The Wisdom of Competitive Crowds

Casey Lichtendahl, Yael Grushka-Cockayne, and Phil Pfeifer

Darden School of Business
University of Virginia

Presented at
DIMACS Workshop on The Science of Expert Opinion
Rutgers University
October 25, 2011

Background

- In 1906, Francis Galton observed that the average of 787 entries was remarkably close to the actual weight of an ox.

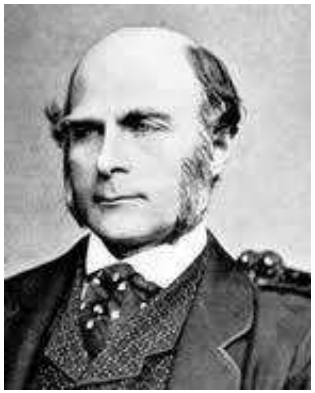

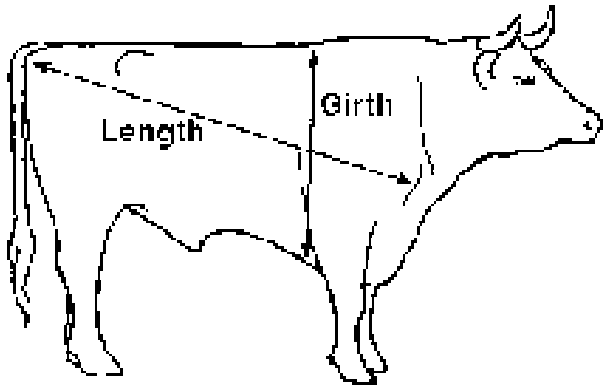

- The average guess was 1197 pounds, whereas the actual weight was 1198 pounds.

Background, Cont.

- Surowiecki (2005) in *The Wisdom of Crowds*, popularized the idea that the crowd's forecast, the average of the individual forecasts, often outperforms any individual forecast.

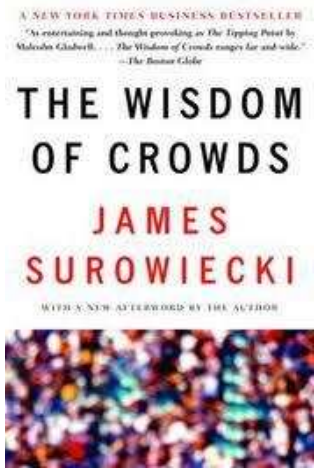

- It has a great deal of empirical support (Clemen and Winkler 1986; Clemen 1989; Armstrong 2001; Page 2007).
- Average point forecasts for GDP growth, etc. are reported by Philadelphia Fed's Survey of Professional Forecasters.

Motivation

- Sometimes overlooked in the retelling of Galton's tale is that "Those who guessed most successfully received prizes" (Galton 1907).

- The purpose of this paper is to examine how competition among forecasters influences the wisdom of the crowd.

The Gold Standard in Forecasting: Truthful Revelation

- "When outcomes are uncertain, planning must be based on forecasts—quite often, on forecasts submitted by others. Naturally, the planner wishes to ensure that these forecasts are prepared honestly and with an appropriate degree of care (Osband 1989, JPE)."
- "Since financial analysts' livelihoods depend on the accuracy of their forecasts...; we can safely argue that these numbers accurately measure the analysts' expectations (Keane and Runkle 1998, JPE)."
- "Prediction markets provide employees with incentives for truthful revelation (Cowgill, Wolfers, and Zitzewitz 2009)."
- Literature on scoring rules is predicated on the idea of truthful revelation (Winkler and Jose 2011).

Introduction

Our plan

- Extend work on competition among forecasters (Ottaviani and Sørensen 2006; Lichtendahl and Winkler 2007; Laster, Bennett, and Geoum 1999).
- Analyze a winner-take-all forecasting competition when forecasters have access to common and private information.
- Develop predictions of play in the competition modeled as a game of incomplete information.
- Show that the competitive crowd's forecast is more accurate than the truthful crowd's forecast and measure its degree of improvement.

Model

Game

- The planner organizes a winner-take-all forecasting competition.
- He invites k forecasters to each report a point forecast r_i for a continuous uncertain quantity x.
- The winner of the competition is the forecaster whose report is closest to the outcome of x.
- To the winner, the planner offers a prize proportional to the size of the crowd. Without loss of generality, we let this prize be equal to $\$k$.
- Each forecaster's objective is to maximize his/her expected prize.

Information Structure

- Each forecaster j receives two signals about x: common s and private s_j.
- $\theta \sim \mathrm{N}(\mu_0, m_0\lambda)$.
- $(s|\theta) \sim \mathrm{N}(\theta, m_1\lambda)$.
- $(s_j|\theta) \sim \mathrm{N}(\theta, n\lambda)$.
- $(x|\theta) \sim \mathrm{N}(\theta, \lambda)$.
- x, s, s_1,...,s_k are conditionally independent given θ.

Information Structure, Cont.

- An important parameter is n.
- As n increases, the private signals become more positively correlated: $Corr[s_i, s_j] = n/(m_0 + n)$.
- Each forecaster's true posterior beliefs are given by

$$(x|s, s_j) \sim \mathrm{N}\left((1 - w_t)\mu + w_t s_j, \frac{m+n}{m+n+1}\lambda\right)$$

where $m = m_0 + m_1$, $\mu = (m_0\mu_0 + m_1 s)/m_1$, and $w_t = n/(m+n)$ is the truthful weight on the private signal.

Pure Strategy Equilibrium

- **Proposition 1.** As the crowd grows large, there exists a limiting pure-strategy equilibrium where each forecaster reports $r_j = (1 - w_e)\mu + w_e s_j$ and exaggerates his private signal (i.e., $w_e > w_t$) using the weight

$$w_e = \frac{1}{2} \frac{\sqrt{n^2 + 4nm(m+1)} - n}{m}$$

if and only if $0 < n \leq 1$.

Pure Strategy Equilibrium, Cont.

- The equilibrium weight on the private signal is greater than the truthful weight.

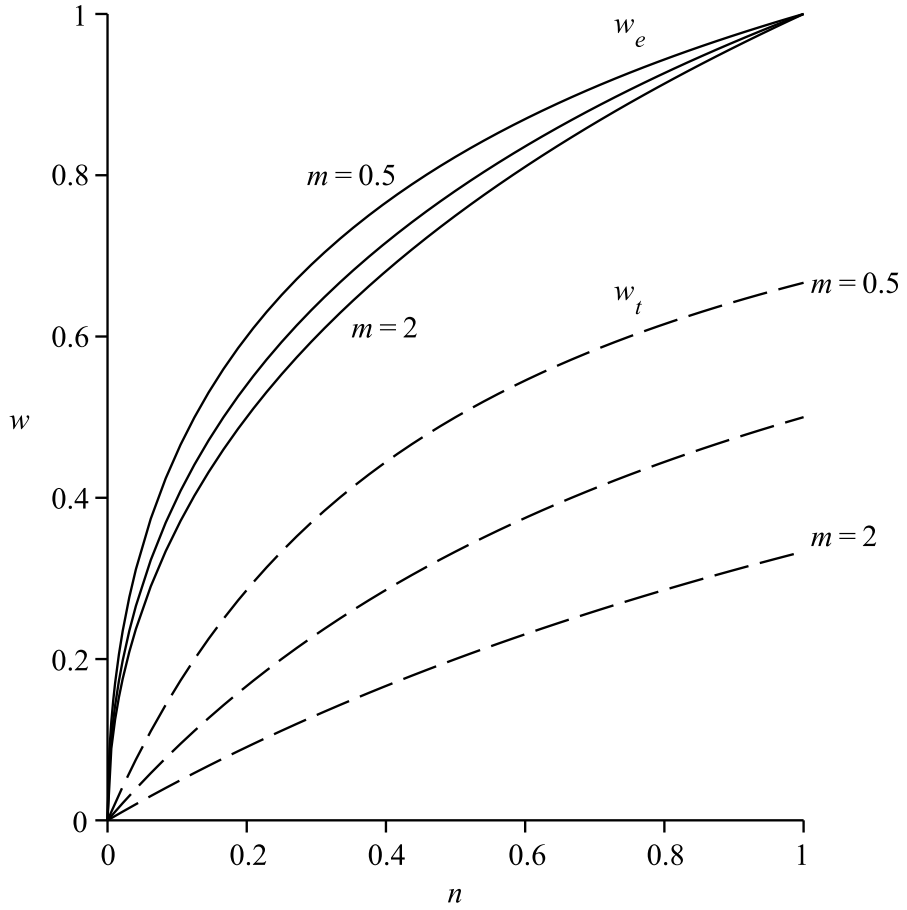

Intuition Behind the Pure Strategy Equilibrium

- Truthful reporting gives a forecaster the best chance to be close to the outcome.
- But with its weight on the common signal, his truthful report is also likely to be close to others.
- By exaggerating his distinguishing characteristic in the competition (i.e., his private signal), a forecaster will, on average, not be as close to the outcome, but when he is close, fewer forecasters are likely to be nearby.

Mixed Strategy Equilibrium

- **Proposition 2.** If $n > 1$, then, as the crowd grows large, there exists a limiting mixed-strategy equilibrium where each forecaster reports $r_j = s_j + ((n-1)/n)^{1/2}\epsilon_j$ and $\epsilon_j \sim \mathrm{N}(0, \lambda)$ for $j = 1, 2, \ldots$ independently.

Mixed Strategy Equilibrium, Cont.

- In this equilibrium, each forecaster ignores the common signal and issues a noisy report centered on his private signal.

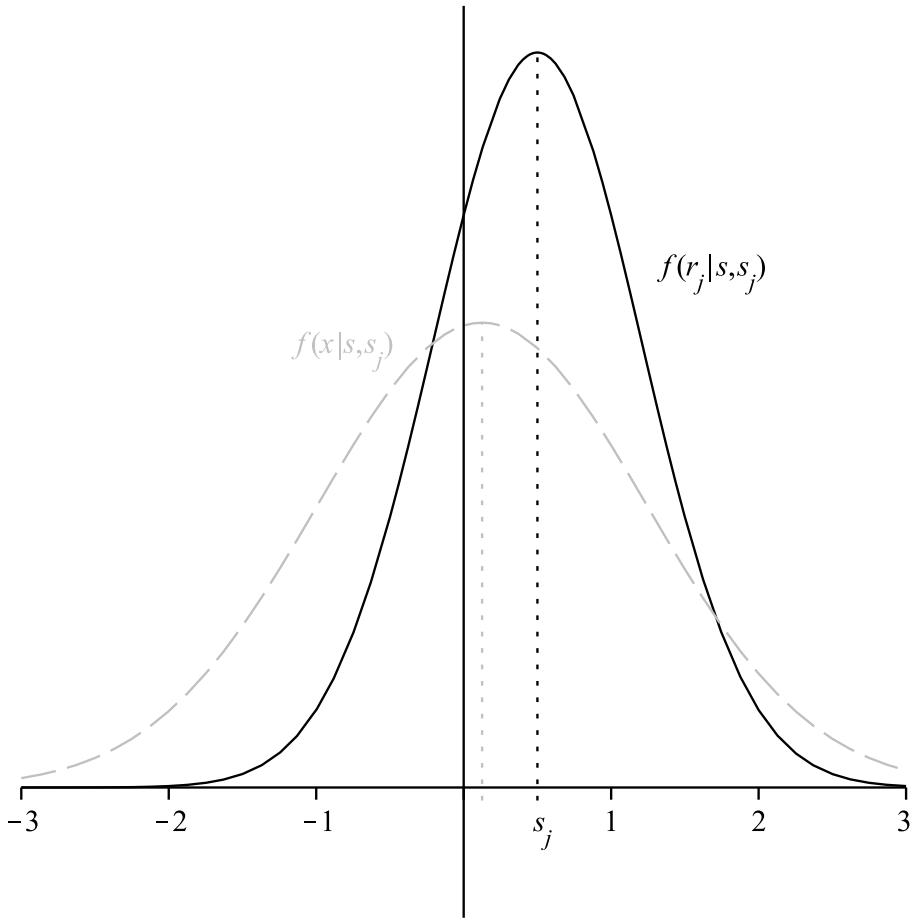

Intuition Behind the Mixed Strategy Equilibrium

- As the correlation among the forecasters' private signals increases beyond a certain threshold (i.e., $n > 1$), their private signals tightly cluster and the forecasters have less room to distinguish themselves with pure exaggerations of their private signal.
- Consequently, each forecaster has an incentive to move farther away from the others than pure exaggeration would entail.
- One stable way to uncluster is for each forecaster to mix around his private signal.

Reinterpreting Our Information Structure

- Recast common and private signals as collections of sample data.

$$\underbrace{x_1, \ldots, x_{m_1}}_{\text{Common sample}}, \underbrace{x_{m_1+1}, \ldots, x_{m_1+n}}_{\text{Forecaster 1's private sample}}, \ldots, \underbrace{x_{m_1+(j-1)n+1}, \ldots, x_{m_1+jn}}_{\text{Forecaster j's private sample}},$$

$$\ldots, \underbrace{x_{m_1+(k-1)n+1}, \ldots, x_{m_1+kn}}_{\text{Forecaster k's private sample}}, \underbrace{x_{m_1+kn+1}}_{\text{Quantity of interest}}$$

- Let $s = (x_1 + \cdots + x_{m_1})/m_1$ be average of common sample data.
- Let $s_j = (x_{m_1+(j-1)n+1} + \cdots + x_{m_1+jn})/n$ be average of private sample data.

Reinterpreting Our Mixed Strategy Equilibrium

- Suppose each forecaster reports the last data point x_{m_1+jn} in his private sample.
- This "report-the-last" strategy is consistent with the availability heuristic.
- Each forecaster "attempts to recall some instances and judges the overall frequency by availability, i.e., by the ease with which instances come to mind" (Tversky and Kahneman 1973, p. 208).

Reinterpreting Our Mixed Strategy Equilibrium, Cont.

- **Proposition 3.** A report-the-last strategy is the pure-strategy equilibrium in Proposition 1 when $n = 1$ and mimics the mixed-strategy equilibrium in Proposition 2 when n is a positive integer greater than one.
- Takeaway: This results suggests the availability heuristic may be well-adapted to competitive forecasting situations.

Accuracy of a Crowd's Forecast

- Define a crowd's forecast as the simple average of the forecasters' reports: $(r_1 + \cdots + r_k)/k$.
- Measure a crowd's accuracy by its MSE: $E[((r_1 + \cdots + r_k)/k - x)^2]$.
- Consider two types of crowds: the truthful crowd and the competitive crowd.

Relative Accuracy of the Competitive Crowd

- **Proposition 5.** As the crowd grows large, the competitive crowd's forecast is more accurate than the truthful crowd's forecast in the limit, and the limiting percentage MSE improvement is

$$\left(1 - \left(\frac{(1-w)^2}{m} + 1\right)\frac{(m+n)^2}{m+(m+n)^2}\right) \times 100\%$$

  where $w = w_e$ for $0 < n \leq 1$ and $w = 1$ for $n > 1$.
- For $w_t < w \leq 1$, this percentage MSE improvement is positive.

Relative Accuracy of the Competitive Crowd, Cont.

- Iso-percentage-improvement curves for the competitive crowd's forecast with $m_0 = 1$.

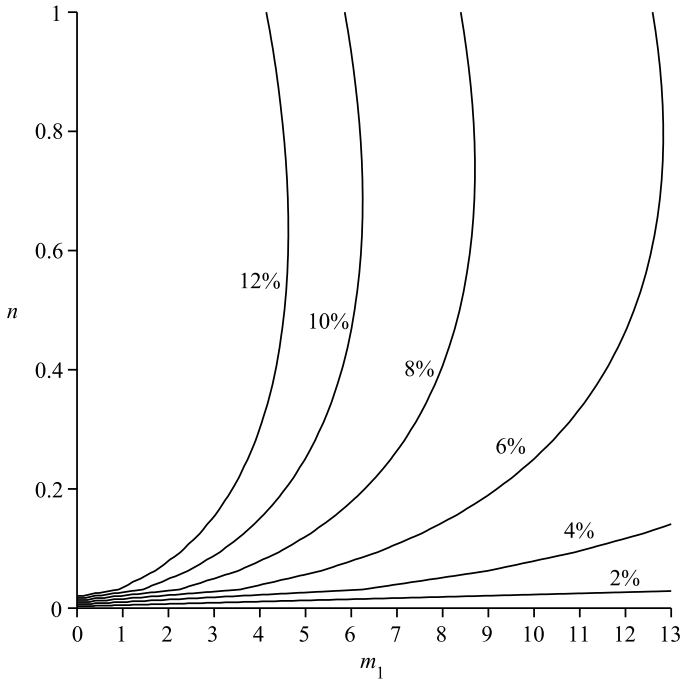

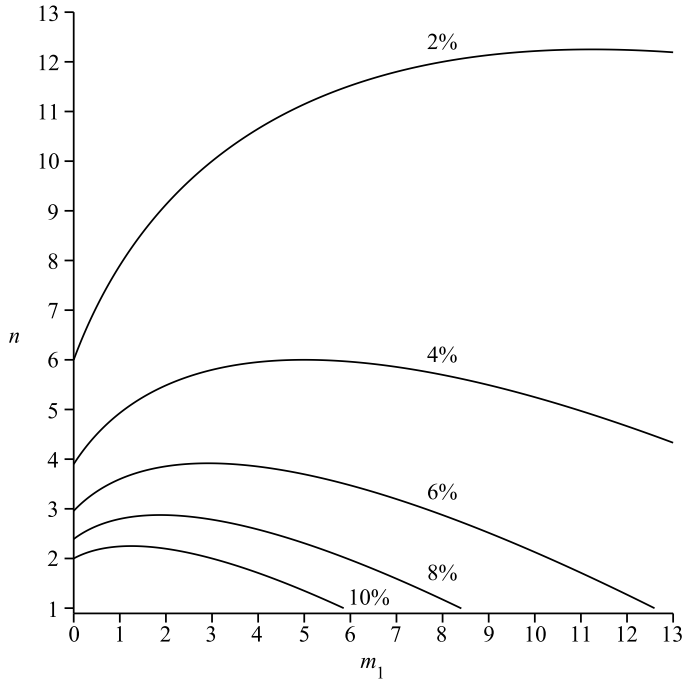

- Left panel has $0 < n \leq 1$, and right panel has $n > 1$.

Intuition Behind Accuracy Improvement

- When forecasters report truthfully and their information sources overlap, the crowd's forecast will contain redundant reports of the common information.
- In competition, each forecaster relies more on his private information in an effort to distinguish himself from the others.
- This results in less emphasis placed on the common information and a crowd's forecast that benefits from the diversity of opinions.
- In other words, the competition reduces the "public knowledge bias" (Chen, Fine, and Huberman 2004).

Implementation Challenges: Tolerance for Individual Errors

- One counterproductive thing a firm could do is to host a winner-take-all forecasting competition and later punish poor individual performance.
- **Proposition 7.** The competitive forecaster is less accurate than the truthful forecaster, and for $w_t < w \leq 1$, the percentage MSE improvement

$$\left(1 - \left(\frac{(1-w)^2}{m} + \frac{w^2}{n} + b^2 + 1\right)\frac{m+n}{m+n+1}\right) \times 100\%$$

is negative where $w = w_e$ and $b = 0$ for $0 < n \leq 1$ and $w = 1$ and $b^2 = (n-1)/n$ for $n > 1$.

Tolerance for Individual Errors, Cont.

- With $m_0 = 1$, the firm must be willing to tolerate individuals who are roughly 60% or more less accurate on an individual basis in order to take full advantage of the competitive crowd.

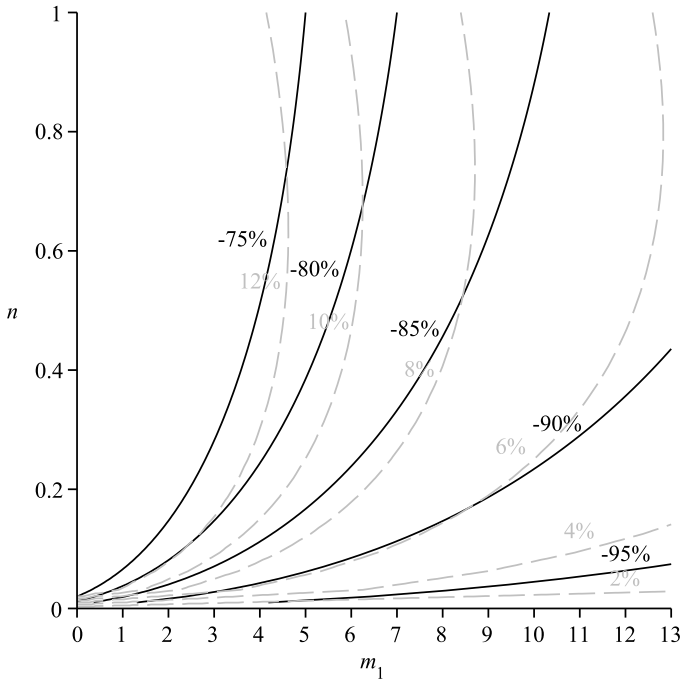

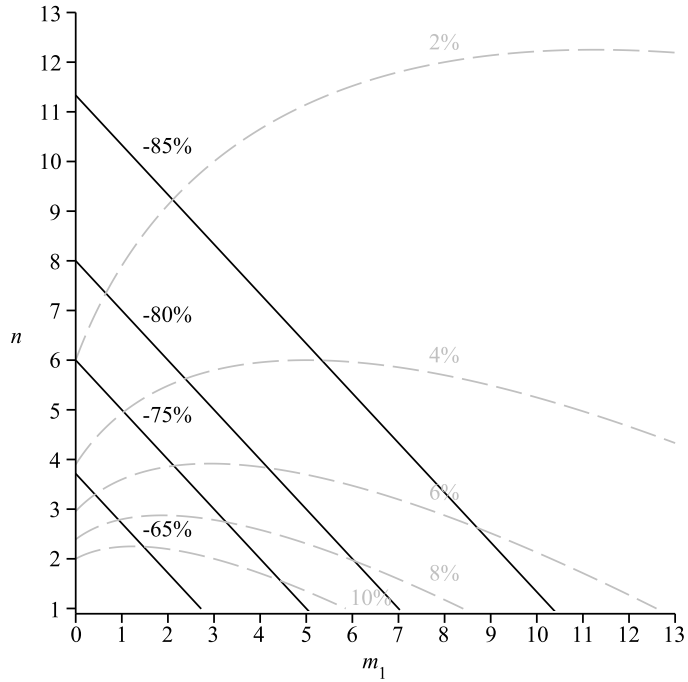

- Left panel has $0 < n \leq 1$, and right panel has $n > 1$.

Implications for Experimental Work

Testable Hypothesis

- A large number of subjects in a winner-take-all competition report differently when n is above and below the threshold 1.
- They purely exaggerate their private signals when $n < 1$.
- They report their last private data point when n is an integer at or above one.
- The competitive crowds in each of the two correlation treatments outperform control groups incentivized to tell the truth.

Prediction Markets vs. Winner-Take-All Competitions

- Leading Edge of Practice: Google prediction markets.

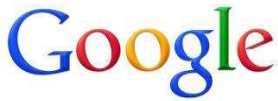

- When predictions markets have low trading volume, their prices may not reveal much information.
- Setup and maintenance costs for a prediction market can be high.
- At a lower cost, prizes can be awarded to winners of forecasting competitions.
- A competition has several advantages: low transaction costs, low cognitive load and low search costs.

Thank You!

(A copy of the paper is available on SSRN.)