool research. We conclude in Section 8.

2 Opinion Pools

An *opinion pool* combines the beliefs of a group into a single, consensus belief. We restrict attention to the case of probabilistic beliefs. Genest and Zidek [1986] and French [1985] provide excellent literature reviews. Three approaches are generally distinguished. The first assumes a single Bayesian decision maker (real or fictitious, within or outside the group), called the "supra Bayesian", with priors over the events in question, and over the participants' beliefs. The supra Bayesian updates his beliefs about the events based on the "evidence" of members' beliefs via Bayes's rule. The resulting posterior is the consensus belief. The second approach is to apply a 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). Consider N participants with beliefs $\Pr_i(A), i = 1 \ldots N$, for an event A, and denote the combined belief $\Pr_0(A)$. The LinOP is a weighted arithmetic mean of the members' probabilities,

$$\Pr_0(A) = \sum_{i=1}^{N} w_i \Pr_i(A), \tag{1}$$

and the LogOP is a weighted geometric mean,

$$\Pr_0(A) = \kappa \prod_{i=1}^{N} [\Pr_i(A)]^{w_i}, \tag{2}$$

where κ is a normalization constant,[2] and the w_i are "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 [Cooke, 1991]. A third, relative 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 between experts [Levy and Delic, 1994, Myung *et al.*, 1996].

The debate over which aggregation method is best continues to rage [Benediktsson and Swain, 1992, Cooke, 1991, French, 1985, Genest and Zidek, 1986]. Several authors (most emphatically Lindley [1985]) argue that, when possible, the supra Bayesian approach is best, as it is grounded in standard normative Bayesian theory [Clemen and Winkler, 1993, Morris, 1977, Winkler, 1981]. However, several conceptual questions arise in characterizing the *single* supra Bayesian responsible for the consensus of a (possibly leaderless) group [Genest and Zidek, 1986]. If the supra Bayesian is real, is he a member of the group or an outsider? If he is an outsider, will the group accept his "authority" to generate the consensus? If he is a member, how is he chosen? If he is fictitious, where do his 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' beliefs are normally distributed and independent—in order to yield tractable results [Clemen and Winkler, 1993, Jacobs, 1995, Lindley, 1985, Winkler, 1981]. 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.[3] Each of many opinion pools, including the LinOP and LogOP, can be shown to be the *only* form possible, if one assumes an appropriate, corresponding set of axioms [Bordley, 1982, Morris, 1983, Genest and Zidek, 1986, West, 1984].

[2] In some treatments, the LogOP is not normalized [West, 1984], in which case the consensus is generally not a legal probability distribution.

[3] For example, Cooke [1991] calls any consensus that does not have what is called the *marginalization property* "downright queer"; on the other hand, Lindley [1985] regards this property as an "adhockery".

It follows that no pool can satisfy a sufficiently large set of potentially desirable properties; these arguments are the so-called "impossibility" results. 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 [Benediktsson and Swain, 1992, French, 1985, Winkler, 1968]. Some more formal methods to derive weights have been proposed, by making assumptions concerning the form of, or interdependence between, participants' beliefs [Cooke, 1991, Degroot and Mortera, 1991, Jacobs, 1995, Morris, 1977], or through iterative self-weighting procedures [Degroot, 1974]. 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.

3 A Securities Market for Pooling Opinion

In this paper, we analyze the use of a securities market as an opinion pooling device. In our context, one unit of a security is essentially a lottery ticket that pays off one dollar contingent on uncertain event, and nothing otherwise. Participants bid on securities according to their own beliefs and utility for money. The market's equilibrium price for each security is the consensus probability for the corresponding event.

Section 3.1 introduces the overall market framework, defines relevant notation, catalogs any assumptions made, and addresses the advantages of (and motivations for) a market approach. Section 3.2 derives several formal properties of the mechanism in general. Subsequent sections will examine two specific cases in depth, where the market yields the same consensus as the LogOP and LinOP.

3.1 Market Setup and Notation

We assume a finite, discrete-probability event space, denoted $\Omega = \{\omega_1, \omega_2, \ldots\}$, where the ω are disjoint, exhaustive events. 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 Ω. Agent i has a subjective joint probability distribution $\Pr_i(\cdot)$ over Ω. His probability for an event A_j is then just $\Pr_i(A_j) = \sum_{\omega \in A_j} \Pr_i(\omega)$. 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).

The market consists of securities "\$1 if A" for every event of interest A. One unit of the security pays off one dollar if the event occurs, and nothing otherwise. When interpretation is clear from context, the security is denoted simply by its corresponding event A. A quantity x of security A is a lottery resulting in $\$x$ if and only if A occurs. If the agent must pay a price p per unit of the security, then investing in x units is equivalent to accepting a lottery with payoff $(1-p)x$ if A occurs, and $-px$ otherwise; this lottery is denoted $L(x) = [\Pr(A), (1-p)x; \Pr(\bar{A}), -px]$. Positive x indicates a quantity to buy, and negative x a quantity to sell; in this framework, *either* action incurs risk. Note that buying A at price p yields exactly the same lottery as selling $\bar{A}$ at price $1-p$.

Throughout this paper, we assume that each participant behaves *competitively*, that is, maximizes expected utility at a given set of prices, or equivalently, ignores any effect his demand may have on prices. We also assume that utility for money does not depend on the uncertain state, and is monotonically increasing. To avoid unbounded demand, we usually make the further assumption that each agent is *risk averse* [Keeney and Raiffa, 1976, Pratt, 1964], that is, always prefers the expected payoff of a lottery to the lottery itself, or equivalently, has concave u. When considering risk averse agents, we also assume that no security can be replicated as a linear combination of other securities. Otherwise an agent could buy one security and sell the equivalent portfolio without risk, and demand may be nonunique and unbounded. Note that this assumption implies that each event is strictly uncertain ($0 < \Pr(A) < 1$). Most of the discussion considers only fixed beliefs, though none of the results depend on it, and in Section 6 we directly consider the case where beliefs are updated based on price information.

Given the assumptions of competitive behavior and state-independent utility, we can define the agent's utility, U, for purchasing a quantity x of the security as his expected u,

$$U(x) = E[u(L(x))] = \Pr(A)u((1-p)x) + \Pr(\bar{A})u(-px).$$

Now consider the case of multiple securities, $A_1, A_2, \ldots, A_S$. Let $x^{(j)}$ represent a quantity of security A_j, and $p^{(j)}$ the security's associated price. The utility of a bundle $\mathbf{x} = \langle x^{(1)}, x^{(2)}, \ldots, x^{(S)} \rangle$ of securities can also be expressed as an expected utility for money,

$$U(\mathbf{x}) = \sum_{\omega \in \Omega} \Pr(\omega) u(Y_\omega), \tag{3}$$

where

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

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

Denote the price vector for all securities as $\mathbf{p} = \langle p^{(1)}, p^{(2)}, \ldots, p^{(S)} \rangle$ and let the agent's *demand function*, $\mathbf{x}(\mathbf{p}) = \langle x^{(1)}(\mathbf{p}), x^{(2)}(\mathbf{p}), \ldots, x^{(S)}(\mathbf{p}) \rangle$, represent the quantities of securities that maximizes utility (3) at those prices.

For risk averse agents, concavity of u entails concavity of U, and hence any critical point of (3) is a global maximum. Furthermore, for a *sufficiently* risk averse agent (one sufficient condition is given in Proposition 2 below), demand is bounded if prices form coherent probabilities; that is, prices do not admit arbitrage.

To this point we have addressed the optimization problem of a single agent faced with the decision to invest in securities. Now consider a market with N participants. Each agent maximizes his own utility (3), trading with other agents until all buy orders (positive $x^{(j)}$) match with sell orders (negative $x^{(j)}$) for each security j. A market is in *competitive equilibrium* at prices $\mathbf{p}$ if and only if

$$\sum_{i=1}^{N} \mathbf{x}_i(\mathbf{p}) = \mathbf{0}. \tag{4}$$

The competitive equilibrium prices in the economy represent the group's consensus probabilities. A market in equilibrium cannot admit arbitrage [Nau and McCardle, 1991] (intuitively, if opportunities for risk-free profit still exist, then rational agents must not be maximizing utility). It follows, from de Finetti's argument, that the equilibrium prices in a securities market must form a legal probability distribution.

There *does* exist a definable consensus at equilibrium: all agents' *risk-neutral* probabilities agree, and equal the equilibrium prices [Nau and McCardle, 1991]. An agent's risk-neutral probabilities are defined as $\widehat{\Pr}(\omega) \equiv \kappa \Pr(\omega) u'(Y_\omega)$, where κ is a normalization constant [Nau, 1995]. A hypothetical agent with transformed probabilities $\widehat{\Pr}(\omega)$ and transformed (state-dependent) *risk-neutral* utility $\hat{u}(y) \equiv u(y)/u'(Y_\omega)$ has equivalent utility U for securities (3), and thus the same optimizing behavior, as an agent with $\Pr(\omega)$ and $u(y)$.[4] Since a risk neutral agent has unbounded demand at any prices that do not equal

[4] U will differ only by the multiplicative factor κ.

his probabilities, the market cannot be in equilibrium unless all agents' risk-neutral probabilities equal the current prices. [Nau and McCardle, 1991].

Most opinion pools require that agents directly, and truthfully, reveal their entire subjective probabilities to a central location. In contrast, agents in a securities market have monetary incentives to participate, to make decisions honestly, and to gather all cost-effective evidence, revealing information only implicitly through their trades. In fact, if measures are taken to ensure the privacy of bids, the *only* information directly revealed about a participant is his risk-neutral probability distribution at equilibrium, which equals the consensus anyway. Since supply and demand are balanced in equilibrium, the mechanism requires no subsidized payments to induce information revelation, and so the only cost involved is that of organizing and running the market.

3.2 General Properties of the Market Mechanism

There are two categories of properties to consider in this market framework: those that deal with a single agent's demand for securities, and those that examine the equilibrium prices arising from group dynamics. We begin with several propositions concerning the former.

Proposition 1 (Incentive to participate) *Any agent with* $\Pr(A_j) \neq p^{(j)}$ *for some* j *has nonzero demand.*

Proof. Without loss of generality, assume $\Pr(A_1) \neq p^{(1)}$. We show that the agent prefers $\mathbf{x} = \langle \epsilon, 0, \ldots, 0 \rangle$ to $\mathbf{x} = \mathbf{0}$, for some nonzero ϵ. For small enough ϵ, utility $U(\epsilon, 0, \ldots, 0)$ (3) can be approximated around $\mathbf{0}$ by a first-order Taylor expansion,

$$\begin{aligned} U(\epsilon, 0, \ldots, 0) &\approx U(\mathbf{0}) + \epsilon \left. \frac{\partial U}{\partial x^{(1)}} \right|_{\mathbf{0}} \\ &= U(\mathbf{0}) + \epsilon \sum_{\omega \in \Omega} \Pr(\omega)(1_{\omega \in A_1} - p^{(1)}) u'(0) \\ &= U(\mathbf{0}) + \epsilon u'(0)(\Pr(A_1) - p^{(1)}). \end{aligned}$$

Utility for money is monotonically increasing, so $u'(0) > 0$. If $\Pr(A_1) > p^{(1)}$, then $U(\epsilon, 0, \ldots, 0) > U(\mathbf{0})$ for positive ϵ. Similarly, if $\Pr(A_1) < p^{(1)}$, then the agent prefers $\epsilon < 0$ to $\epsilon = 0$. □

Note that Proposition 1 does *not* prove that optimal demand includes positive or negative $x^{(j)}$, only that $\mathbf{x} = \mathbf{0}$ is strictly dominated.

Proposition 2 (Bounded demand) *Any agent with utility for money such that*

$$\lim_{y \to \infty} \frac{u(y)}{u(-y)} = 0$$

has bounded demand if prices form coherent probabilities.

Proof. Consider a nonzero, finite portfolio of securities $\delta = \langle \delta^{(1)}, \delta^{(2)}, \ldots, \delta^{(S)} \rangle$. Let $\Delta_\omega \equiv \sum_{j=1}^{S} (1_{\omega \in A_j} - p^{(j)})\delta^{(j)}$ be the payoff of this portfolio in outcome ω. Let Ω_+ be the set of outcomes with nonegative payoff ($\Delta_\omega \geq 0$ for all $\omega \in \Omega_+$), and Ω_- the set of outcomes with strictly negative payoff. Prices are coherent (do not admit arbitrage) if and only if, for every portfolio δ, there is at least one negative-payoff outcome $\omega \in \Omega_-$. The agent invests in $z > 0$ "shares" of this portfolio so that $\mathbf{x} = z\delta$. Utility is given by (3):

$$\begin{aligned} U(z\delta) &= \sum_{\omega \in \Omega} \Pr(\omega) u(z\Delta_\omega) \\ &= \sum_{\omega \in \Omega_+} \Pr(\omega) u(z\Delta_\omega) + \sum_{\omega \in \Omega_-} \Pr(\omega) u(z\Delta_\omega) \end{aligned}$$

The condition on u implies that $\lim_{z \to \infty} u(z\Delta_{\omega_+})/u(z\Delta_{\omega_-}) = 0$ for any $\omega_+ \in \Omega_+$ and $\omega_- \in \Omega_-$. Thus, in the limit as $z \to \infty$, the negative payoff terms dominate U. Mathematically,

$$\lim_{z \to \infty} U(z\delta) = \sum_{\omega \in \Omega_-} \Pr(\omega) u(z\Delta_\omega),$$

which is decreasing in z. Thus, regardless of the choice for δ, optimal z is bounded. $\Box$

Lemma 3 (No negative expected payoffs) *A risk neutral or risk averse agent chooses demand* $\mathbf{x}$ *to yield nonnegative expected payoff; moreover, if* $\Pr(A_j) \neq p^{(j)}$ *for some* j*, then he chooses demand with strictly positive expected payoff.*

Proof. A risk neutral agent, by definition, maximizes expected payoff. A risk averse agent would prefer zero demand (a lottery with zero certainty equivalent), to any nonzero (risky) demand that yields negative or zero expected payoff (a lottery with negative certainty equivalent). If $\Pr(A_j) \neq p^{(j)}$ for some j, then Proposition 1 implies nonzero demand and thus strictly positive expected payoff. □

Proposition 4 (Qualitative partial demand) *A risk neutral or risk averse agent with* $\Pr(A_j) > (<)\ p^{(j)}$ *for some* j*, and* $\Pr(A_k) = p^{(k)}$ *for all* $k \neq j$ *has positive (negative) demand* $x^{(j)}$ *for security* A_j*. In the boundary case, a risk averse agent with* $\Pr(A_j) = p^{(j)}$ *has optimal demand* $\mathbf{x} = \mathbf{0}$*.*

Proof. Note that, regardless of his holdings of other securities, the agent's expected payoff is $(\Pr(A_j) - p^{(j)})x^{(j)}$. Then, by Lemma 3, any demand $\mathbf{x}$ must include $x^{(j)} > (<)\ 0$ to ensure strictly positive expected payoff. If $\Pr(A_j) = p^{(j)}$, *every* demand yields zero expected payoff; thus any nonzero demand for a risk averse agent is precluded by Lemma 3. □

Propositions 1 and 4, at first glance, may seem to violate the intuition of risk aversion. After all, a risk averse agent always values a lottery at a price strictly less than its expected payoff. Then why will he buy a security at any price arbitrarily smaller than his belief? The key point here is that demand is continuous. The participant may not wish to buy a fixed amount (say, one unit) of the security at a price slightly below his belief, but he will always be willing to buy some (possibly very small) amount of it.

We turn now to analysis of equilibrium prices. The next proposition establishes that the market has one of the properties often considered desirable for opinion pools.

Proposition 5 (Unanimity) *If all agents are risk averse or risk neutral with equal beliefs across the events of interest (*$\Pr_h(A_j) = \Pr_i(A_j)$ *for all* $h, i = 1 \dots N$ *and* $j = 1 \dots S$*), then the unique competitive equilibrium prices equal these beliefs (*$p^{(j)} = \Pr(A_j)$ *for all* j*).*

Proof. At prices $p^{(j)} = \Pr(A_j)$, $\mathbf{x} = \mathbf{0}$ is optimal for risk averse agents (by Proposition 4) and *any* demand, including $\mathbf{x} = \mathbf{0}$, is optimal for risk neutral agents. Thus excess demand can be zero, (4) is satisfied, and these prices form a competitive equilibrium. We prove by contradiction that, at any other prices $p^{(j)} \neq \Pr(A_j)$, the market cannot be in equilibrium. By Lemma 3, each agent i has demand such that $E[L(\mathbf{x}_i)] > 0$. The equilibrium condition (4) and identical beliefs imply that $E[\sum_{i=1}^{N} L(\mathbf{x}_i)] = \sum_{i=1}^{N} E[L(\mathbf{x}_i)] = 0$. But this implies that $E[L(\mathbf{x}_i)] \leq 0$ for some i. □

4 Exponential Utility and the LogOP

In this section, we investigate the consequences stemming from exponential utility for money, $u(y) = -e^{-cy}$. This utility form is synonymous with *constant* risk aversion, where c is the risk aversion coefficient. We find that, when events are disjoint, the competitive equilibrium prices are equivalent to the consensus derived by the logarithmic opinion pool (LogOP), and the expert weights are normalized risk tolerances.[5]

As before, we break up the discussion into two parts: Section 4.1 considers each agent's optimization problem, and Section 4.2 examines the effects of their collective actions on equilibrium prices.

4.1 Individual Demand

To facilitate discussion, first consider the simplest case of a market in only one security, "\$1 if A", priced at p. Each participant must decide what quantity, x, to buy ($x > 0$) or sell ($x < 0$) given his belief, $\Pr(A)$, and utility for money, $u(y) = -e^{-cy}$. The decision to purchase x units is equivalent to the acceptance of a lottery $L(x)$ with payoff $(1-p)x$ if A occurs, and $-px$ otherwise. In this case, his utility (3) becomes

$$U(x) = E[u(L(x))] = -\Pr(A)e^{-c(1-p)x} - \Pr(\bar{A})e^{cpx}.$$

Solving for optimal demand yields a closed form,

$$x(p) = \frac{1}{c}\ln\left(\frac{\Pr(A)}{p} \cdot \frac{(1-p)}{\Pr(\bar{A})}\right). \tag{5}$$

Consistent with Proposition 4, the agent's demand is directly related to his belief in the probability of A, and inversely related to the price of the security. At $p = \Pr(A)$, his demand is zero. As his risk aversion approaches zero (approximating risk neutrality), he is willing to buy or sell increasing amounts of the good, assuming that price does not equal belief.

Now consider the slightly more general case of two arbitrary events A_1 and A_2, and two corresponding securities "\$1 if A_1" and "\$1 if A_2". Following the same progression above, the utility for the two securities is the expected

[5]Risk tolerance is the reciprocal of risk aversion.

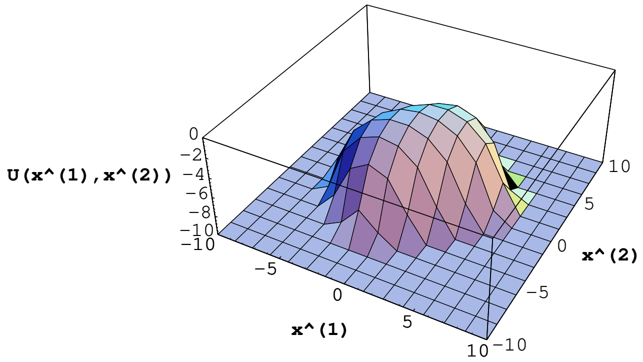

Figure 1: Utility (6) versus $x^{(1)}$ and $x^{(2)}$ for the parameters $\Pr(A_1) = \Pr(A_2) = 0.5$, $\Pr(A_1A_2) = 0.1$, $p^{(1)} = p^{(2)} = 0.5$, and $c = 1$. Utility is maximized at $\{x^{(1)}, x^{(2)}\} = \{0, 0\}$, as expected when price equals belief.

utility of the corresponding lottery, which for constant risk aversion is given by

$$\begin{aligned} U(x^{(1)}, x^{(2)}) = E[u(L(x^{(1)}, x^{(2)}))] = \\ &- \Pr(A_1A_2)e^{-c\left[(1-p^{(1)})x^{(1)}+(1-p^{(2)})x^{(2)}\right]} \\ &- \Pr(A_1\bar{A}_2)e^{-c\left[(1-p^{(1)})x^{(1)}-p^{(2)}x^{(2)}\right]} \\ &- \Pr(\bar{A}_1A_2)e^{-c\left[-p^{(1)}x^{(1)}+(1-p^{(2)})x^{(2)}\right]} \\ &- \Pr(\bar{A}_1\bar{A}_2)e^{-c\left[-p^{(1)}x^{(1)}-p^{(2)}x^{(2)}\right]}. \end{aligned} \tag{6}$$

The decision variables in this optimization problem are coupled; optimal demand for A_1 may depend on the demand for A_2, and thus on the price of A_2 (and vice versa). Figure 1 graphs utility (6) as a function of $x^{(1)}$ and $x^{(2)}$ for a particular instantiation of beliefs, risk coefficient, and prices. By solving the first-order condition, we can again find a closed form for optimal demand:

$$x^{(1)}(p^{(1)}, p^{(2)}) = \frac{1}{c}\ln\left(\frac{Q(1-p^{(1)}-p^{(2)}) + R(p^{(2)}-p^{(1)}) + \sqrt{Z}}{2\Pr(\bar{A}_1A_2)\Pr(\bar{A}_1\bar{A}_2)p^{(1)}}\right), \tag{7}$$

where

$$\begin{aligned} Z &\equiv Q^2(1-p^{(1)}-p^{(2)})^2 + 2QR[p^{(1)}(1-p^{(1)}) + p^{(2)}(1-p^{(2)})] + R^2(p^{(2)}-p^{(1)})^2, \\ Q &\equiv \Pr(A_1\bar{A}_2)\Pr(\bar{A}_1A_2), \text{ and} \\ R &\equiv \Pr(A_1A_2)\Pr(\bar{A}_1\bar{A}_2). \end{aligned}$$

Optimal demand for A_2 is analogous, with security subscripts switched.

Proposition 6 (Separation of independent events) *Given exponential utility for money, if A_1 and A_2 are believed independent, then the optimization problem is separable with respect to $x^{(1)}$ and $x^{(2)}$, and each reduces to the single-security case (5).*

Proof. If $\Pr(A_1 A_2) = \Pr(A_1)\Pr(A_2)$, then (7) becomes

$$\begin{aligned}
& x^{(1)}(p^{(1)}, p^{(2)}) = \\
& \quad \frac{1}{c}\ln\left(\frac{\Pr(A_1)\Pr(\bar{A}_1)\Pr(A_2)\Pr(\bar{A}_2)(1-2p^{(1)})+\sqrt{Z}}{2\Pr(\bar{A}_1)\Pr(\bar{A}_1)\Pr(A_2)\Pr(\bar{A}_2)p^{(1)}}\right) \\
& \quad \text{where } Z \equiv [\Pr(A_1)\Pr(\bar{A}_1)\Pr(A_2)\Pr(\bar{A}_2)]^2 \\
= & \quad \frac{1}{c}\ln\left(\frac{\Pr(A_1)}{p^{(1)}}\cdot\frac{(1-p^{(1)})}{\Pr(\bar{A}_1)}\right),
\end{aligned}$$

and similarly for $x^{(2)}$. □

When the events are dependent, demands for the two goods are correlated in the opposite direction of the dependence. That is, $x^{(1)}$ is increasing in $x^{(2)}$ when $\Pr(A_1|A_2) < \Pr(A_1|\bar{A}_2)$, and decreasing in $x^{(2)}$ when $\Pr(A_1|A_2) > \Pr(A_1|\bar{A}_2)$. This behavior arises because negatively correlated goods provide *insurance* for each other, whereas positively related events increase the exposed risk, and because both securities pay off in a common currency (dollars) for which the agent is risk averse.

For a moment, let us relax the assumption that no security is equivalent to a combination of other securities, in order to examine its implications. Consider two securities representing equivalent events, $A_1 = A_2$. If $p^{(1)} \neq p^{(2)}$, then the agent can achieve unbounded utility through arbitrage, buying infinite amounts of the cheaper security to sell in the higher priced market. If the prices coincide, then the agent demands the same total amount as if there were only one security, splitting his demand arbitrarily between the two available. Put another way, he still desires the same total amount of $A_1 = A_2$, but there also exists a risk-free portfolio (z of A_1 and $-z$ of A_2) for which the agent is indifferent between any z. Figures 2 and 3 graph expected utility (6) versus $x^{(1)}$ and $x^{(2)}$ for this situation. The former illustrates the case when $p^{(1)} = p^{(2)}$, the latter when $p^{(1)} > p^{(2)}$.

The same basic principles hold when $A_1 = \bar{A}_2$. If $p^{(1)} \neq 1 - p^{(2)}$, then an arbitrage opportunity exists. Otherwise, there exists a risk-free investment option (buy equal amounts of both), and optimal demand is not unique.

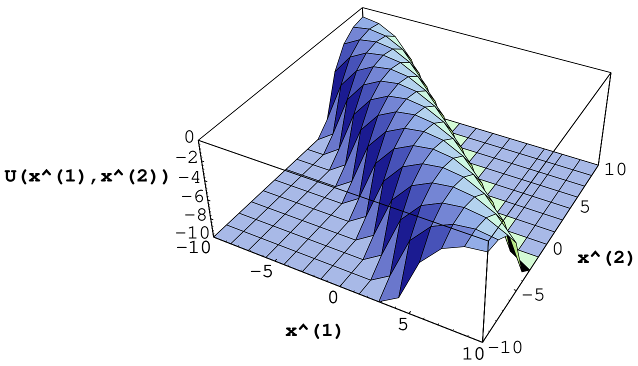

Figure 2: Utility (6) versus $x^{(1)}$ and $x^{(2)}$ for the parameters $\Pr(A_1) = \Pr(A_2) = \Pr(A_1 A_2) = 0.5$ (thus implying that the events are equivalent), $p^{(1)} = p^{(2)} = 0.5$, and $c = 1$. Utility is a maximum along the line $x^{(1)} + x^{(2)} = 0$; the agent's total demand for $A_1 = A_2$ is zero, but he splits purchases arbitrarily between the equivalent markets.

In general, any prices that form incoherent probabilities (that allow arbitrage) cannot be part of an equilibrium. Moreover, if any participant, 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 equivalent to another, then he will effectively dictate that relationship in the consensus, by demanding infinite quantities if prices suggest otherwise.[6]

Formal analysis for more than two arbitrary securities becomes difficult. Although it is a straightforward task to write down each agent's expected utility for money, the number of terms is in general exponential in the number of events. Numeric maximization methods such as gradient descent or Newton's method will be well behaved, since utility for securities is concave, though will still suffer from exponential blowup in the size of the objective function.

Finally, consider a set of S *disjoint* securities, $A_1, A_2, \ldots, A_S$. To ensure that no security is equivalent to a portfolio of other securities, the events are not exhaustive. For notational purposes, we define the event $A_0 \equiv \bar{A}_1 \bar{A}_2 \cdots \bar{A}_S$, though it does not have a corresponding security. (A_0 can be replicated by selling equal amounts of all securities.) Mutually exclusive events are certainly a special case, however any event space can partitioned

[6]In this treatment, we impose no restrictions on demand magnitude, either by budget constraints, or through market rules.

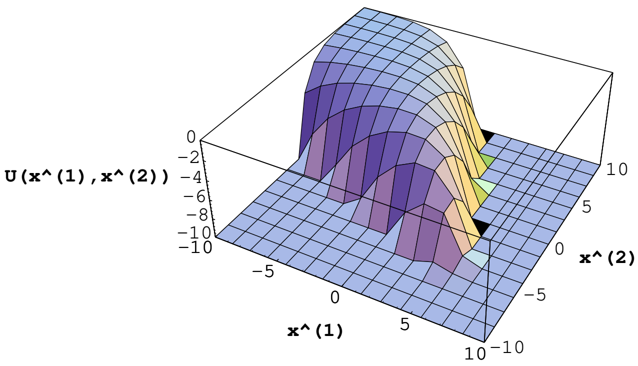

Figure 3: Utility (6) versus $x^{(1)}$ and $x^{(2)}$ for the parameters $\Pr(A_1) = \Pr(A_2) = \Pr(A_1 A_2) = 0.5$ (thus implying that the events are equivalent), $p^{(1)} = 0.7$, $p^{(2)} = 0.3$, and $c = 1$. The prices are inconsistent with the logical relationship $A_1 = A_2$, and the agent maximizes utility by selling A_1 and buying A_2 in infinite quantities.

into securities representing disjoint events.[7] Optimal demand is now:

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

This solution is unique, since demand is bounded (Proposition 2) and utility is concave.

Note that, as expected, (8) reduced to (5) when $S = 1$, and (7) coincides with (8) when A_1 and A_2 are disjoint, and $S = 2$.

4.2 Equilibrium Prices

We next examine how the collective decisions of all participants affect the market's equilibrium prices, which, by construction, form the group's consensus probabilities.

We continue with the scenario of S disjoint securities, generalizing to N agents with constant risk aversion. The condition for competitive equilibrium (4) is $\sum_{i=1}^{N} x_i^{(j)}(\mathbf{p}) = 0$, for all $j = 1 \ldots S$. Each demand $x_i^{(j)}$, of agent i for security j, is given by (8).

[7] This can be achieved by setting $A_0 \leftarrow \bar{B}_1 \cdots \bar{B}_s, A_1 \leftarrow \bar{B}_1 \cdots B_s, \ldots, A_S \leftarrow B_1 \cdots B_s$, where $S = 2^s - 1$ and the Bs are arbitrary events, or by simply creating one security for all but one $\omega \in \Omega$.

Proposition 7 (Exponential utility $\Rightarrow$ LogOP) *In a market consisting of S disjoint securities and N agents, each with constant risk aversion c_i, the prices*

$$p^{(j)} = \frac{\prod_{i=1}^{N}[\Pr_i(A_j)]^{\alpha_i}}{\prod_{i=1}^{N}[\Pr_i(A_0)]^{\alpha_i} + \prod_{j=1}^{N}[\Pr_i(A_1)]^{\alpha_i} + \cdots + \prod_{j=1}^{N}[\Pr_i(A_S)]^{\alpha_i}} \qquad (9)$$

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

Proof. These prices can be shown to satisfy (4) by substitution, and some lengthy algebra. □

These prices are the same probabilities as derived by the LogOP (2), where the expert weights are normalized risk tolerances. Our market model, therefore, provides one way to ground a well-known centralized pooling mechanism in terms of individual behavior.

It also provides a decision-theoretic interpretation for the notoriously slippery concept of expert weights. In the usual interpretation, the exponents in (9) encode some sort of degree of expertise, confidence, or reliability, and are almost always chosen in an ad hoc manner [Benediktsson and Swain, 1992, French, 1985, Genest and Zidek, 1986, Ng and Abramson, 1994, Winkler, 1968]. In arguing against the use of opinion pools, French [1985] presents as his first reason that

> ...they all introduce weights ...which are not operationally defined. How should they be chosen? To say that one expert is twice as good as another is a figure of speech, not an arithmetic statement. The problem is further complicated by the likelihood of correlation between the expert's [sic] opinions. Several pragmatic solutions have been proposed ..., but to my knowledge none avoid a certain arbitrariness.

In our model, the weights are *derived*, and are directly related to a well-understood quantity, namely risk aversion. Note also that the weights sum to unity, as is the standard convention.

When considering any model of group dynamics, it is often desirable to show that the group as a whole behaves rationally. The following proposition establishes such a characteristic in our system, for the case of disjoint events and constant risk aversion.

Proposition 8 (Representative "super-agent", Part I) *Suppose that agents $i = 1 \ldots N$ with beliefs $\Pr_i(A_j)$ for disjoint securities $j = 1 \ldots S$, and utility functions $u_i(y) = e^{-c_i y}$, effect equilibrium prices $q^{(j)}$. Then their aggregate (total) demand for each security is equal to the demand of a single representative "super-agent", with beliefs $\Pr_0(A_j) = q^{(j)}$ and utility $u_0(y) = e^{-c_0 y}$, such that $1/c_0 = \sum_{i=1}^{N} 1/c_i$.*

Proof. For each security, A_j, the total demand of all agents is:

$$
\begin{aligned}
&= \sum_{i=1}^{N} \frac{1}{c_i} \ln \left(\frac{1 - p^{(1)} - \cdots - p^{(S)}}{p^{(j)}} \cdot \frac{\Pr_i(A_j)}{\Pr_i(A_0)} \right) \\
&= \frac{1}{c_0} \ln \left(\frac{1 - p^{(1)} - \cdots - p^{(S)}}{p^{(j)}} \right) + \frac{1}{c_0} \ln \left(\frac{\prod_{i=1}^{N} [\Pr_i(A_j)]^{c_0/c_i}}{\prod_{i=1}^{N} [\Pr_i(A_0)]^{c_0/c_i}} \right) \\
&= \frac{1}{c_0} \ln \left(\frac{1 - p^{(1)} - \cdots - p^{(S)}}{p^{(j)}} \cdot \kappa \prod_{i=1}^{N} [\Pr_i(A_j)]^{\alpha_i} \cdot \frac{1}{\kappa \prod_{i=1}^{N} [\Pr_i(A_0)]^{\alpha_i}} \right) \\
&\quad \text{where } \kappa \equiv \frac{1}{\prod_{i=1}^{N} [\Pr_i(A_0)]^{\alpha_i} + \cdots + \prod_{i=1}^{N} [\Pr_i(A_S)]^{\alpha_i}} \\
&= \frac{1}{c_0} \ln \left(\frac{1 - p^{(1)} - \cdots - p^{(S)}}{p^{(j)}} \cdot \frac{q^{(j)}}{1 - q^{(1)} - \cdots - q^{(S)}} \right) \text{ (from (9))} \\
&= \frac{1}{c_0} \ln \left(\frac{1 - p^{(1)} - \cdots - p^{(S)}}{p^{(j)}} \cdot \frac{\Pr_0(A_j)}{\Pr_0(A_0)} \right)
\end{aligned}
$$

□

To an outside observer, the group behavior is indistinguishable from that of a rational individual, optimizing as delineated in Section 4.1—total demand can be rationalized as maximization of expected utility (with constant risk aversion, at that). In the same way, any subset of the agents can be aggregated according to Proposition 8, interacting with the rest of the system as if it were an individual. 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 super-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.

5 Logarithmic Utility and the LinOP

This section makes a different assumption regarding each participant's utility for money, namely that $u(y) = \ln(y + b)$. This is a standard *decreasingly* risk averse form. Note that the function 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. For this reason, we call b the *maximum loss bound*. The reader is free to interpret b as a "budget" of sorts; however the bound arises directly from the utility function, not from any external constraint, and applies only for money that may possibly be lost, not total amount invested.

We follow a progression similar to the previous section, arriving at the conclusion that the equilibrium prices for disjoint securities equal the probabilities from the linear opinion pool (LinOP).

5.1 Individual Demand

We start again at the simplest case: one agent investing in one security. Utility (3) is

$$U(x) = -\Pr(A)\ln((1-p)x + b) - \Pr(\bar{A})\ln(-px + b).$$

The limit on maximum dollars at risk implies that $-b/(1-p) < x < b/p$, since outside this range $U = -\infty$. These bounds can also be understood as follows: if an agent sells a quantity $-b/(1-p)$ of the security, he stands to lose \$$b$ in the worst case; similarly if he buys b/p units he also could lose \$$b$. Within this range, his optimal demand is,

$$x(p) = b\left(\frac{\Pr(A)}{p} - \frac{\Pr(\bar{A})}{1-p}\right).$$

For more than one security, demands are correlated due to dependencies between events, and also due to wealth-dependent risk aversion. Even for the case of independent events, optimization of demand is *not* separable. Two securities could be priced attractively enough that an agent would be willing to place \$$b/2$ or more at risk on either security, if offered by itself; however, he would never put \$$b$ or more at risk on both together.

For the more general case of S disjoint securities, the unique optimal demand is

$$x^{(j)}(\mathbf{p}) = b\left(\frac{\Pr(A_j)}{p^{(j)}} - \frac{\Pr(A_0)}{1 - p^{(1)} - p^{(2)} - \cdots - p^{(S)}}\right). \tag{10}$$

5.2 Equilibrium Prices

Now we generalize to a market of N agents.

Proposition 9 (Logarithmic utility $\Rightarrow$ LinOP) *In a market consisting of S disjoint securities and N agents, each with utility for money $u_i(y) = \ln(y + b_i)$, the prices*

$$p^{(j)} = \sum_{i=1}^{N} \beta_i \Pr_i(A_j) \tag{11}$$

form a competitive equilibrium, where $\beta_i = b_i / \sum_l b_l$ is agent i's normalized maximum loss bound.

Proof. These prices satisfy (4) by substitution. □

This is exactly the LinOP formula (1), where the expert weights are again directly related to risk tolerance. Thus both the linear and logarithmic opinion pools arise as special cases of our market mechanism; choosing still other utility forms seems a natural way to catalog a variety of aggregation functions.

The market mechanism diverges from the LinOP in one very important circumstance. The derivations of (10) and (11) hold only when no security can be expressed as a combination of other securities—otherwise, demand may be nonunique and unbounded. If two equivalent portfolios have incoherent prices (e.g., $A = \emptyset$ and $p^A > 0$), then there is a risk-free portfolio with positive expected profit, and demand will be infinite—as a result, the equivalence will effectively be enforced in equilibrium.

Finally, we have proved that, in the current scenario, the group's collective investment behavior is indistinguishable from that of a rational individual.

Proposition 10 (Representative "super-agent", Part II) *Suppose that agents $i = 1 \ldots N$ with beliefs $\Pr_i(A_j)$ for disjoint securities $j = 1 \ldots S$, and utility functions $u_i(y) = \ln(y + b_i)$, effect equilibrium prices $q^{(j)}$. Then their aggregate (total) demand for each security is equal to the demand of a single representative "super-agent", with beliefs $\Pr_0(A_j) = q^{(j)}$ and utility $u_0(y) = \ln(y + b_0)$, where $b_0 = \sum_{i=1}^{N} b_i$.*

Proof. Analogous to the proof of Proposition 8. □

6 Conditioning Beliefs on Market Prices

All of the results in this paper are *equilibrium* results, and we have so far ignored any changes that might occur before an equilibrium state is reached, for example the revelation of evidence bearing on the likelihood of events. In particular we seem to have endowed each agent with a *fixed* probability for each event. In fact, there is nothing in our derivations that depends on static beliefs. The equations are all equally valid if we substitute $\Pr(A|E)$, were E is some possibly changing evidence. The going prices of securities make for natural evidence variables, as they reveal information about the other participants' beliefs. One could imagine that an agent's posterior belief depends only on the corresponding price, $\Pr(A|p^A)$, or depends on all prices, or even depends on the entire time history of all prices. In this section, we take an initial look at the first possibility.

Consider the case described in Section 5, where the market consists of S disjoint securities, and each agent has logarithmic utility for money, so that his optimal demand is of the form (10). We now add a simple generalization: each agent's posterior probability is a weighted average of his prior probability and the current price. Mathematically,

$$\Pr(A_j|p^{(j)}) = w\Pr(A_j) + (1-w)p^{(j)}.$$

We have no Bayesian justification for this update rule in terms of a likelihood function; we adopt it strictly for its simplicity, and to illustrate adjusted demands and prices in closed form. Now (10) becomes,

$$\begin{aligned}
x^{(j)}(\mathbf{p}) &= b\left(\frac{\Pr(A_j|p^{(j)})}{p^{(j)}} - \frac{\Pr(A_0|p^{(j)})}{1-p^{(1)}-\cdots-p^{(S)}}\right) \\
&= b\left(\frac{w\Pr(A_j)}{p^{(j)}} + (1-w)\right) - \\
&\quad\; b\left(\frac{1+w(-\Pr(A_1)-\cdots-\Pr(A_S))+(1-w)(-p^{(1)}-\cdots-p^{(S)})}{1-p^{(1)}-\cdots-p^{(S)}}\right) \\
&= b\left(\frac{w\Pr(A_j)}{p^{(j)}} + (1-w)\right) -
\end{aligned}$$

$$
\begin{aligned}
& b\left(\frac{1+w(1-\Pr(A_1)-\cdots-\Pr(A_S))-w+(1-w)(1-p^{(1)}-\cdots-p^{(S)})+w-1}{1-p^{(1)}-\cdots-p^{(S)}}\right) \\
= \;& wb\left(\frac{\Pr(A_j)}{p^{(j)}}-\frac{\Pr(A_0)}{1-p^{(1)}-\cdots-p^{(S)}}\right). \qquad (12)
\end{aligned}
$$

Thus the relative weight that an agent places on public prices, as compared to his own prior, simply attenuates demand by a multiplicative factor.

We could alternatively work directly with the equilibrium price (11), which in the current discussion becomes an implicit equation:

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

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

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

which is consistent with (12), and again has the form of the LinOP.

We can go one step further, by allowing agents to associate different weights to their priors for different events. Mathematically, each agent's posterior for event A_j is $\Pr(A_j|p^{(j)}) = w^{(j)}\Pr(A_j) + (1-w^{(j)})p^{(j)}$. In this case, the competitive equilibrium price is

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

which is a derivative LinOP where the weights can vary across both agents and events. Although such generalized weights seem natural (an agent may have more expertise or confidence in assessing some probabilities over others), and have been suggested by other researchers [Bordley and Wolfe, 1981, French, 1985], there is no known axiomatic justification for them [Genest and Zidek, 1986]. They emerge in our market context from the simple consideration that agents may learn from public prices.

Finally, analogous derivations apply to the situation in Section 4, where agents have constant risk aversion, if the posterior update rule is a weighted *geometric* average of prior and price:

$$
\Pr(A_j|p^{(j)}) = \kappa\left([\Pr(A_j)]^{w^{(j)}}[p^{(j)}]^{(1-w^{(j)})}\right),
$$

where κ is a normalization constant.

7 Related Work

The idea of using markets as a belief aggregation mechanism was proposed as early as forty years ago by Eisenberg and Gale [1959]. Inspired by 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 H obtains, agents share the total amount bet in proportion to their bets on H. 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 [1967] 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 [1986], 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.

Another exploration of markets for belief aggregation is Hanson's proposal for "Idea Futures" markets to encourage honest revelation of beliefs about future developments in science, technology, or other arenas of public interest [Hanson, 1995]. Although the mechanism has not been analyzed as a formal protocol, it has been operational as a game (i.e., no real money) on the world-wide web since 1994. The current version is called the "Foresight Exchange", operating at `http://www.ideosphere.com/`. Hanson has also described a scenario for employing a similar mechanism for coordinating computational agents [Hanson, 1991].

A third popular example of market-based belief aggregation is the Iowa Electronic Markets (IEM) system (`http://www.biz.uiowa.edu/iem/`), which runs real-money markets in uncertain political and financial propositions. Their market in the 1996 US Presidential election, for example, attracted wide participation and following. *Slate* magazine even used IEM as its main index of the election "horse race".

Although most work in economics and finance on securities markets assumes (or entails) homogeneous beliefs, Varian [1989] presents a model where agents have divergent probability distributions. He considers the effects of belief dispersion on asset valuation and trading volume, under various preference models.

Wilson [1968] obtains results similar to Propositions 7, 8, and 10 in a different context. He considers an explicitly cooperative group, called a *syndicate*, that faces a common decision under uncertainty, analyzing the case of single, continuous uncertain event ξ. A decision α yields a total payoff to the group of $p(\xi, \alpha)$, which is shared among members according to a prespecified *sharing function*. He derives conditions on the sharing function to ensure Pareto optimality. He shows that in some cases a *surrogate*, or representative, decision maker exists, and its risk tolerance is the sum of the individual members' risk tolerances. For the case of constant risk aversion, the surrogate probability is a normalized, weighted geometric mean, where the weights are normalized risk tolerances.

West [1984] presents one axiomatic justification for the (unnormalized) LogOP. He defines S_i as the value for which individual i is indifferent between D_i units of security A, and S_i dollars for sure. From this and u_i, one can derive $\Pr_i(A)$. 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 a geometric mean:

$$\Pr_0(A) = \prod_{i=1}^{N} [\Pr_i(A)]^{w_i},$$

which is an *unnormalized* LogOP. In contrast to the normalized version (2), 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.[8]

[8] One of the purposes of West's investigation was to demonstrate the pitfalls of blindly aggregating probabilities according to simple formulas (Mike West, personal communication). We share this motivation, and agree with the idea of deriving aggregate measures from behavioral postulates.

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.

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 system 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 exponential and logarithmic utility for money. For a market in disjoint securities, the equilibrium prices are the same as the probabilities derived by the LogOP and LinOP, respectively. The group's aggregate demand for securities is exactly that of a rational "super-agent", whose beliefs equal the market's equilibrium prices. 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 have discussed the ramifications of some simple adaptive agents that update beliefs based on price information.

Future theoretical work will aim to catalog which aggregation functions can and cannot be supported by a market framework, and under what conditions a representative "super-agent" exists. We plan to analyze when price equilibria are unique, and when the market is guaranteed to converge via a price adjustment protocol. We also plan concurrent computer simulations in more complex markets, including those with noncompetitive agents, to test robustness of convergence and properties of the aggregate prices. 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 [Wurman *et al.*, 1998].

A natural practical application for pooling opinions is as a sub-procedure within the more general context of *group decision making*, where agents' beliefs and utilities are combined to enable inference of optimal group decisions. Future plans include identifying and evaluating appropriate generalizations in pursuit of a market-based approach to group decision making in situations involving both asymmetric uncertainty and heterogeneous preferences.

Acknowledgments

Thanks to Didier Dubois, Robin Hanson, Jeff Mackie-Mason, Robert Nau, Steven Pollock, Ross Schachter, Jim Smith, Mike West, and Fredrik Ygge for providing helpful insights, feedback, encouragement and relevant pointers.

References

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

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

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

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

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

[de Finetti, 1974] Bruno de Finetti. *Theory of probability; a critical introductory treatment*, volume 1. J. Wiley, New York, 1974.

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

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

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

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

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

[Hanson, 1991] Robin Hanson. Even adversarial agents should appear to agree. In *IJCAI-91 Workshop on Reasoning in Adversarial Domains*, 1991.

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

[Hylland and Zeckhauser, 1979] 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.

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

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

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

[Lindley, 1985] Dennis V. Lindley. Reconciliation of discrete probability distributions. In J. M. Bernardo, M. H. DeGroot, D. V. Lindley, and A. F. M. Smith, editors, *Bayesian Statistics 2*, pages 375–390, Amsterdam, 1985. North Holland.

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

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

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

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

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

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

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

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

[Pennock and Wellman, 1996] David M. Pennock and Michael P. Wellman. Toward a market model for Bayesian inference. In *Proceedings of the 12th Annual Conference on Uncertainty in Artificial Intelligence (UAI–96)*, pages 405–413, Portland, OR, USA, August 1996.

[Pennock and Wellman, 1997] David M. Pennock and Michael P. Wellman. Representing aggregate belief through the competitive equilibrium of a securities market. In *Proceedings of the 13th Annual Conference on Uncertainty in Artificial Intelligence (UAI–97)*, pages 392–400, Providence, RI, USA, July 1997.

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

[Saari, 1995] Donald G. Saari. Chaotic exploration of aggregation paradoxes. *SIAM Review*, 37(1):37–52, March 1995.

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

[Varian, 1989] 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.

[Weerahandi and Zidek, 1981] 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.

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

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

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

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

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

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

[Wurman *et al.*, 1998] 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 (Agents-98)*, May 1998. To Appear.

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