Corporate Prediction Markets: Evidence from Google, Ford, and Koch Industries

Preliminary

Bo Cowgill
Haas School of Business, UC Berkeley

Eric Zitzewitz
Dartmouth College and NBER

Despite the popularity of prediction markets among economists, businesses and policymakers have been slow to adopt them in decision making. Most studies of prediction markets outside the lab are from public markets with large trading populations. Corporate prediction markets face additional issues, such as thinness, weak incentives, limited entry and the potential for traders with ulterior motives – raising questions about how well these markets will perform. We examine data from prediction markets run by Google, Ford and Koch Industries. Despite theoretically adverse conditions, we find these markets are relatively efficient, and improve upon the forecasts of "experts" at all three firms by as much as 25% reduction in MSE. Where inefficiencies exist, they disappear over time, as traders gain experience and as less skilled traders exit the market. The most notable inefficiency is an optimism bias in the markets at Google and Ford; this bias is larger on days that the company stock rises, suggesting it is at least partly behavioral and not strategic. Lastly, we provide a combination of anecdotal and statistical evidence that managers make use of the markets in decision making, and that the markets increased the visibility of firm goals and improved the analysis capabilities of participants.

Corporate Prediction Markets: Evidence from Google, Ford, and Koch Industries

The success of public prediction markets such as the Iowa Electronic Markets has led to considerable interest in running prediction markets inside organizations. Interest is motivated in part by the hope that prediction markets might help aggregate information that is trapped in hierarchies for political reasons, such as perceptions that messengers are punished for sharing bad news (e.g., Prendergast, 1993).

Markets in organizations face issues distinct from public and laboratory markets. If markets are run on topics of strategic importance, there is often a need to limit participation for confidentiality reasons. Limited participation makes markets thinner. In thinner markets, biases in participants' trading may have more influence on prices. Employees may optimistically bias their trading in order to influence management's view of their projects' performance or prospects. In addition to strategic biases, members of an organization may not be sufficiently dispassionate when making predictions. Optimism about a firm may make an individual more likely to select into employment, and belonging to an organization may likewise engender an favorable view of its prospects. Employees may suffer from other biases, such as probability misperceptions or loss aversion. Whereas in public prediction markets arbitrageurs may enter to eliminate any resulting inefficiencies, in corporate prediction markets, this entry may be less feasible.

This paper examines the efficiency of corporate prediction markets by studying markets at three major companies: Google, Ford Motor Company, and Koch Industries. These firms' markets were chosen because they span the many diverse ways that other companies have employed prediction markets in their businesses (summarized in Table 1).Our sample includes markets on a variety of topics about company performance, external events and markets meant to inform specific decisions. It includes markets into which the entire company was invited to trade, markets available only to hand-picked insiders and markets available to select divisions. It also includes a diversity of in in strength of incentives and in market mechanisms and design. Table 1 summarizes these characteristics and shows which other major corporations have used markets similar to those in our sample.

Despite large differences in market design, operation, participation and incentives, we find that prediction market prices at our three companies are well-calibrated to probabilities and improve upon alternative forecasting methods by as much as 25%. In the Google and Ford

markets, we find inefficiencies towards the beginning of the sample that disappear over time. Both companies' markets exhibit an initial optimism bias and exhibit overreaction to new information, and the Google markets exhibit an under pricing of small probabilities. Improvement over time is driven by two mechanisms: First, more experienced traders trade against the identified inefficiencies and earn higher returns, suggesting that traders become better calibrated with experience. Second, traders (of a given experience level) with higher past returns earn higher future returns, trade against identified inefficiencies, and trade more in the future. These results together suggest that traders differ in their skill levels, that they learn about their ability over time, and that self-selection causes the average skill level in the market to rise over time.

We also compare the predictions of the markets with alternatives available to management. Ford employs experts to forecast weekly vehicle sales, and we show that contemporaneous prediction market forecasts outperform the expert forecast, achieving an 25% lower mean squared error. Google and Koch did not have formal expert forecasts of the variables being predicted by its markets, but for markets forecasting continuous variables, expert opinion was used in the construction of the securities. Google and Koch created securities tracking the probability of the outcome falling into one of 3-5 bins, and an expert was asked to create bin boundaries that equalized ex ante probabilities. Koch also ran binary markets on whether a variable would be above or below an "over/under" median forecast. We find that at both Google and Koch market-based forecasts outperform those used in designing the securities, using market prices from the first 24 hours of trading so that we are again comparing forecasts of roughly similar vintage. Ford's in-house researchers separately created a statistical forecast of car sales, which the prediction market beat in 69% of trials[1].

With the Google data, we can also examine the role played by insiders in corporate markets. If we define an insider narrowly, as a team member for a project that is the subject of a market, or as a friend of a team member (as reported on a social network survey), we find that insiders account for 10 percent of trades, that insiders are more likely to be on the optimistic side of a market, and that insiders' trades are not systematically profitable or unprofitable. If we instead defined insiders as those traders we would expect to be most central to social and professional networks at Google (software engineers located at the Mountain View headquarters with longer tenure), we find that these traders are less optimistic and more profitable than other traders. So while a small number of insiders may trade optimistically in

1" Ford would not share the underlying data from this algorithm, so we cannot perform our own statistical tests on it.

markets on their own projects, perhaps reflecting either overconfidence or ulterior motivations, they are offset by a larger group of traders who also have relevant expertise and fewer professional reasons to be biased.

Taken together, these results suggest that despite limited participation and the potential for ulterior trading motives, corporate prediction markets perform well, and appear to do so for reasons anticipated by theory. Equilibrium market prices reflect an aggregation of the information and any subjective biases of their participants (Grossman, 1976; Grossman and Stiglitz, 1980; Ottaviani and Sorensen, 2007). Traders with an outside interest in manipulating prices may attempt to do so (Allen and Gale, 1992; Aggarwal and Wu, 2006; Goldstein and Guembel, 2008), but as emphasized by Hanson and Oprea (2009), the potential for manipulation creates incentives for other traders to become informed. Similar logic would apply to traders with subjective biases -- their presence creates incentives for participation by informed traders. Our results of initial inefficiency disappearing, with more experienced and skilled traders trading against the inefficiencies, are consistent with this set of predictions.

Our paper contributes to an increasingly extensive empirical literature on prediction markets and a much smaller literature describing experimental markets run at companies. Forsythe, et. al. (1992) and Berg, et. al. (2008) analyze the empirical results from the Iowa Electronic Market on political outcomes, finding that markets outperform polls as predictors of future election results. Wolfers and Zitzewitz (2004) and Snowberg, Wolfers, and Zitzewitz (2005, 2012) examine a broader set of markets, again concluding that prediction markets at least weakly outperform alternative forecasts. A series of papers have used prices from public prediction markets to estimate the effects of policies and political outcomes (Rigobon and Sack, 2005; Snowberg, Wolfers and Zitzewitz, 2007a and 2007b; Wolfers and Zitzewitz, 2009).

A smaller literature has documented experiments run by economists in corporate settings. Ortner (1998) reports on markets run at Siemens on whether projects would meet their deadlines. Chen and Plott (2002) and Gillen, Plott, and Shum (2013) report on sales forecasting markets run inside Hewlett-Packard and Intel, respectively. Hankins and Lee (2011) describe three experimental prediction markets run at Nokia, including one predicting smart phone sales. The sales forecasting experiment at Intel is of roughly comparable scale to the sales forecasting portion of the markets run at Ford;[2] the other corporate experiments discussed were on a much smaller scale than the markets at Google, Ford, and Koch.

[2] Gillen, Plot, and Shum (2013) analyze parimutuel markets that forecast sales for 46 product*period combinations. We analyze markets that forecast car sales for 78 product*period combinations (6 models times 13 weeks).

Our study differs from these prior and concurrent studies in several ways. First, the larger scale of the markets we analyze allows us to test for market inefficiencies with great statistical power, as well as to characterize differences in efficiency over time and across types of markets. Second, the microdata available on Google participants allows us to identify the characteristics of employees who trade with and against inefficiencies. Third, the markets we analyze are non-experimental in the sense that they were initiated by the companies themselves.[3] They are thus more field than field experiment. While there are downsides to this in that some research opportunities may have been missed, the advantage is that the markets we study are more likely to be representative of prediction markets as companies will implement them in the future.

Prior research informs our analyses of the specific inefficiencies we examine. Building on Ali's (1977) analysis of horseracing paramutuel markets, Manski (2006) shows that two common features of prediction markets -- budget constraints and the skewed payoff structure of binary securities -- can combine to cause a longshot bias in which prices on low-priced securities will be upwardly biased relative to median beliefs. Gjerstad (2005), Wolfers and Zitzewitz (2005), and Ottaviani and Sorensen (2007a) generalize this result to a broader set of risk preferences and information environments, showing that the sign of any bias is ambiguous.

The optimistic bias we document could either arise from genuine optimism or a conscious effort to manipulate prices. As Hanson and Oprea (2009) argue, the extent to which a (consciously or unconsciously) biased trader will affect prices depends on the ability of other traders to become informed and enter the market. In past episodes of apparent price manipulation in public prediction markets, other traders entered and traded against the apparent manipulation, reducing its impact on prices.[4] The price impact of manipulators in experimental markets is examined by Hanson, Oprea, and Porter (2006) and Jian and Sami (2012), with the former concluding that manipulation does not affect the accuracy of prices and the latter concluding that effects depend on the correlation of signals given to participants. In the field, the robustness of a corporate prediction market may depend on the ability and willingness of unbiased traders to enter the market and become informed, which may be constrained by limited participation.

3" The markets at Google were created by a group that included an author on this paper (Cowgill), but several years prior to his beginning his career as an economist.

4" See Wolfers and Zitzewitz (2004 and 2006), Rhode and Strumpf (2004 and 2008), Hansen, Schmidt, and Strobel (2004), and Newman (2012) for discussions.

To the extent that the optimistic bias we document is behavioral, our results also speak to the growing literature about overconfidence and excess optimism in organizations. Recent work shows that worker overconfidence has significant economic consequences for workers and firms. A theoretical literature explores how optimism may improve motivation of employees (Benabou and Tirole, 2002 and 2003; Compte and Postlewaite, 2004), reduce compensation costs (Oyer and Schaefer, 2005; Bergman and Jenter, 2007) or lead to risk-taking that generates positive externalities (Bernardo and Welch, 2001; Goel and Thakor, 2008). Empirical work suggests that the benefits of optimism-induced risk taking may be mixed. Hoffman (2011) [COMPLETE SENTENCE] acquire additional training. The consequence to the trucking firms is [AN INCREASE OF] nearly $7,000 in profits per truck, and [AN EFFECT IN WHICH DIRECTION] 1.5% of worker welfare. At the same time, Larkin and Leider (2012) show that overconfident employees select into suboptimal incentive contracts and Malmendier and Tate (2008) find that overconfident CEOs undertake mergers that perform poorly. Corporate prediction markets provide tools for both measuring and correcting employee optimism. Indeed, Koch Industries told us that a primary motivation for running markets was a desire to help employees become better calibrated forecasters.

The remainder of the paper is organized as follows. The next section provides background on the markets at Google, Ford, and Koch Industries. The following section presents our empirical analysis of the efficiency and inefficiencies of these markets. A discussion concludes.

1. Background on the Corporate Prediction Markets

The three companies whose prediction markets we examine, Google, Ford, and Koch Industries, are in different industries, have distinct corporate cultures, and took different approaches in their prediction market implementations. We will describe them in turn, and then discuss commonalities and differences.

1.1 Background on the companies and their use of prediction markets

Google is a software company, headquartered in Mountain View, CA, with an highly educated workforce and a high level of internal transparency. Its prediction markets began as a "20% time project" initiated by a group of employees that was lead by a co-author of this paper (Cowgill). Google opened its prediction markets to all employees.

The focus of Google's markets were whether specific quarterly "Objectives and Key Results" (OKRs) would be achieved. OKRs are goals of high importance to the company (e.g., the number of users, third-party quality rating, or on-time completion of key products). The attainment of OKRs was widely discussed within the company, as described by Levy (2011):

> OKRs became an essential component of Google culture. Four times per year, everything stopped at Google for division-wide meetings to assess OKR progress. [...]
>
> It was essential that OKRs be measurable. An employee couldn't say, ``"I will make Gmail a success"" but, ``"I will launch Gmail in September and have a million users by November."" ``"It's not a key result unless it has a number,"" says [senior executive] Marissa Mayer.

Google's markets were run with twin goals: 1) aggregating information for management about the success of an important project and 2) further communicating management's interest in the success of the project. [INSERT GOOGLE.ORG ANECDOTE, INFO ABOUT INTRANET PLACEMENT AND REFERENCES TO MEETING MINUTES/TEXT]

While the project aspired to cover every company-wide OKR, information on some projects was or needed to be too compartmentalized for them to be appropriate for a prediction market with mass participation. Thus a cost of wide participation was that some topics were necessarily off limits. Despite this, over X% of quarterly OKRs were covered by markets.

The markets on OKRs spanned the topics typically covered in other corporate prediction markets, including demand forecasting, project completion, and product quality (Table 1). Demand forecasting markets typically involved an outcome captured by a continuous variable (e.g., "How many Gmail users will there be by the end of Q2?"). An expert was asked to partition the continuum of possible outcomes into five equally likely ranges. In contrast, project completion and product quality OKRs were more likely to have binary outcomes (e.g., would a project be completed by the announced deadline), and these markets had two outcome securities. In addition to markets on OKRs, Google also ran markets on other business-related external events (e.g., will Apple launch a computer based on Intel's Power PC chip) and on fun topics that were designed to increase participation in the other markets.

Ford is an automobile maker, headquartered in Dearborn, Michigan, with operations in many countries. Ford chose to focus its prediction markets on two topics of especially high

importance: forecasting weekly sales volumes and predicting which car features would be popular with customers (as proxied in the interim by focus groups). Ford limited participation to employees with relevant expertise (in marketing, engineering, and vehicle design departments).

Sales forecasting is a common application for prediction markets: some of the Google OKRs involved future use of its products, and sales forecasts were the subject of markets at Hewlett-Packard (Chen and Plott, 2002) and Intel (Gillen, Plott, and Shum, 2013). Like H-P and Intel, Ford has an expert make official sales forecasts, with which we can compare the contemporaneous forecast of the market for accuracy. Unlike in the Google markets, in the Ford sales forecasting markets, a single security was traded with a payoff that was a linear function of the weekly sales for a particular model.

The features markets run by Ford provide an example of so-called decision markets (Hanson, 2002), in which markets ask about the success of a particular decision conditional on it being taken. Rather than defining success as a feature's long-term success in the marketplace, Ford chose feedback from focus groups as an more immediate outcome measure. A market asked whether a series of potential car features (e.g., an in-car vacuum) would reach a threshold level of interest in a focus group at a specified price. Each feature was represented by a binary security that paid off if the threshold was reached.

In the decision markets that have been discussed elsewhere in the literature, trades are unwound if the decision is not taken, so the price of a security reflect an expected outcome conditional on the decision. Ford's decision markets differed in that the features were always shown to focus groups eventually, so trades were never unwound. Rather than being used to decide whether to show a feature to focus groups, the market's expectation of the focus group outcome was combined with the actual outcome to provide a less noisy assessment. Focus groups are expensive to run, sample sizes are necessarily small, and so sampling errors can be meaningful. In contrast, opinions of employees may be cheaper to obtain, but employees are potentially biased, which is the reason focus groups are used in the first place. By asking employees to predict the results of focus groups, Ford sought to increase sample sizes while mitigating any biases in employee opinion.[5]

5" After the feature markets we analyze, Ford ran subsequent markets in which security payoffs were based on final markets prices, rather than focus group outcomes. These markets were beauty contests, with multiple equilibria. While the relative predictive performance of these markets would be an interesting research question, we were unable to obtain access to the focus group outcome data that corresponded to these markets. We therefore exclude these markets from our analysis.

Koch Industries is a diversified basic materials firm headquartered in Wichita, Kansas. It refines crude oil, transports oil and petroleum products, manufactures products including chemicals, building materials, paper products, and synthetic fibers like spandex. Many of its businesses are very sensitive to the macro-economy and/or to commodity prices, both of which were quite volatile during the period their markets ran (March 2008 to present). Koch decided to focus its prediction markets on macroeconomic and commodity prices that were relevant to its business. Some of these variables were already priced by existing futures markets (e.g., the future level of the Dow Industrials index or the West Texas Intermediate crude oil price) and some are the subject of macroeconomic forecasting (e.g., the unemployment rate and general price inflation), but others were not (e.g., the Spandex price in China, the Kansas City Fed's Financial Stress Index). In addition, markets were run on policy and political outcomes of interest to Koch, such as bailouts, health care reform, and the midterm and Presidential elections.

Koch's markets were started by a Senior Manager in Koch's strategic planning department, and participation was limited to a hand-selected group of employees with relevant expertise. While the number of participants in the Koch markets was much smaller than at Google or Ford, 57 out of 58 invitees participated, and the average participant placed 220 trades (compared with 48 at Google and 10 at Ford).

Koch's market creation had an additional motivation in creating the markets beyond obtaining forecast from them. "People are overconfident in their predictions," he says. "They either say 'X will happen' or 'X won't happen.'" They fail to think probabilistically, or confront their mistakes when they happen. The market therefore changes the way participants think, and I believe this not only improves our forecasts but has a positive spillover on everything else our team does." This stated goal is particularly interesting in light of our results, which suggest that markets are initially optimistic and overconfident (e.g., they display a bias away from a naive prior), but that these biases decline over time and more experienced traders trade against them.

Just under 60 percent of Koch's markets predicted a continuous variable. About one-fifth of these markets divided the continuum of possible future outcomes up into 3-10 bins as in Google's markets, while the other 80 percent specified a single "over/under" threshold. A very small number of markets (18 out of 1345) used the linear payoffs used by Ford. For markets predicting a discrete event (e.g., would President Obama be re-elected), usually there was a single security, which paid off if the specified event occurred.

1.2 Commonalities and differences

Table 1 summarizes the types of markets run by the three companies, and provides examples of a few other companies we are aware of that have run related markets. All six types of markets we are aware of being run at other firms were run at our three firms. Google ran markets of all varieties, while Ford focused on sales forecasting and decision makers, and Koch focused on external events. A few other firms have run many types of prediction markets (e.g., Eli Lilly, Best Buy), while others have run more focused experiments with one particular type of market.[6]

Table 2 contrasts the scale and some key features of our three markets. One important difference was the structure of the securities in the markets. As discussed above, Google used multiple bins for continuous outcomes (e.g., demand) and two bins for discrete outcomes (e.g., deadlines). In contrast, Ford used securities with linear payoffs for continuous outcomes and single binary securities for discrete outcomes. With a very small number of exceptions, Koch used single binary securities for discrete outcomes and an either bins or a single binary security combined with an "over/under" threshold.

The choice between two bins and single binary securities for discrete outcomes can potentially affect market efficiency if some participants exhibit "short aversion" (i.e., prefer to take positions by buying rather than selling). With bins, choices of boundaries can affect efficiency if participants take cues from them, as the literature on partition dependence suggests some do (Fox and Clemen, 2005; Sonnenman, et. al., 2011). We will test whether pricing suggests bias towards buying, as well as whether there is a bias towards pricing each of N bins at 1/N.

Two other important differences were the market making mechanism and the incentives provided to participants, which we discuss in turn.

1.2.1 Market mechanism

Google used internally developed software that implemented an approach similar to the Iowa Electronic Markets, in which the range of possible future outcomes is divided into a set of mutually exclusive and completely exhaustive answers, and securities are offered for each. For

---

6 We base these statements on public comments made at conferences by firms, as well as on interviews. In the latter case, we do not identify specific firms (e.g., the reference to "other pharma") unless we have received permission to, and we omit some examples we are aware of for brevity. It is of course possible that firms have run markets we are unaware of.

continuous variable outcomes, such as feature demand for a product, five bins were typically used, with the boundaries chosen by an expert to roughly equalize ex ante probability. For OKRs with discrete outcomes, like whether a deadline or quality target will be met, there are generally two outcomes, and no reason to expect the ex ante probability to be 0.5 (indeed, Google's official advice on forming OKRs is that they should be targets that will be met 65% of the time).

As on the Iowa Markets, participants can exchange a unit of artificial currency for a complete set of securities or vice versa. This approach should limit the impact of any bias towards buying on prices, since participants can buy any of the possible outcomes, and other participants can simultaneously sell all outcomes if their bid prices sum to more than one. Google did not have an automated market maker, although traders were observed placing arbitrage trades (e.g., selling all possible outcomes when their bid prices summed to greater than one or, more rarely, buying when their ask prices summed to less than one).

Ford and Koch used prediction market software developed by Inkling Markets (http://www.inklingmarkets.com/). Inkling's software uses an automated market maker described in Hanson (2003). The market maker allows trading of infinitesimal amounts at zero transaction costs, and moves its price up or down in response to net buying or selling. The automated market maker ensures that traders can always place trades, which helps avoid frustration and is particularly important when participation is limited. In cases where securities are linked to mutually An issue with an automated market maker is that it must be set at an initial price, and if the initial price differs from a reasonable prior, then easy returns can be earned by being the first to trade. If relative performance (e.g., "bragging rights") is a source of motivation for trades, having performance depend to heavily on simply being the first to trade against an obviously incorrect price can be counterproductive. In addition, market prices will be at least initially biased towards initial prices, and trading may be insufficient to eliminate a bias towards initial values. As a result, Inkling users take some care in setting initial prices, or in setting bin boundaries so that initial prices of 1/N are appropriate.

Thus the use of Inkling mechanism could potentially reinforce potential biases toward pricing at 1/N discussed above. We will test whether prices at Ford and Koch are biased towards their initial starting values, particularly early in markets' life or in markets that attract few trades, when compared with the markets at Google.

1.2.2. Incentives

Incentives for successful trading were provided at all three firms. Google endowed traders with equal amounts of an artificial currency at the beginning of each quarter, and at the end of each quarter, this currency was converted linearly into raffle tickets for traders who placed at least one trade. The prize budget was $10,000 each quarter, or about $25-100 per active trader. The raffle approach creates the possibility that a poorly performing trader may win a prize through chance, but has the advantage of making incentives for traders linear in artificial currency. Awarding a prize to the trader with the most currency would create convex incentives which could make low priced binary securities excessively attractive, potentially distorting prices.

Ford also used a lottery that created incentives that were linear in the currency used by the marketplace. For legal and regulatory reasons, it was not able to offer prizes to participants based outside the U.S., but we are told that this was a small share of participants in the markets we analyze.[7] Ford's incentives in North America were smaller than Google's, consisting of several $100 gift certificates.

Koch did not offer monetary incentives for its traders, but publicized the most successful traders. The high participation rate of eligible Koch traders suggests that the prediction markets were emphasized by management, and thus reputational incentives to perform should have been meaningful. Unfortunately, the convexity of the relationship between reputational incentives and performance is unclear. Given that Koch publicized its best performing traders, it may be reasonable to expect traders to have preferred low-priced binary securities, which offered opportunities for large payoffs.

Given the potentially convex reputational incentives used at Koch, we will test whether low-priced binary securities are overpriced at Koch, relative to Google. With smaller linear incentives for most of its participants, Ford might be expected to be an intermediate case between Google and Koch.

Results

2.1 Markets vs Experts

Supplying well-calibrated forecasts are the core purpose of prediction markets. Markets are just one of many sources of by which a company might obtain data needed for decisions. Companies also routinely turn to experts and statistical models to inform decisions. In this section and in Table 3, we compare the predictive accuracy of the markets to the firms' forecasting alternatives.

7" We have thus far been unable to obtain a precise percentage.

Our sample contains 2208 instances of comparable forecasts by both a market and an internal expert. This is roughly 40% of the total markets in our sample. In some cases, alternative forecasts were easily supplied by other institutions in the company, and in other cases this data was collected as part of the creation of these markets. We have organized these 2208 contests into four groups described below.

Table 3 contains “horse race” regressions (cite Fair/Shiller) in which we regress the realized outcome against the forecasts from both the market and experts. The resulting coefficients can be interpreted as the weight placed on one forecasting method relative to the other. In all cases, the regressions place more weight on the market's forecast. In no instance did the regression place zero weight on the expert's forecast; however, in three instances the regression suggested placing negative weight on the expert's forecast (controlling for the market's forecast).

The first set of contests comes from Ford, where we compare market and expert forecasts weekly autosales. The expert forecasts were provided by the company's human business forecaster. Before and after the arrival of the market, this individual's job was to provide “official” forecasts for use in company decisionmaking. This individual used a combination of human judgement and extrapolation from existing data, without formally training a statistical model.

Our sample of market and expert forecasts covers 13 weeks of six models of Ford cars in the United States: The Escape, F-150, Focus, Fusion, Super Duty and Lincoln (all models). Forecasts by the market and the expert were offered simultaneously, eleven days prior to the beginning of the week in question. The dataset thus contains a total of 78 observations for the human forecaster, prediction market and realized outcome. The results of our horse race regressions for these markets can be seen in Table 3, column one.

The official sales forecasts were not made public to traders within the company and was regarded a sensitive, compartmentalized figure. By contrast, the prediction market's forecast was visible within the company to all members of the related departments -- a number as large as 10K employees.

As a result, our forecasting competition is therefore (if anything) unfair to the prediction market because the human forecaster could have mimicked the market, but not vice-versa. Ford's human forecaster denies any mimicking took place, but he was the 6th most active participant in the prediction market out of 1,465 traders. His trading was net profitable, but he

was in the 25th percentile of earnings (75% of participants profited more). FOOTOTE: How is it possible for someone in the 25% percentile to be net-profitable? Ford's mechanism allowed for infinite liquidity at the expense of the market maker losing money.

Statistical forecasts are also commonly used inside companies for forecasting. Ford separately developed statistical forecasts for the same outcomes, based on a combination of historical car sales and search engine query volume (from Google Trends, http://trends.google.com). Ford would not share data from this algorithm, but its internal analysis of showed the market beating the algorithm in 69% of head-to-head trials – thus providing an additional data point about the relative usefulness of the markets against alternative forecasts.

Our second and third set of contests came from Google and Koch, respectively. Google and Koch created several securities to forecast the probability of a continuous variable falling into one of 3-5 bins, and an expert was asked to create bin boundaries that equalized ex ante probabilities. In Google's case, these market were usually about demand forecasting.

For example: To forecast the number of new Gmail signups in the upcoming quarter, the product manager for Gmail would be consulted to set five number ranges – designed so that each range had an equal shot at being true. The experts were encouraged to use whatever sources they preferred to set these ranges, and they often consulted historical data or statistical forecasts.

In these contests, we coded the realization of each bin as zero or one depending on if the outcome fell into the bin or not. The expert's forecast is coded as equal probabilities for all number ranges. In order to use a contemporaneous forecast by the market, we use the average first day forecast from the prediction market. There were 257 such markets at Google, and 1330 at Koch. The results can be seen in Table 3 (columns two for Google and three for Koch).

The forth and final set of contests come from Koch. In these markets, Koch sought to forecast a point estimate of a continuous variable (no bins were involved). The Koch markets featured an automated market maker to provide liquidity into the market. The market maker had to be initialized at a starting value, representing the firm's prior belief about the outcome. In these contests, this initialized value was used as the expert forecast. We used the average from the first day of trading as the market's forecast. 748 of Koch's markets fell into this category. Table 3 summarizes the results. PROVIDE COMMENT.

REPEAT SUMMARY

Biases and their tapering

The results above show a remarkable improvement above alternative forecasts. We now turn to studying biases in the markets' predictions. The specific biases we study regard optimism, aversion to extreme outcomes and collective overconfidence.

To study optimistic bias, we show that traders on average lose money by betting on their own companies' success. X% of the markets in our sample tried to predict an outcome with a clear directional effect on the company's success. An example might be, “How many new advertising clients will Google serve in Q2 2006?” Google would prefer this number be as high as possible.

To study this, we encoded the trades in our markets as being in favor or against the host firm's goals and preferences. In our encoding, the worst possible outcome for the firm is coded as -1, and the best is coded as +1 – with intermediate outcomes spaced evenly in between. In the example market above, the options for betting might be (a) “Zero to Ten new advertisers,” (b) “Ten to Twenty,” and (c) “Above Twenty.” (FOOTNOTE: The actual numbers have been concealed). In our encoding, the optimism of bets on (a), (b) and (c) would be encoded as -1, 0, and +1, respectively. For traders betting against these outcomes, the signs would be reversed (+1, 0 and -1 for sellers).

Many of the markets in our sample sought to predict a continuous variable, discretized into bins. For these markets, we can test whether the bets on the “extreme” outcomes (those on upper and lower tails) are consistently more profitable – reflecting an overpricing of events in the middle of the distribution of outcomes. In our regressions, we encoded the “extreme” events as on either side as +1, and the median event as zero (with outcomes in between spaced evenly).

Bibliography

Aggarwal, Rajesh K., and Guojun Wu. "Stock Market Manipulations." *The Journal of Business* 79.4 (2006): 1915-1953.

Ali, Mukhtar M. "Probability and utility estimates for racetrack bettors." *The Journal of Political Economy* (1977): 803-815.

Allen, Franklin, and Douglas, Gale (1992). "Stock-Price Manipulation." *Review of Financial Studies*, 5(3), 503–529.

Bergman, Nittai K., and Dirk Jenter. "Employee sentiment and stock option compensation." *Journal of Financial Economics* 84.3 (2007): 667-712.

Forsythe, R., F. Nelson, G. R. Neumann, J. Wright. 1992. Anatomy of an experimental political stock market. Amer. Econom. Rev. 82(5) 1142–1161.

Fox, Craig R., and Robert T. Clemen. "Subjective probability assessment in decision analysis: Partition dependence and bias toward the ignorance prior."*Management Science* 51.9 (2005): 1417-1432.

Grossman, Sanford J. (1976) "On the Efficiency of Competitive Stock Markets Where Traders Have Diverse Information." *Journal of Finance*, 31(2), 573–585.

Gillen, Benjamin J., Charles R. Plott, and Matthew Shum. "Inside Intel: Sales Forecasting using an Information Aggregation Mechanism." (2013).

Gjerstad, Steven. "Risk Aversion, Beliefs, and Prediction Market Equilibrium." (2005).

Goldstein, Itay, and Alexander Guembel. "Manipulation and the allocational role of prices." *The Review of Economic Studies* 75.1 (2008): 133-164.

Hankins, R., and Alison Lee. "Crowd sourcing and prediction markets." *CHI*. Vol. 11. 2011.

Hansen, Jan, Carsten Schmidt, and Martin Strobel. "Manipulation in political stock markets–preconditions and evidence." *Applied Economics Letters* 11.7 (2004): 459-463.

Hanson, Robin D. "Decision markets." *Entrepreneurial Economics: Bright Ideas from the Dismal Science* (2002): 79-85.

Hanson, Robin, Ryan Oprea, and David Porter. "Information aggregation and manipulation in an experimental market." *Journal of Economic Behavior & Organization* 60.4 (2006): 449-459.

Hanson, Robin, and Ryan Oprea. "A manipulator can aid prediction market accuracy." *Economica* 76.302 (2009): 304-314.

Jian, Lian, and Rahul Sami. "Aggregation and manipulation in prediction markets: effects of trading mechanism and information distribution."*Management Science* 58.1 (2012): 123-140.

Newman, Alexandra Lee. "Manipulation in Political Prediction Markets." *The Journal of Business, Entrepreneurship & the Law* 3.2 (2012): 1.

Oprea, R., Porter, D., Hibbert, C., Hanson, R., & Tila, D. (2007). Can manipulators mislead market observers. *Manuscript, University of California Santa Cruz*.

Ortner, Gerhard. "Forecasting markets—An industrial application." *Universität Wien* (1998).

Ottaviani, Marco, and Peter Norman Sørensen. "Outcome manipulation in corporate prediction markets." *Journal of the European Economic Association*5.2−3 (2007a): 554-563.

Ottaviani, Marco, and Peter Norman Sørensen. "Aggregation of Information and Beliefs in Prediction Markets." (2007b).

Rhode, Paul W., and Koleman S. Strumpf. "Historical presidential betting markets." *The Journal of Economic Perspectives* 18.2 (2004): 127-141.

Rhode, Paul W., and Koleman S. Strumpf. "Manipulating political stock markets: A field experiment and a century of observational data." *University of Arizona, mimeo* (2006).

Rigobon, Roberto, and Brian Sack. "The effects of war risk on US financial markets." *Journal of banking & finance* 29.7 (2005): 1769-1789.

Snowberg, Erik, Justin Wolfers, and Eric Zitzewitz. "18 Information (in) efficiency in prediction markets." *Information efficiency in financial and betting markets* (2005): 366.

Snowberg, Erik, Justin Wolfers, and Eric Zitzewitz. "Partisan impacts on the economy: evidence from prediction markets and close elections." *The Quarterly Journal of Economics* 122.2 (2007a): 807-829.

Snowberg, Erik, Justin Wolfers, and Eric Zitzewitz. "Party Influence in Congress and the Economy." *Quarterly Journal of Political Science* 2.3 (2007b): 277-286.

Snowberg, Erik, Justin Wolfers, and Eric Zitzewitz. "Prediction Markets for Economic Forecasting." *Handbook of Economic Forecasting, Vol 2A*. (2013)

Sonnemann, U., Camerer, C. F., Fox, C. R., & Langer, T. (2011). Psychological biases affect economic market prices. *Unpublished manuscript, University of Münster, Germany*.

Tziralis, George, and Ilias Tatsiopoulos. "Prediction markets: An extended literature review." *The journal of prediction markets* 1.1 (2012): 75-91.

Wolfers, Justin, and Eric Zitzewitz. *Prediction markets*. No. w10504. National Bureau of Economic Research, 2004.

Wolfers, Justin, and Eric Zitzewitz. *Five open questions about prediction markets*. No. w12060. National Bureau of Economic Research, 2006.

Table 1. Summary of corporate prediction markets at Google, Ford, Koch Industries and selected other companies

Market Topic	Example	Google	Ford	Koch	Other companies running similar markets
Company performance					
Demand forecasting	Ford F-150 sales next week	X	X		Arcelor Mittal, Best Buy, Chrysler, Eli Lilly, HP, Intel, Nokia
Project completion	Will chat be launched within Gmail by end of quarter?	X			Best Buy, Electronic Arts, Eli Lilly, Microsoft, Nokia, Siemens
Product quality	Google Talk sound quality rating	X			Electronic Arts, Eli Lilly, other pharma
External events	Spandex Price in China	X		X	Eli Lilly, other pharma
Decision markets	If feature X is offered, what will demand be?	X	X		Best Buy, GE, Motorola, pharma, Qualcomm, Rite Solutions, Starwood
Fun	Will the Interns win at the Koch company picnic?	X		X	Electronic Arts

Incentive Type	Example	Google	Ford	Koch	Other companies running similar markets
Monetary Prizes	$1 cash awards, credit in company store	X			
Non-Monetary Prizes	T-shirts, plaques	X	X		
No Incentives			X	X	

Market Mechanism	Example	Google	Ford	Koch	Other companies running similar markets
Decentralized					
Continuous Double Auction		X			Hewlett Packard, Intel, Nokia, Siemens
Centralized					
Hansonian Market Maker			X	X	
Other market maker					Electronic Arts

Criteria for Participation	Example	Google	Ford	Koch	Other companies running similar markets
Open to all employees		X			Unfinished
Open to selected employees			X	X	Unfinished

Contains Beauty Contests	Example	Google	Ford	Koch	Other companies running similar markets
No Beauty Contests					
Beauty Contests					

Particiation Among Invited	Example	Google	Ford	Koch	Other companies running similar markets
Unfinished					

Table Notes: In this table, we show that the markets in our sample span a wide diversity of implementations of corporate prediction markets. The details about firms outside of our sample come from public comments by firms and interviews. In some cases, the firm asked not to be identified. W e omit some examples we are aware of for brevity. It is of course possible that firms have run markets we are unaware of.

Table 2. Summary statistics on prediction markets in our sample

	Google	Ford	Koch
Industry	Software/Internet	Automobile	Basic materials
Ownership	Public (Ticker: GOOG)	Public (Ticker: F)	Private
Sample begins	April 2005	May 2010	March 2008
Sample ends	September 2007	December 2010	January 2013
Markets (questions)	270	149	1,345
Securities (answers)	1,116	149	4,278
Trades	70,706	9,258	12,655
Unique traders	1,465	902	57
Market mechanism	IEM-style CDA	LMSR	LMSR
Software	Internally developed	Inkling	Inkling
Style of market			
One continuous outcome (e.g., how many F-150s?)		68%	1.3%
One binary outcome (e.g., Project X done by Sep 30?)		32%	59%
Two outcomes (e.g., Yes and No securities)	29%		0.7%
3+ outcomes (e.g., bins)	71%		39%
Topic of market			
Demand forecasting	20%	68%	
Project completion	15%		
Product quality	10%		
External news	19%		96%
Decision	2%	32%	
Fun	33%		4%
Share for which optimism can be signed	58%	100%	71%

Notes: IEM-style CDA = continuous double auction with separate securities for each outcome (Forsythe, et. al., 1992)
LMSR = Logarithmic Market Scoring Rule (Hanson, 2003)

Table 3. Markets vs. Experts

Company	Ford	Google	Koch	
Prediction market type	One continuous outcome	3-5 bins	3-10 bins	One binary outcome
Expert forecast source	Expert forecaster	Derived from Bins	Derived from Bins	Contract over/under
Market topic	Auto sales	Demand	Macro numbers	Macro numbers
Timing	Simultaneous	Within one day	Within one day	Within one day
Prediction market forecast	0.67	1.12	1.01	1.16
	(0.10)	(0.12)	(0.19)	(0.19)
Expert forecast	0.38	-0.38	-0.11	-0.27
	(0.08)	(0.17)	(0.57)	(0.17)
Observations	78	257	1330	748
Unique markets	6	52	185	296
Time periods	13	10	45	58
MSE (prediction market)/MSE(expert)	0.742	0.727	0.924	0.908
P-value of difference with 1	0.104	0.00004	0.002	0.002

Notes: In this table, we run horse race regressions (Fair & Schiller) in which the dependent variable is the realized outcome. The independent variables include both the market's forecast as well as the expert's forecast. The coefficients for these are reported above, and tests against zero and for equality of the MSEs.

Table 4: Biases

Panel A: Google

	(1) Returns	(2) Returns	(3) Returns	(4) Returns
Optimism	-0.105***	-0.103**	-0.104**	-0.210***
	(0.040)	(0.041)	(0.041)	(0.061)
Optimism x Date				0.272**
				(0.114)
Favorite		-0.226**	-0.211**	-0.368***
		(0.090)	(0.089)	(0.134)
Favorite x Date				0.365
				(0.286)
Extreme			0.043*	0.041
			(0.023)	(0.041)
Extreme x Date				0.005
				(0.073)
Date (Scaled 0 to 1)				0.002
				(0.031)
Constant	-0.013***	-0.006	-0.022***	-0.024
	(0.004)	(0.004)	(0.008)	(0.019)
Observations	37,910	37,910	37,910	37,910
R-squared	0.043	0.067	0.074	0.097

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Panel B: Ford

	(1) Returns (All)	(2) Returns (Sales)	(3) Returns (Sales)	(4) Returns (Sales)	(5) Returns (Features)
Favorite			-0.585***	-1.195***	
			(0.141)	(0.321)	
Favorite x Date				0.920***	
				(0.326)	
Date (scaled 0-1)				0.016	
				(0.033)	
Constant	-0.073**	-0.013	0.005	-0.005	-0.106**
	(0.036)	(0.012)	(0.012)	(0.018)	(0.051)
Observations	9,258	3,262	3,262	3,262	5,996
R-squared	0.000	0.000	0.237	0.254	0.000

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Panel C: Koch

	(1) Returns	(2) Returns	(3) Returns	(4) Returns	(5) Returns
Favorite		0.021	0.026	0.017	0.003
		(0.046)	(0.050)	(0.050)	(0.053)
Optimism				0.021	0.026
				(0.021)	(0.022)
Extreme					-0.052
					(0.041)
Constant	0.017	0.017	-0.003	-0.010	0.032
	(0.012)	(0.011)	(0.013)	(0.014)	(0.030)
Observations	12,655	12,655	8,910	8,910	8,910
R-squared	0.000	0.000	0.000	0.002	0.003

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: In these tables, we study biases in prediction market trading. To be completed.

Table 5: Biases by Subsample. Dependent Variable: Returns

Panel A: Google

	(1) Demand Forecasting	(2) Company News	(3) Company Performance	(4) Decision Markets	(5) External News
Optimism	-0.043	-0.184**	-0.200***	0.116	0.108**
	(0.046)	(0.081)	(0.064)	(0.110)	(0.051)
Favorite	-0.212*	-0.062	-0.372***	-0.028	-0.446*
	(0.125)	(0.218)	(0.118)	(0.287)	(0.243)
Extreme	0.064	0.112**	-0.062	0.134	0.028
	(0.041)	(0.046)	(0.041)	(0.108)	(0.066)
Constant	-0.024***	-0.074***	0.045*	-0.032	-0.010
	(0.008)	(0.023)	(0.024)	(0.021)	(0.020)
Observations	12,387	7,430	10,057	489	6,898
R-squared	0.041	0.260	0.221	0.125	0.106

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Panel B: Koch Industries

	(1) All	(2) Politics	(3) Policy	(4) Stocks	(5) Growth	(6) Jobs	(7) Commodities	(8) Exchange Rates	(9) Eurozone	(10) Energy	(11) Inflation
Optimism	0.021	0.163*	0.034	0.001	0.049	0.175***	-0.019	0.012	-0.050	0.020	-0.095**
	(0.021)	(0.090)	(0.120)	(0.086)	(0.032)	(0.044)	(0.043)	(0.066)	(0.057)	(0.051)	(0.041)
Favorite	0.017	-0.058	-0.081	-0.257	-0.013	-0.453	0.699***	0.129	0.452***	0.412***	0.034
	(0.050)	(0.221)	(0.086)	(0.158)	(0.125)	(0.278)	(0.123)	(0.090)	(0.090)	(0.116)	(0.122)
Constant	-0.010	-0.026	-0.004	0.028	-0.019	0.043**	0.019	-0.059	0.002	-0.062**	-0.006
	(0.014)	(0.055)	(0.075)	(0.060)	(0.029)	(0.017)	(0.038)	(0.062)	(0.017)	(0.024)	(0.018)
Observations	8,910	449	382	1,309	2,205	657	290	462	166	492	1,429
R-squared	0.002	0.119	0.005	0.018	0.003	0.088	0.143	0.005	0.120	0.055	0.032

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 6: Biases Reduced Over Time

Table 7: Optimism and Stock Returns

Panel A: Google

	(1) Returns	(2) Returns	(3) Returns	(4) Returns
Optimism	-0.086**	-0.087**	-0.117**	-0.116**
	(0.035)	(0.034)	(0.050)	(0.048)
Extreme	0.043*	0.043*	0.043*	0.043*
	(0.022)	(0.022)	(0.022)	(0.022)
Favorite	-0.223***	-0.222***	-0.224***	-0.226***
	(0.082)	(0.083)	(0.080)	(0.079)
Optimism*Goog	-0.880	-0.878	-0.989	-1.045
	(0.719)	(0.714)	(0.742)	(0.722)
Optimism*Goog	-1.162	-1.149	-1.059	-0.854
	(0.797)	(0.772)	(0.760)	(0.712)
Optimism*Goog	-2.060***	-2.072***	-1.984***	-1.458*
	(0.752)	(0.762)	(0.673)	(0.749)
Optimism*Goog	-0.690	-0.689	-0.680	-0.647
	(0.437)	(0.436)	(0.436)	(0.447)
Constant	-0.020***	-0.020***	-0.017	-0.016
	(0.008)	(0.008)	(0.015)	(0.014)
Google stock r	Y	Y	Y	Y
Interactions of Google stock		Y	Y	Y
Day of week fixed effects and interactio			Y	Y
S&P and Nasdaq returns (t+1, t, t-1, t-2) and interactions				Y
Observations	37,910	37,910	37,910	37,910
R-squared	0.095	0.096	0.100	0.104

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Panel B: Ford

	(1) Ford Returns	(2) Ford Returns	(3) Ford Returns
Favorite	0.506***	0.465***	0.488***
	(0.140)	(0.136)	(0.136)
1 day future Diff between for		-0.028	0.010
		(1.016)	(1.031)
today's diff		0.795	0.875
		(1.504)	(1.329)
1 day lag diff		0.777*	2.240**
		(0.412)	(1.116)
2 day lag diff		0.230	0.461
		(0.422)	(0.576)
Forward One	0.234		
	(0.670)		
Today's Ford s	0.618		
	(1.187)		
1 day lag	1.065***		
	(0.256)		
2 day lag	1.488***		
	(0.414)		
Day of Week D	N	N	Y
Constant	-0.134***	-0.137**	-0.129
	(0.049)	(0.056)	(0.077)
Observations	9,258	9,258	9,258
R-squared	0.101	0.079	0.090

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table studies the relationship between profits and optimism in our corporate prediction market samples and the performance of each company's publicly-traded stock.

Standard errors are heteroskedasticity robust and adjust for clustering of outcomes within market and month.

Table 8: Biases, Experience and cumulative returns

Panel A: Google

	(1) Optimism	(2) Favorite	(3) Extreme	(4) Buy	(5) Returns
Cumulative Returns	-1.516***	-0.191**	-0.162	-0.584***	1.699***
	(0.270)	(0.078)	(0.169)	(0.223)	(0.096)
Experience (Log Trades)	-0.009	-0.031***	-0.119***	-0.119***	0.000
	(0.009)	(0.003)	(0.017)	(0.020)	(0.002)
Observations	75,820	141,412	75,820	141,412	141,412
R-squared	0.034	0.053	0.080	0.056	0.152

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Panel B: Ford

	(1) Sales Buy/Optimism	(2) Sales Favorite	(3) Sales Returns
Cumulative Returns	-0.204	-0.032***	0.025***
	(0.139)	(0.006)	(0.006)
Experience (Log Trades)	-0.117***	-0.008***	0.006***
	(0.023)	(0.001)	(0.002)
Constant	0.464***	0.030***	-0.018***
	(0.100)	(0.005)	(0.006)
Observations	2,810	2,810	2,810
R-squared	0.023	0.040	0.019

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Panel C: Koch Industries

	(1) Optimism	(2) Favorite	(3) Extremeness	(4) Buy	(5) Returns
Cumulative Returns	-8.135	0.992	-2.680	8.463	6.732***
	(5.320)	(1.681)	(6.671)	(9.194)	(2.368)
Experience	0.003	0.019***	-0.018	-0.061*	-0.012***
	(0.011)	(0.005)	(0.028)	(0.035)	(0.004)
Time (scaled 0-1)	0.007	-0.045*	0.012	0.194	0.035
	(0.093)	(0.025)	(0.107)	(0.127)	(0.026)
Constant	0.086	-0.017	0.105	0.277	0.101***
	(0.073)	(0.032)	(0.157)	(0.211)	(0.020)
Observations	8,696	12,318	8,696	12,318	12,318
R-squared	0.001	0.018	0.001	0.010	0.005

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: This table measure trading biases as a function of individuals experience, cumulative returns and the passage of time. These factors generally correlate with lower bias and higher trading profits.

Each observation is one side of a trade, and the dependent variable is characteristics of the bet. All regressions include trade fixed effects and a dummy variable for one particular extremely prolific trader. Standard errors are heteroskedasticity robust and adjust for clustering of outcomes within market and person.

Optimism is scaled between -1 and +1 based on whether the traders is betting for or against good outcomes for the firm. “Favorite” is coded as price-1/N in the Google markets an price-initial price for the Ford/Koch markets. “Extreme” is the absolute value of optimism. “Buy” refers to whether the trade was a buy or sell.

Table 9: Trader Characteristics, Biases and Returns (Google)

	(1) Optimism	(2) Favorite	(3) Extreme	(4) Buy	(5) Returns
Market Insider	0.126	0.066*	0.315***	0.404***	-0.013
	(0.099)	(0.034)	(0.101)	(0.153)	(0.046)
Friend of Insid	0.143**	0.034*	0.023	-0.194	-0.027
	(0.068)	(0.019)	(0.096)	(0.180)	(0.019)
Coder / Engine	-0.015	-0.157***	-0.371***	-0.492***	0.088***
	(0.085)	(0.028)	(0.095)	(0.129)	(0.033)
Hire Date	0.067**	-0.016	-0.059	-0.174***	-0.003
	(0.031)	(0.010)	(0.037)	(0.062)	(0.013)
NYC-Based	-0.150	-0.117***	0.090	-0.085	0.067*
	(0.115)	(0.032)	(0.108)	(0.159)	(0.037)
Mountain-View	-0.177*	-0.037	-0.045	-0.119	0.038
	(0.094)	(0.026)	(0.082)	(0.104)	(0.025)
Observations	75,820	141,412	75,820	141,412	141,412
R-squared	0.009	0.046	0.102	0.065	0.005

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: This table measure trading biases as a function of individual characteristics. The dependent variables are the characteristics of trades. Each observation is a trade. Each observation is one side of a trade, and the dependent variable is the characteristics of the trade for that trader.

All regressions include trade fixed effects and a dummy variable for one particular extremely prolific trader. Standard errors are heteroskedasticity robust and adjust for clustering of outcomes within market and person.

Optimism is scaled between -1 and +1 based on whether the traders is betting for or against good outcomes for the firm. “Favorite” is coded as price-1/N. “Extreme” is the absolute value of optimism. “Buy” refers to whether the trade was a buy or sell.