

Please provide contact information, including email.

Bo Cowgill. bo.cowgill@gsb.columbia.edu. (650) 492-8252.

- Please provide a short description of the course of study pursued during graduate work, including the type of course work pursued, major and minor focus of study, year of advance, and so on.

I am an incoming Assistant Professor at Columbia Business School, where I start in January 2016. I am an empirical microeconomist interested in digitization, organizations, personnel, labor and platform markets. Prior to my PhD at UC Berkeley I was an engineer at Google for six years.

- Please provide a short description of general research interests or other projects pursued. If the student has proposed a dissertation, and formed a committee, please provide a short description, expected date of completion, and general details.

Below are a set of relevant completed and in-progress projects.

Behavioral Biases and Machine Learning: Evidence from Résumé Screening

Extended abstract at the end of this document. Field experiment in progress w/ preliminary results.

Corporate Prediction Markets: Evidence from Google, Ford, and Firm X

With [Eric Zitzewitz](#).

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 Firm X (a large private materials company). 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 a 25% reduction in mean squared error. The most notable inefficiency is an optimism bias in the markets at Google and Ford. The inefficiencies that do exist

generally become smaller over time. More experienced traders and those with higher past performance trade against the identified inefficiencies, suggesting that the markets' efficiency improves because traders gain experience and less skilled traders exit the market.

Forthcoming at [Review of Economic Studies](#). Abstract in [EC'14](#).

Related regulatory activity: [CFTC Request for Comment](#) . [Reply from Google/Yahoo](#). [Reply from Microsoft](#). [Replies from others](#). Local copies: [Request for comment](#). [Google/Yahoo](#). [Microsoft](#).

Press discussion: [New York Times \(html\)](#), [FT Magazine \(html\)](#), [InfoWorld \(html\)](#). A digestable summary via the [Freakonomics blog](#). Earlier version: [Using Prediction Markets to Track Information Flows: Evidence from Google](#).

The Value of Hiring through Employee Referrals

With [Stephen Burks](#), [Mitchell Hoffman](#) and [Michael Housman](#).

Using unique personnel data from nine large firms in three industries, we provide stylized facts about the benefit to firms of hiring through employee referrals. Compared to non-referred applicants, referred applicants are more likely to be hired and more likely to accept offers, even though referrals and non-referrals have similar skill characteristics, both observable-to-the-firm (e.g., schooling) and unobservable-to-the-firm (e.g., cognitive and non-cognitive ability). With some notable exceptions on some metrics, referred workers tend to have similar productivity compared to non-referred workers. Compared to non-referred workers, referred workers are substantially less likely to quit and earn slightly higher wages. In call-centers and trucking, the two industries where we can calculate worker-level profits, referred workers yield substantially higher profits per worker than non-referred workers. These profit differences are driven by lower turnover and lower recruiting costs for referrals. For trucking, the value of hiring through referrals is larger for high-productivity referrers and in strong local labor markets.

Forthcoming, [The Quarterly Journal of Economics](#). Press discussion: [New York Times once \(html\)](#), [twice \(html\)](#).

Television Advertising and Online Search

With [Mingyu Joo](#), [Kenneth C Wilbur](#) and [Yi Zhu](#).

Despite a 20-year trend toward integrated marketing communications, advertisers seldom coordinate television and search advertising campaigns. We find that television advertising for financial services brands increases both the number of related Google searches and searchers' tendency to use branded keywords in place of generic keywords. The elasticity of a brand's total searches with respect to its TV advertising is .17, an effect that peaks in the morning. These results suggest that practitioners should account for cross-media effects when planning, executing, and evaluating both television and search advertising campaigns.

Management Science, Volume 60 Issue 1, January 2014, pp. 56-73.

Winner of the 2015 [John D. C. Little Award](#), given annually to the best marketing-related paper published in *Management Science* or *Marketing Science*.

Finalist of the 2015 [Frank M. Bass Dissertation Paper Award](#), given annually to the best marketing paper derived from a Ph.D. thesis published in *Management Science* or *Marketing Science*.

Competition and Productivity in Employee Promotion Contests

Why do firms use incentives that encourage anti-social behavior among employees? Rank-based promotion schemes are among the most widespread forms competition and incentives, despite encouraging influence-peddling, sabotage and anti-social behavior. I study a natural experiment using rich administrative data from a large, white collar firm. At the firm, competitors for promotions depend partly on dates-of-hire. I utilize the date-of-hire assignment as a source of exogenous variation in the intensity of intra-worker competition. I use the firm's multidimensional timestamped productivity logs as "time diaries" to study the amount, character and allocation of output across tasks. I find that competition has significant incentives for effort and efficiency -- as well as lobbying- and sabotage- like behaviors -- without affecting the quality and innovativeness of output. I also find that employees facing high competition are more likely to quit and join other companies, particularly higher-performing employees. Lastly, I show that competition induces workers to differentiate and specialize by concentrating effort into a smaller set of tasks. These results show that while workers respond to incentives from competition, they also seek to avoid it through sorting and differentiation strategies. The productivity gains from differentiation and specialization may partly explain the common use of these incentives by firms.

Working paper.

Mood Swings at Work: Stock Price Movements, Effort, and Decision Making

With [Eric Zitzewitz](#).

We show that daily stock price movements affect the mood, effort level, and decision making of employees. Positive current-day stock returns are accompanied by greater reported economic confidence and job satisfaction, shorter working hours, more optimistically biased beliefs about firm performance, tougher grading of innovative ideas, and tougher evaluation of interviewees. These effects are very short lived, lasting one or two business days. The effects on mood and some types of behavior are larger for employees with larger prior stock

and option grants. We show that the short-term effects of the (plausibly exogenous) shock to moods are generally opposite in sign to the cross-sectional correlations of mood and behavior. Whereas happier employees perform better and are more lenient evaluators, shocks that increase happiness are accompanied by lower work effort and tougher evaluation.

Revision requested by *The Journal of Economics, Management and Strategy*.

Incentive Effects of Equity Compensation: Employee-Level Evidence from Google

With [Eric Zitzewitz](#).

We examine whether differences in stock exposure motivate differences in performance using data from Google. Employees with more initial exposure to Google stock than peers hired at the same time and job grade perform slightly better across a variety of measures. To attempt to determine the direction of causality in this relationship, we use the fact that Google assigns option strike prices based on the stock price during an employee's first week to identify variation in initial stock exposure caused by luck rather than negotiation. We find that employees who begin work when the Google stock price is lower than surrounding weeks (and thus receive more initial stock exposure) do not perform better than their peers. In contrast, we find evidence that employees do respond to incentives tied to individual performance, such as those created by the promotion cycle. We also find that unvested equity compensation may indirectly affect performance by encouraging retention, since employees appear to make retention and performance decisions jointly.

Working paper.

The US-Canada Border Effect in Online Commerce

This is the first study to examine the US-Canada border effect using online trade data. A proprietary data set from Google allows us not only to measure the effect of the US-Canada border effect in the virtual world, but also to test Hillbery and Hummels (2003, 2008)'s "intermediate goods" theory of border effects. We find the border effect to have a large and statistically significant impact on online trade. At the aggregate level, our estimates for the average border effect indicate that intranational trade on Google is 5.8 times higher than international trade. When we disaggregate our data to the sector level, we find that Retail Trade, a category that encompasses cross-border transactions in final goods, is negatively affected by an average border effect of 3.8. Our results suggest not only that the US-Canada border hinders online trade, but also that theoretical explanations for the existence of the border effect should look beyond intermediate inputs.

Working paper. Based on earlier draft, [Gravity and Borders in Online Commerce](#).

Worldwide Gravity in Online Commerce

We analyse geographical patterns of cross-country Internet transactions using proprietary data from Google. Spanning all countries from 2008-2011 and over 10 billion online transactions, the data allow us to examine a fast growing, but neglected area of trade. We find the elasticity of distance to be around -0.5, indicating that distance effects persist in the virtual marketplace. We also discover that cultural characteristics, such as shared languages or religions, have a large impact in e-commerce, while economic ties, such as a common currency, have an insignificant effect. Our results improve upon an unpublished study of eBay transactions by significantly expanding the set of countries and correcting known econometric biases.

Working paper. Based on earlier draft, [Gravity and Borders in Online Commerce](#).

Behavioral Biases and Machine Learning: Evidence from Résumé Screening (Extended abstract)

How much do human behavioral biases distort hiring decisions? I share results from a large natural experiment and from a separate large scale field experiment. In the natural experiment, screeners are randomly assigned to job applicants. I show evidence that outcomes for applicants are highly dependent on the screener, and particularly the characteristics of the applicant-screener match.

Next, I evaluate a large scale field experiment attempting to manage these behavioral biases. In the experiment, an employer has developed software to parse resumes into a structured dataset containing over 1M variables. It then developed an algorithm to predict human-driven hiring outcomes using these variables using the past ~5 years of applicant data.

The algorithm is then deployed experimentally. Candidates above a threshold score in the algorithm are randomly assigned into treatment and control groups. Machine candidates are guaranteed an interview, while control candidates proceed through the status-quo, human process.

I find large differences (about 20%) in the numbers of candidates who proceed to interviews between treatment and control. Treatment candidates are ultimately 7% more likely to be extended an offer and accept the offer than control candidates -- this difference is driven not only by differences in interviewing rates but also in performance during the interviews. The

coefficient on the treatment is slightly higher than the coefficient on receiving a job referral from an incumbent employee at the firm.

The results demonstrate the ability of modern algorithms to automate, improve and displace white-collar labor. Hiring is a particularly interesting area for this result, as prior studies have emphasized the nuanced judgements of character, interpersonal skills, career trajectory and other intangible characteristics in screening workers.

Several results suggest that the the effect is driven by a decrease in homophily. The effect is particularly strong for non-traditional candidates with atypical academic backgrounds. These results are also surprising particularly because the machine was trained using human data. My results also suggest that the machine's performance improvement is the result of aggregating individual screeners' tastes, which may cancel out idiosyncratic tastes that would otherwise affect human screeners' decisions.