

INTRODUCTION TO SPECIAL ISSUE ON CORPORATE APPLICATIONS OF PREDICTION MARKETS

*Koleman Strumpf**

1. INTRODUCTION

Most academic journals are destined to be read by a few specialists and then quickly archived to dusty bookshelves. This special issue has a very different aim. It is primarily intended to be a gentle introduction for business practitioners interested in corporate applications of prediction markets.¹ While it is easy to find laudatory reports on these markets in the popular press, such stories are typically based on only one or two case studies. This special issue seeks to bridge this gap, by collecting the experiences of several pioneers in corporate prediction markets.

The authors in this issue are especially well-suited guides. They all have intimate first hand knowledge of these markets, whether organizers at their companies or as academics who have studied the actual workings of these markets. The articles are based on their presentations at the “Conference on Corporate Applications of Prediction/Information Markets,” held at the Kansas City’s Kauffman Foundation on 1 November 2007.²

Though the papers here highlight the potential for prediction markets to assist in firm decision-making, this is not a Pollyannaish assessment. Rather the special issue seeks to provide a sober assessment of both the benefits and

* Koch Professor of Economics, University of Kansas School of Business

¹Prediction markets utilize the knowledge of a pool of individuals to help forecast questions of importance to companies, such as whether a sales target will be reached or whether a project will be completed in a timely manner. A more recent development is the use of such markets to generate and evaluate new ideas, such as new products or cost saving procedures

²The sponsors of the conference were the University of Kansas School of Business, the Center for Applied Economics, and the Kauffman Foundation. I served as the conference organizer. The archive for the conference, which includes the complete schedule as well as notes on each talk, is available at:

http://people.ku.edu/~cigar/PMConf_2007/PredictionMarketsConference.html

possible costs of these markets. Prediction markets are not a magic elixir and will have a limited impact when they are applied to the wrong sort of questions or embedded in a dysfunctional environment. At the same time, under the right conditions these markets can be an extremely valuable aid for company decision-makers. The papers here glean such general lessons from actual markets.

In the next section I lay out why companies are a particularly interesting test-bed for prediction markets. While firms are well positioned to utilize the information which prediction markets generate, they also face special challenges such as opposition of employees which currently control information flow within the company. The third section provides an executive summary of each article, along with a recap of the accompanying commentary pieces. As my summary points out, the papers not only provide basic information but also raise a number of thought-provoking questions which are currently the subject of active research. The future development of corporate markets will require the teamwork of both practitioners and academics to solve these problems. Finally, the last section concludes

2. WHY COMPANIES?

The first known corporate applications of prediction markets occurred about a decade ago. Since then, over a hundred companies have run internal markets. These firms span the size spectrum, including some of the largest in the world to those with only a handful of employees, and they cover a broad range of sectors, including those whose products are abstract ideas to others which manufacture the most low-tech products. Why have such a broad range of firms become interested in prediction markets? The answer lies in a common problem facing firms.

Prediction markets have been employed to help combat a leading factor in bad decision-making, the isolation of executives from the views and insights of the company's workforce. Such seclusion is no accident but instead reflects one of the reasons companies exist in the first place, namely to avoid information overload for already busy executives. To reach this goal firms developed a hierarchy structure, and assigned to middle management the task of deciding how much and what information was transmitted from employees to higher-level decision-makers. But this system has its costs, as potentially useful information may be filtered out if it reflects poorly on those who control the information flow. At the same time, lower-level employees have little incentive to make reports which conflict with their managers, who may

later punish them for their candor. The net result is that executives may only receive one-sided information, and flawed and unproductive decisions can result.

This is where prediction markets come in. Suppose the CEO must decide whether to continue funding a research project, but he is concerned that he has been receiving overly optimistic reports on its prospects from managers who will benefit from the project continuing. A market on the project's prospects would allow front-line employees to convey more realistic information, and they could do so without fear of reprisal so long as trading is anonymous. Prediction markets are also likely to function better than other approaches currently in use. For example, group meetings are less likely to have frank discussions while suggestion boxes do not scale well (prediction markets tend to perform better when there are more participants). And while most workers dread the thought of meetings, markets are often considered a fun and typically do not require much incentives to generate active employee involvement.

This motivation aside, there are also particular reasons to think that companies are well positioned to utilize the information generated from prediction markets:

- company divisions often serve as standalone silos, and markets can be a means of integrating the pockets information contained in each
- executives may be interested not just in market aggregates, such as prices, but also the trades of particular groups of employees. For example, one could examine whether members of certain divisions are less prone to making biased forecasts.
- companies need real-time information about the many uncertain events surrounding their decision-makers
- firms can internalize the informational benefits of the market. A company can profit from the information generated from prices, since the market can be kept private and outside of the purview of competitors.

The last point is particularly important. Since the benefits of the markets largely accrue to the company, we should expect many prediction market innovations to first arise in a corporate setting. For example, companies have been the first to experiment with using these markets to generate and evaluate ideas. One of the surprising results of these markets is that the ideas come not just from research staff but also (to borrow a phrase from Jim Lavoie, a contributor to this volume) the “quiet geniuses” whose suggestions are typically ignored.

INTRODUCTION

At the same time corporate applications of prediction markets provide several unique challenges:

- they face special legal requirements such as avoiding insider trading laws
- they must overcome investor reluctance to a project with upfront costs and possibly delayed benefits
- there are impacts on employees, both detrimental (markets may distract staff away from their main responsibilities) and beneficial (there is often a gain in morale, as workers feel empowered because their market-mediated suggestions are impacting corporate decisions)
- employees may sabotage a project in order to profit from a market on its prospects. There is an unfortunate asymmetry here, since it is far easier for an employee to engineer an unfavorable outcome for the firm than it is to create a favorable one.
- the markets may overwhelm executives with too much information
- market organizers must allay concerns of middle management and others whose current role in the company is threatened by the market

The last two points highlight the delicate balance needed with internal markets. While the markets can provide executives with timely and unfiltered information, too much of this will be difficult for already busy executives to process and effectively utilize. This suggests that the markets should focus on targeted topics of particular interest in the executive suite. Similarly if the markets are too much of a threat to those who currently regulate information flows in the corporate hierarchy, then these managers will lobby hard to eliminate them or prevent them from gaining traction in the first place. Markets which focus on the creation of ideas, rather than forecasts, are likely to generate less of this kind of opposition, and thus may be easier to implement.

Such general principles aside, it is important to stress that prediction markets may only be appropriate for certain kinds of questions. For example, prediction markets have had difficulty in forecasting outcomes in which the fundamental information is quarantined from market participants (for example, in forecasting the decision of a secretive government committee).

As this discussion points out, prediction markets are likely to be a better fit for some companies (and for some questions) than others. To evaluate whether to use markets, it is important to weigh the costs and benefits listed above. The tradeoff may be favorable only in for certain kinds of firms, such as those where the corporate culture already embraces experiments.

3. PREVIEW OF PAPERS IN THIS ISSUE

The papers in this special issue span a wide range of topics related to corporate markets. The subject matter includes tips on how to establish markets, on how to tweak a market to obtain certain goals, what pitfalls to avoid, as well as highlights of some of the exciting state-of-the-art applications occurring in companies today. Accompanying each paper is a commentary which helps put the work in context. I briefly note some of the key highlights of each paper, but I highly recommend reading each in their entirety.

Jim Lavoie, a co-founder of Rite-Solutions, writes about his personal experience with internal markets (this paper is based on the keynote address at the 2007 conference). Jim lays out the reasons why he has made such markets an integral part of his company, and he provides a host of practical tips on how to get markets up and running. Jim also touches on some of the successes he has had with these markets, such as boosting employee morale and improving innovation. He also describes the novel use of these markets to both elicit and evaluate new ideas, an application which is now known as an Idea Market (the next paper in this volume discusses Idea Markets at GE). This is a wonderful paper for those curious about whether an internal market would be a good match for their company, and given its cutting edge application to idea markets it a fount of ideas for researchers and experienced practitioner alike. Jim has also set up a website, <http://artofinnovation.net/>, which allows readers to sample some of his Rite-Solutions markets.

Art Hall, who heads the Center for Applied Economics at University of Kansas, comments on Lavoie's piece. Art points out some of the potential pitfalls in the application of markets within companies. He shows how Rite-Solutions has been able to incorporate the markets into their corporate culture, which has allowed them to liberate pent-up creative instincts. At the same time, he notes the managerial challenges which must be met to get markets to operate successfully.

The second paper in the volume is by Brian Spears, Christina LaComb, John Interrante, Janet Barnett, and Deniz Senturk-Dogonaksoy, and it discusses GE's experience with Idea Markets in its Energy division (these authors organized and operated GE's markets). The authors show how GE has used these markets to generate and rank new ideas. A wealth of information is provided on topics ranging from participation rates to trader satisfaction with the markets to a quantitative appraisal of how well the markets elicit and rate ideas. There is also a very detailed description of the mechanics of the markets, such as the incentives given to traders and to creators of new ideas, which will be of particular interest to those looking for tips on how to

INTRODUCTION

structure a market at their own company. The authors also chart out some important avenues for future work, such as how to deal with the unusual trading behavior of idea creators. The numerical analysis in this paper serves as nice complement to the qualitative discussion in the earlier Rite-Solutions paper. Given that Idea Markets are one of the most exciting applications of corporate markets and GE and Rite-Solutions are the unquestioned leaders in this area, these papers are essential reads.

Marco Ottaviani offers commentary on the GE paper. Marco notes that the creation and evaluation of ideas are fundamentally thorny issues both within companies and more generally. He also highlights two leading challenges to the use of Idea Markets. Since the market expires before any of the ideas in the market are implemented, it is difficult to validate which idea is in fact best (in contrast the truth is observed in standard prediction markets, e.g. we know whether printer sales exceeded their goals for this quarter). A second issue involves the potential manipulation of the prices by an idea creator, who can potentially benefit if the market identifies his idea as the best. Marco offers some tentative solutions to these issues, and notes that working on a fuller solution is an important topic for both academics and practitioners.

The third paper in the volume is from Henry Berg and Todd Proebsting, who have jointly run a variety of prediction markets at Microsoft. Their paper focuses on some practical issues associated with the implementation of an automated market maker (AMM). An AMM is a mechanism which overcomes the illiquidity problem of many prediction markets: what happens if a market participant shows up and no one else stands ready to trade with him? The AMM stands ready to accept all trades on either side of a contract, and so allows the participant to immediately execute his trades. Such AMMs are now commonly employed in most corporate prediction markets as well as many public ones. Henry and Todd provide formulae which they have used in their implementation of an AMM at Microsoft. They also provide a range of practical tips for how these formulae can be tailored to other applications, as well as provide suggestions for other important aspects of market design which encourage both participation and appropriate trading behavior.

Robin Hanson, whose research laid out the foundations for most implementations of AMM's in prediction markets, is the commenter on the Microsoft paper. Robin provides an intuitive introduction to AMM's, which will be of particular use to those who have not encountered them before. He also discusses various practical issues with the AMM, and how it is flexible enough to be adjusted to suit a market administrator's particular goals.

In the fourth paper Adam Siegel, a co-founder of Inkling Markets, writes about his experience as a vendor and organizer of prediction markets (Inkling has helped a blue-chip list of corporate clients, such as Johnson & Johnson,

Chevron, and CNN, set-up and run markets. Inkling also runs several public markets which may be found at <http://home.inklingmarkets.com>). This article serves as an excellent how-to guide for companies just starting to experiment with prediction markets. Adam first lays out the case for why companies should consider using markets as a means for improved decision-making on topics such as project planning and resource allocation. He also provides a list of practical tips for market administrators, such as the need to promote interest among potential traders/employees and to provide a constant stream of new markets. A set of challenges which the new market is likely to confront is also provided. Adam concludes with a practical check-list which will allow companies to decide whether prediction markets are a good fit for them. A real bonus for practitioners is the insights and lessons provided by two of Inkling's clients, Cisco and a global consumer products company.

Paul Rhode's commentary on the Inkling paper points out the close connection of the benefits and challenges from corporate applications of prediction markets. One potential benefit from these markets is the facilitation of information flows within the company hierarchy, mitigating the censoring and potential bias from middle management. At the same time, Paul points out one of the reasons we have companies is to economize on transaction costs and information overload among executives. Middle management's winnowing of information can be a crucial time-saver for busy executives. And from a less altruistic perspective, these middle managers will strongly oppose markets if they view them as a challenge to their information hegemony. Paul points out that a successful corporate market will only be possible if these and other issues are adequately addressed.

Tom Bell authors the last paper in this issue. Tom's paper is a bit different from the others in that he is writing not about the workings of a corporate market, but rather he is addressing the rather murky legal environment which surrounds them (Tom is one of the leading law scholars studying firm markets). The legal issues are far more than an academic curiosity, as some firms have decided to steer clear of prediction markets to avoid exposure to litigation or government regulation. As Tom shows, understanding the legal status of a company's market is quite complicated. While issues related to gambling laws can be mitigated through the use of play money, any market potentially violates securities laws through the creation of new classes of insiders. Solutions to the latter questions are more subtle, and some (such as the public disclosure of market outcomes) may reduce or even eliminate the benefit of the information which the market collects. Tom goes on to suggest better solutions to these and other problems. He also discusses the wide range of regulatory changes currently being considered, any of which will have important implications for the legality of corporate markets. Clearly the legal

INTRODUCTION

issues surrounding corporate markets are in need of further study, and it will be leading scholars such as Tom who will help provide that scholarship. This article is a must read for anyone interested in using such markets.

Robert Litan's comment on this paper also notes the rather uncertain legal environment surrounding corporate markets. While many companies have braved these issues and started their own markets, the potential legal repercussions may have dissuaded others. Robert points out that regulators seem to have placed a premium on avoiding speculation rather than encouraging innovative new mechanisms such as prediction markets. He further notes that the overlapping jurisdiction of several regulators such as the SEC and CFTC itself creates legal anxiety and further discourages the spread of markets. Even if corporate markets do not receive a clear stamp of legality, Robert points out several ways in which current regulations could be improved.

4. CONCLUSION

My hope is that this volume will serve as a guide for future applications of prediction markets. At the same time, the papers here raise several thought provoking challenges which practitioners and scholars will grapple with for some time to come. It will be exciting to see the solutions and also to see whether prediction markets deliver on their potential in a corporate setting.

Readers interested in even more details on the topics raised here should consult the archive of the original Kansas City conference which includes presentations from the authors here as well as from other companies involved with prediction markets (a link is provided Section 1).

Finally, if you are interested in starting your own corporate market and have questions, please feel free to get in touch with me directly at cigar@ku.edu. If I am not able to directly answer your question, I will try and connect you with others in the field that can.