Information Systems Frontiers 5:1, 95–105, 2003
© 2003 Kluwer Academic Publishers. Manufactured in The Netherlands.

The Effect of Electronic Markets on Forecasts of New Product Success

Thomas S. Gruca, Joyce Berg and Michael Cipriano
Tippie College of Business, University of Iowa, S356 Pappajohn Business Bldg., Iowa City, IA 52242-1000, USA
E-mail: thomas-gruca@uiowa.edu

Abstract. *In this paper, we extend field experiments of real money prediction markets to the problem of forecasting the success of a new product. We collect forecasts using a traditional survey mechanism and a market mechanism. Our results suggest that market prices summarize the information contained in survey forecasts and improve those forecasts by reducing the variability of the forecast. However, we find no evidence of a "crystal ball" equilibrium. Our markets have considerable variability and predict only as well as the public signal provided by the HSX movie market game.*

Key Words. *electronic markets, new product success, forecasting, marketing research*

Introduction

Recently there has been increased interest in the possibility of using electronic markets to replace traditional marketing research techniques in the area of new product forecasting. For example, the Internet game "The Hollywood Stock Exchange" (HSX.com) allows traders to buy and sell "MovieStocks" whose value (in fictitious hsx dollars or $H) is based on a movie's four-week U.S. box office performance. A recent report by Pennock et al. (2000) claims that the resulting prices in the market are good predictors of the movies' actual performance in the market place. Other examples include CMXX.com, a web site of markets intended to predict the success of movies, music CD's and video games in Germany and Incentivemarkets.com which aims to predict the success of new pharmaceuticals using a market mechanism.

There are three primary motivations for harnessing the power of markets to assist with the important and difficult task of new product forecasting. First, traditional marketing research techniques can be expensive in terms of money and time spent to gather information from consumers (Urban and Hauser, 1993). In addition, typical marketing research studies are subject to a wide number of potential measurement, sampling and response biases (Churchill and Iacobucci, 2001).

Second, the Internet has swiftly developed from a means of communication to a means of exchange. Trading sites such as ebay.com and exchanges such as Chemdex.com are much more than electronic versions of traditional retailers. These sites can create and expand markets by bringing buyers and sellers together from around the world, any time day or night. By using an electronic market to gather information about potential new product success, an organization may be able to save time and money. Furthermore, an electronic market can more easily gather information from far-flung participants than any traditional marketing research technique.

The third motivation has been the progression of research on asset markets[1] as information aggregation mechanisms over the past 25 years from theory to lab experiment validation to successful field-testing. This stream of research provides compelling evidence that markets may be used effectively to gather and summarize information from a wide variety of sources through the interaction of traders. In the next section, we summarize this research briefly.

Prior Research on Markets as Information Aggregation Mechanisms

Rational expectations theory suggests that asset markets are able to distill disparate sources of public and private information into a single measure: price. In other words, price alone is a sufficient statistic for all information available to traders in a market (Lucas, 1972; Grossman, 1981). Consequently, through the mechanism of price, all traders in a market (and any observers) can come to share information held by each and every trader.

Therein lies the appeal of asset markets as information gathering tools. Using the market mechanism, a researcher could conceivably summarize all the information regarding the success of a new product in the single statistic, price. The only problem for the researcher is to create asset markets in which the value of the assets being traded are tied to an important marketing measure such as the sales or market share of a new product. By summarizing all trader information, the resulting prices in the asset market should correctly reflect the probability of various levels of sales or market share, depending on the focus of the market.

Two key tasks that markets must perform in order for prices to summarize the information available to traders are dissemination (traders must be able to learn from the market price) and aggregation (initial information and subsequent revisions of beliefs must be incorporated in price). In a series of lab experiments, Plott and Sunder (1982) show that asset markets indeed have the capability to disseminate information from informed to uninformed traders. A later series of lab experiments by Plott and Sunder (1988) shows that markets are also able to aggregate information so that all traders end up with the same information. These seminal papers show that asset markets are capable of the information processing tasks required by rational expectations theory.

Subsequent to the lab experiments were field studies that examined how well markets process information outside of the control of the experimenter. The best-known examples are markets used to forecast election outcomes. Beginning in 1988, the Iowa Political Stock Market (now known as the Iowa Electronic Markets or IEM) has accurately predicted a number of national, local and international elections (Forsythe et al., 1992; Berg et al., 2000). In many instances, the market's prediction of the popular vote was more accurate than that available from pre-election "trial heat" polls (Berg, Forsythe, and Reitz, 1997).

An important finding from these field studies is that traders need not be a representative sample of voters in order for the market to provide accurate forecasts. For example, in 1988 the traders were primarily students at the University of Iowa. Even though these traders were a distinctly non-representative sample of voters, the prices in the IEM markets predicted the election vote shares better than polls of a representative sample of voters.

This characteristic is important to the use of prediction markets for new product forecasting. If a representative sample of consumers were required to serve as traders in order to yield accurate results, much of the cost advantage of market-based forecasting would be eliminated. Many of the same costs incurred in traditional market research to identify, contact, and motivate a sufficiently large, representative sample of consumers to provide information to the researcher would also then be incurred in setting up a market. In fact, if real incentives are necessary for an accurate market forecast, then the cost of the market mechanism could exceed that of traditional marketing tools.

Limitations of current research

Clearly, the research discuss above provides compelling motivations to use asset markets in the area of new product forecasting. However, there are two important issues to consider in the implementation of market-based forecasting for new products. First, while lab experiments such as Plott and Sunder (1982, 1988) tell us much about dissemination of perfect information (information with no uncertainty) or the aggregation of complete information (instances in which traders in the aggregate have perfect information), little is known about the aggregation of information in markets where traders do not have perfect information either individually or collectively. Indeed, lab experiments by Forsythe and Lundholm (1990) and Lundholm (1991) found relaxation of information completeness or perfection leads to prices for assets that are quite far from their true underlying value. This is a particular concern in predicting new product success because those settings by definition involve incomplete information.

Second, the realm of successful field applications of market-based forecasting has been limited primarily to political events. In addition to the small number of published applications, the situation is very different from the typical new product situation. There is a great deal of historical information available (e.g., Rosenstone, 1983) as well as a high level of public interest in the outcome of major elections. These sources of information can be assumed to be superior to those available to traders trying to predict the success of a new product.

Study Overview

In this paper, we focus on an interesting and difficult new product forecasting problem: forecasting the box office receipts for a new movie. Academic interest in this area is reflected in the number of recent studies on this topic (e.g., Eliashberg and Sawhney, 1994; Krider and Weinberg, 1998; Neelamegham and Chintagunta, 1999; Sawhney and Eliashberg, 1996; Zufryden, 1996).

While the movie business may seem unique, there is an important parallel with other new product situations. Each movie is a unique entity, but it also shares characteristics with others introduced in the past (e.g., star, director, plot line). The same is true with most new products, which are a mixture of new and established benefits for consumers.

One of the important contributions of this study is our comparison of a forecast resulting from a traditional marketing research method (i.e. a survey) with one provided by prices in an electronic market. While prior research on political markets compared their results with polls of people not necessarily (and most likely not) trading in the markets, our survey data comes from the traders in the market, as they are required to submit a forecast of box office receipts before trading commences.

Since we have both types of data (survey results and market prices), we can examine a crucial question: How does the interaction of traders affect the accuracy of forecasts? Without knowing how trading affects forecast accuracy, there is no way to determine which alternative method is better. From the experience of the IEM political markets, we might expect that trading improves the accuracy of the prices as forecasts since traders have to "back up" their opinions about the election's outcome with a real money investment. Only those traders who perceive that they hold superior information would be willing trade in order to move prices in the direction consistent with their information, towards the eventual outcome.

On the other hand, over-confident traders might affect prices in the market if they stick to their forecasts irrespective of the (perhaps better informed) opinions of their fellow traders. Such a situation could result in prices that are less accurate than the aggregate of individual forecasts resulting from a survey. In laboratory experiments with informed and uninformed traders, Plott and Sunder (1982) have documented that traders with superior information set prices. However, the type of information available in a new product forecasting setting differs substantially from the structured information available in a laboratory market. Information about new products is less structured and the quality of individual pieces of information held by individual traders is not verifiable ex ante.

The second major issue we address is the ability of traders to arrive at a forecast that is better than any available from public information. In political markets, the benchmark for accuracy is pre-election polls. In these markets, traders seem to be able to improve on the information that is publicly available since the resulting market prices are able to forecast vote shares more accurately than pre-election polls. This is known as a "crystal ball equilibrium" (Plott and Sunder, 1982:678), a market that provides a forecast that is significantly better than one available from public sources of information. In this study, we examine whether electronic markets have the general characteristic of finding the crystal ball equilibrium or whether such a result is unique to political markets. As our analogue to pre-election polls, we use the forecasts available from the Internet game known as the Hollywood Stock Exchange (HSX.com).

In the next section, we describe the IEM in general and the Movie Box Office Market, in particular, which is the setting for our study.

The Iowa Electronic Markets

The IEM is an asset market, in contrast to a goods market in which physical goods are exchanged. The assets being bought and sold are futures contracts whose values are tied to a future event. While the best-known markets focus on the outcomes of political elections, there are also markets to predict stock price levels, corporate earnings, stock returns, changes in Federal Reserve policy and other economic events.

Participants in the IEM can act as both buyers and sellers, so we refer to them as traders. Trading occurs through a computerized, anonymous double auction,[2] accessible 24-hours a day through the Internet (www.uiowa.edu/iem). To trade, a trader must open an account, investing a maximum of $500 in real money. Traders then face the real possibility of losing or increasing their investment through their trading behavior. No transaction fees are charged.

To buy a contract, a trader can execute a market order and buy at the current best price available (lowest ask from another trader) in the market. Alternatively, the trader can submit a limit order. This would include an offer to buy (bid) at a higher price and a time limit on the offer. (An analogous process can be followed to sell or offer to sell contracts.) The limit orders (bids/asks) are queued by price and submission times. The best prices in each queue are displayed to traders. All trading of individual contracts and the resulting prices in the market are determined by exchanges between individual traders who remain anonymous to each other.

Traders can also acquire contracts by purchasing a bundle consisting of one of each of the contracts in the market. These bundles can be purchased from or sold to the IEM exchange at any time for $1, the guaranteed aggregate liquidation value (payoff at the end of the market) of the bundle. This feature of the market allows contract supply to expand and shrink as traders demand without contaminating the individual contract prices as set by the traders.

In general, there are two types of IEM markets: winner-take-all and share markets. Because the IEM Movie Box Office Market is a winner-take-all market, we omit discussion of share markets in this paper.[3] In a winner-take-all market, contracts are constructed as a set of mutually exclusive and collectively exhaustive outcomes. At the end of the market, exactly one of these contracts pays off while the others expire worthless. Since the contracts are designed to pay off $1 or $0, the corresponding market prices of the contracts can be interpreted as an estimate of the probability of the corresponding event occurring.

At the end of the market, the IEM liquidates the contracts held by traders. The process consists of an exchange of $1 for each winning contract in winner-take-all markets.[4] Since the IEM charges no transaction fees, this is a zero-sum market in which all investments by traders are returned to the traders collectively.

IEM movie box office markets

We began movie prediction markets in the 1995–1996 holiday movie season with markets tied to the box office receipts of "Money Train" and "Nick of Time." These initial markets met with a great deal of interest from traders in other (i.e. political) prediction markets. In the next section, we describe our field experiments using the IEM movie markets in detail.[5]

The outcomes we attempted to predict were the U.S. box office performances of two movies over a four-week period. For each movie, the related market consists of four to six contracts. Each contract is associated with a mutually exclusive and collectively exhaustive range of box office receipts. The liquidation values of these contracts are based on the movie's box office performance within the specified time frame. If the box office receipts fall within a given contract's range, then the contract pays $1; all other contracts pay $0. For example, in the Spring 1998 Movie Market, there were 4 contracts associated with the movie "Lost in Space." Their definitions are given in Fig. 1.

These contracts are a set of outcome-spanning Arrow-Debreu securities. Prior research by Plott and Sunder (1988) suggests that asset markets using this contract framework have been successful in aggregating information from individual traders.

Description of traders

The traders in these markets were predominantly masters-level business students at a major Midwestern university. In some markets, which we identify as "closed" markets, all of the traders are students enrolled in the same marketing course. In "open" markets, any IEM trader with an academic affiliation, including those enrolled in the course, may trade in the market.

LIS20L	$1.00 if <u>Lost in Space</u>'s official box office receipts for the 4/3-4/30 period are lower than or equal to $20 million; zero otherwise.
LIS30L	$1.00 if <u>Lost in Space</u>'s official box office receipts for the 4/3-4/30 period are higher than $20 million and lower than or equal to $30 million; zero otherwise.
LIS40L	$1.00 if <u>Lost in Space</u>'s official box office receipts for the 4/3-4/30 period are higher than $30 million and lower than or equal to $40 million; zero otherwise.
LIS40H	$1.00 if <u>Lost in Space</u>'s official box office receipts for the 4/3-4/30 period are higher than $40 million.

Fig. 1. *Example IEM movie market contracts.*

As part of their course, the student traders were provided with a $10 trading account that they could redeem for cash after the market liquidated. In exchange, the students were required to submit a forecast of the four-week box office performance for the two movies that would be traded on the IEM during the semester. The students were also required to execute at least two trades while the market was open. They could choose to trade in one or both of the IEM movie markets open at the time.

Market timeline

The timeline of the market is provided in Fig. 2.

A class assignment required that each student submit forecasts of the two movies' box office receipts prior to the opening of the IEM movie markets. Students' forecasts included a point estimate of the movie's box office receipts as well as an explanation of how they arrived at this forecast. More details including a sample assignment may be found in Gruca (2000).

After the forecast assignments were turned in to the experimenter, the IEM market opened. Trading in the markets began from 4–14 days before the opening of the movie in theaters (all of the movies opened on a Friday). Once the movie opened in theaters, trading continued for four weeks. Of their two required trades, students were to make one trade before and one after the movie opened in theaters.

Nielsen/EDI (entdata.com) tracks movie box office performance on a weekly basis. Daily estimates are also available at other web sites, e.g., the-numbers.com. After the final four-week receipts are available in print (through Variety), the markets are liquidated. This entails exchanging $1 for each winning contract held by a trader. Nothing is paid for losing contracts.

Overview of markets

In Table 1, we present some descriptive data about the eight prediction markets.

With respect to size and duration, these markets fall between the typical small sample, short-duration lab experiment and the large-scale, multi-month IEM political markets (Forsythe, Rietz, and Ross, 1999).

Forecasts

Our data provides two different forecasts of the four-week performance of the movie in question: a survey-based forecast and a market-based forecast. Both these forecasts are made before the movie opens. While the survey precedes the market, traders do not know the aggregate survey results when they are trading.

The survey-based forecast uses the forecasts provided in the students' written assignments as input. Since outlier forecasts can have a strong influence on the mean of the forecast distribution, we used the median of the students' forecasts in our analyses. Note that this forecast is generated before the IEM movie market opens for trading.

So that our two forecasts use comparable measures, we compute an implied point forecast from the IEM market prices. Recall that contracts in the IEM movie markets represent ranges of box office receipts. To generate the implied IEM market point forecast, we start with the midnight last trade prices for all contracts in the market. We then normalized the prices to sum to one. Finally, we interpolate to find the point forecast implied by these normalized prices. In making this

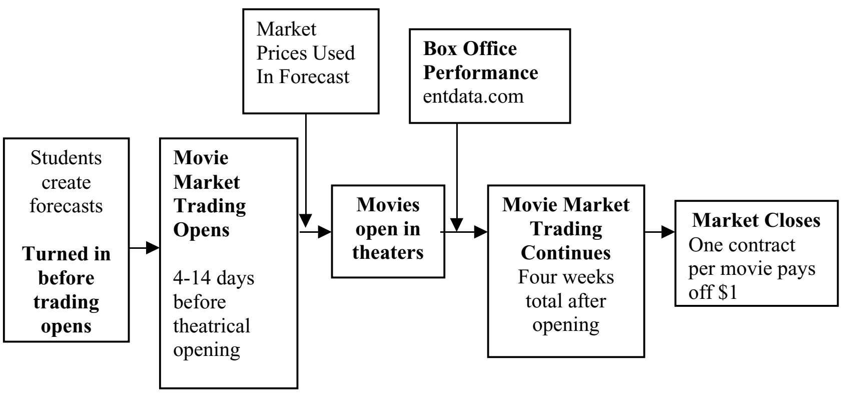

Fig. 2. *Market time line.*

Table 1. *Overview of markets*

Movie	Date market began trading	Date movie opened	Number of forecasts (Type of market)
Lost in Space	3/27/98	4/3/98	44 (open)
Mercury Rising	3/27/98	4/3/98	44 (open)
Enemy of the State	11/9/98	11/20/98	88 (closed)
I Still Know What You Did Last Summer	11/9/98	11/13/98	88 (closed)
Sleepy Hollow	11/5/99	11/19/99	106 (open)
The World is Not Enough	11/5/99	11/19/99	106 (open)
The 6th Day	11/3/00	11/17/00	91 (closed)
How the Grinch Stole Christmas	11/3/00	11/17/00	91 (closed)

interpolation, we assume a uniform distribution over the range of outcomes contained in a particular contract. This process is illustrated in Fig. 3.

Consider the contract prices for the Fall 1999 IEM market associated with the movie, "Sleepy Hollow," shown in the top panel of Fig. 3. The median (the point with 0.50 cumulative distribution) falls within the range of the SH70L contract, so it is somewhere between $50 and $70 million. Using the assumption that point outcomes are uniformly distributed over the range of the contract, we compute an implied median forecast of $53.1 million.[6]

When the median of the distribution lies in the upper-most contract, we must make an assumption about the range of that contract before we can interpolate to find the median. For instance, consider the

Sleepy Hollow (Fall 1999 IEM Movie Market)

Contract (Range)	Pre-opening Price (Norm.)	Cumulative Distribution
SH30L (≤ 30)	0.099	0.099
SH50L (30+, 50)	0.356	0.456
SH70L (50+, 70)	0.280	0.735
SH90L (70+, 90)	0.240	0.975
SH90H (> 90)	0.024	1.000

Forecast = 20 * (0.5 — 0.456) / 0.28 + 50
Forecast = 53.1

Mercury Rising (Spring 1998 IEM Movie Market)

Contract (Range)	Pre-opening Price	Cumulative Distribution
MR10L (≤ 10)	0.02	0.02
MR20L (10+, 20)	0.13	0.15
MR30L (20+, 30)	0.28	0.42
MR30H (> 30)	0.58	1.00

Forecast = 10 * (0.5 — 0.42) / 0.58 + 30
Forecast = 31.4

Fig. 3. *Interpolation examples.*

1998 Spring Movie market for the movie, "Mercury Rising" shown in the lower panel of Fig. 3. Note that the median falls in the range of the MR30H contract. This contract is unbounded above. To determine the median, we assume that the size of the range of the upper-most contract is the same as the size of the range of the next lower contract. In this case, MR30L has a range of $10 million, so we assume the range of the MR30H contract is also $10 million. Applying the formula described above, we obtain a forecast of $31.4 million.

In the next section, we discuss the results of our eight markets.

Results

Effects of trading on forecasts

To determine how well the market is able to aggregate the trader's private information about the movie's performance, we plot the median of the survey forecast (x-axis) versus the median forecast implied by the market prices (y-axis). The results are presented in Fig. 4.

We see that most of the points lie near the 45-degree line indicating a close correspondence between the two forecasts. The correlation is 0.98.

This result suggests that markets are able to accurately aggregate the student's private information reported in the survey forecasts. In controlled laboratory studies of information aggregation, experimenter generated private information is explicitly given to traders. Here we see that this same aggregation capability is present in an applied field setting where the information that each subject holds is uncontrolled.

The correspondence between the median of the survey forecasts and the median market forecast does not provide any guidance about the superiority of one approach (market v. survey) over the other. If the costs of conducting a survey are comparable to operating a market, then the choice cannot be made on the basis of the central tendency of the two forecasts. To provide further insight into the benefits of each method, we next compare the variability of market forecasts to the variability of the survey forecasts.

To examine the effect of market trading on forecast variance, we need to estimate the expected prices of the contracts based on the students' forecasts. Recall that in the IEM Movie Markets, a contract's price is the probability of that contract paying off. So, to construct an implied price from the survey data, we will determine the probability of payoff implied by the survey data. We estimate the probability of a contract paying off (and hence its price) by determining the frequency of student forecasts in the range of each contract. We then normalize these frequencies to sum to one. For example, if 10% of the forecasts lie within the range of a given contract, then its expected price should be $0.10.

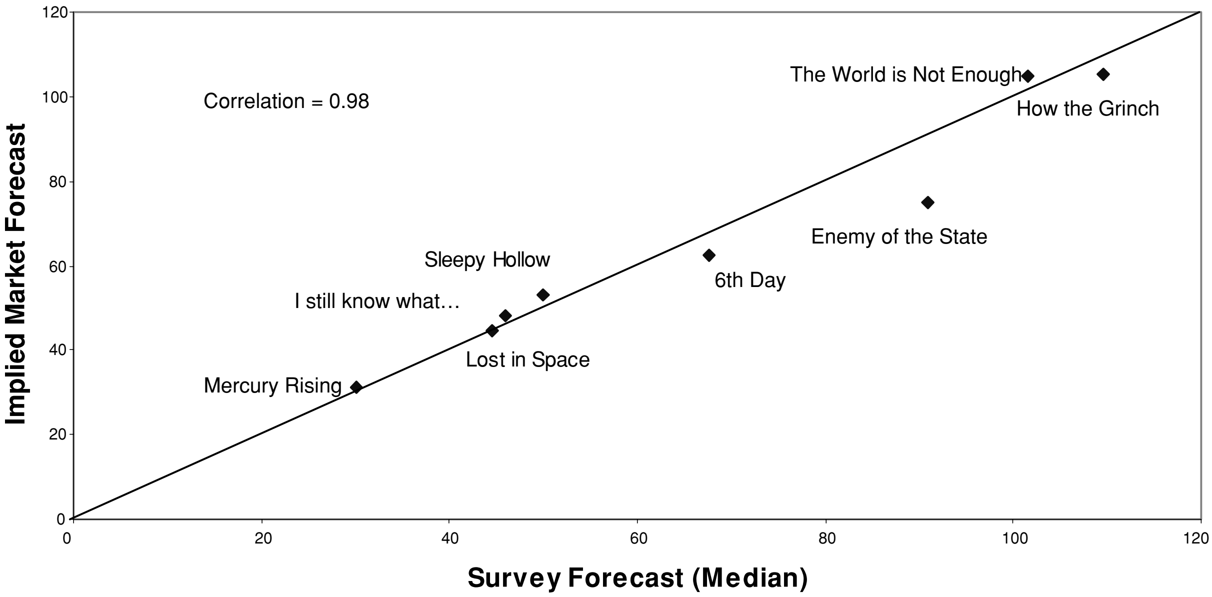

Fig. 4. *Forecast comparison.*

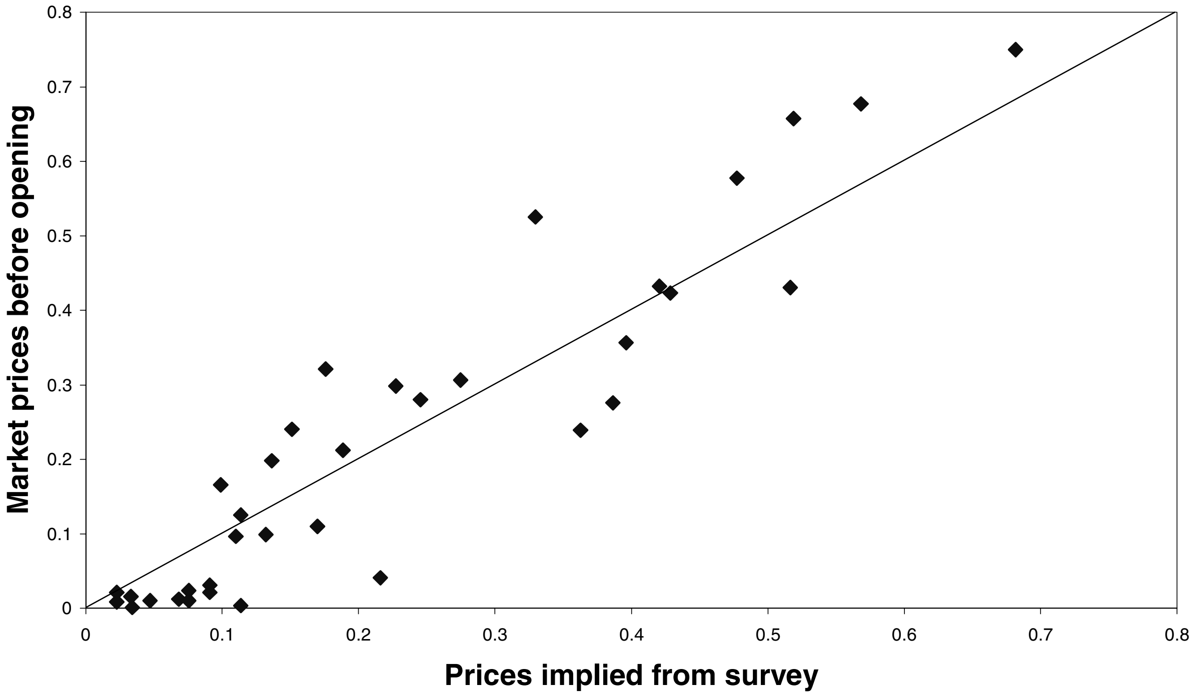

Fig. 5. *Contract prices implied from survey vs. market prices.*

In Fig. 5, we plot the expected prices based on survey forecasts on the x-axis versus the actual market prices on the y-axis. The market prices are the normalized prices from midnight the Thursday night before the movie opened in theaters.

We see that most of the contracts with expected prices lower than \$0.10 are driven lower during market trading while most of the higher-priced (>\$0.10) contracts tended to increase over that same time period. This means that the least likely outcomes, as indicated by the surveys, are considered even less likely by the market. The market appears to be solving a problem that is generally left for the researcher to solve in a survey: what to do with outlier observations. In a survey, the researcher must determine what weight, if any, to give apparent outliers, here the market has decided to drive their weights to almost zero.

We measure the variability of the two forecasts using the sum of the squared prices of all contracts. This statistic ranges from $1/N$ to 1 where N is the number of contracts available in the market. A higher sum is indicative of a lower variance with the upper limit of 1 for the case where one contract's price has risen to \$1 and the others have dropped to \$0.

The average sum of squared contract prices implied by the survey forecasts is 0.36 while the statistic is 0.43 for the forecasts implied by market prices. Comparing the results using a paired t-test, we find that the variability of the implied market forecast is significantly lower than the variability of the survey forecast ($p < 0.05$ level).

Thus, the IEM Movie Market data suggests that markets can improve survey forecasts of new product success by reducing the variance of the forecast. Our markets are able to aggregate the traders' private information in the form of the median forecast with no significant difference in the medians of the two forecasts. However, the difference in variability of the two forecasts is statistically significant, with the market appearing to "drop" outliers.[7]

Effects of trading on forecast accuracy

Trading appears to improve the forecast by reducing variance, but is the market forecast more accurate than the survey forecast? To answer this question, we have to first choose a measure of forecast accuracy.

In the IEM political markets, accuracy is measured using the absolute deviation between the normalized market prices and the actual share of the vote recorded on election day. This makes sense because both measures are bounded between 0 and 1.

In forecasting movie box office results, there is no natural notion of "share." Instead, outcomes are measured by actual four-week box office receipts. This means that the importance of a particular absolute error size can differ across markets. For example, a forecast that is off by $15 million for a movie that eventually makes $150 million in its first 4 weeks of release is more acceptable than the same absolute error for a movie that only earns $20 million in its first 4 weeks of release. Therefore, we use a relative measure of accuracy: absolute percentage error (APE = absolutevalue[actual − forecast]/actual).[8]

We computed the mean absolute percentage error (MAPE) across all eight markets for both the survey forecasts and the market forecasts. The MAPE for the survey forecasts is 0.33. The MAPE for the market forecasts is 0.30. A paired comparison t-test showed no difference between these two measures.

Do open markets predict better?

One of the purported advantages of using markets to predict new product success is the ability to tap into the expertise of traders from around the world using the connectivity of the Internet. However, the presence of a large number of uninformed traders could also introduce noise in the market prices.

Across our eight prediction markets, we have two different types of trader populations. In closed markets, all traders are students who had submitted forecasts before the opening of the market. In open markets, other traders can joint the market (so that the students are a subset of all traders in the market). Comparing the forecast accuracy of the open and closed markets, we find that the MAPE for the closed markets was much higher (0.42) than for the open markets (0.18). These results support the conclusion that open markets do a better job of prediction than closed markets.[9]

One possible explanation lies in the students' commitment to their individual forecasts. The students have formally committed themselves to a particular forecast in the form of a written estimate. They might be less willing to learn from other traders, including other students, who may have superior knowledge. In open markets, there may be a number of traders informed only by public information provided by web sites such as boxofficeguru.com or speculative traders seeking to exploit arbitrage opportunities. Such traders may add market liquidity, allowing prices to be more efficient aggregators of information.

Yet another explanation is self-selection. In closed markets, all traders were students recruited as part of a classroom exercise and each was provided with a $10 trading account. On the other hand, non-student traders joining the open markets put their own money at risk. We expect that such traders would not participate in the market unless they believed that they had superior information. Furthermore, we suspect that the requirement that traders invest real money in the market (as opposed to fictional money as is used in some Internet games) is an important factor in the efficiency of market prices. Consequently, these self-selected traders may have been instrumental in correctly moving prices in the direction of the actual outcome.

Unfortunately, our data regarding the impact of the trader pool can only be regarded as suggestive. Our limited number of observations and our experimental design prevents us from making a statistically valid conclusion. While our data suggest that open markets are more accurate than closed markets, it may also be the case that some movies are inherently more predictable than others. Because open and closed markets were not run simultaneously on the same movies, we cannot eliminate the hypothesis that it is this inherent predictability of a particular movie that is driving our result. The difference in accuracy as a function of market openness remains an interesting and important area for future research.

Does the market improve on the publicly available information?

In U.S.-based IEM political markets, traders have access to polling information that should be helpful in predicting the eventual outcome of the election (Kou and Sobel, 2001). There is a parallel source of information available to traders in the IEM Movie Box Office Markets, an Internet game known as the Hollywood Stock Exchange (HSX.com). In this popular simulation, players buy and sell "movie stocks" which pay owners in fictitious "Hollywood $" based on the four weekend box office performance of a given movie. There are incentives for the top performing players who receive HSX-theme merchandise based on overall performance over a given time period (monthly, quarterly, etc.). While there are some 400,000 registered players, only a small proportion of these buy or sell a particular movie stock. A recent report by Pennock et al. (2000) claims that the resulting prices in the market are good predictors of the movies' actual performance in the market place. For further details, see www.HSX.com.

While the time frames for the IEM and HSX differ slightly (four-week total versus four-weekend total), traders in the IEM should be able to use the public information from HSX to improve their own trading performance and, consequently, to improve the forecasting performance of the market.

We examined the overall forecasting performance of the HSX movie stocks associated with the same movies used in the IEM markets. There is a high correlation (0.81) between the percentage errors of the HSX and IEM forecasts. The MAPE for the eight HSX movie stocks (using Thursday night prices) is 0.31. Since the MAPE of the IEM market forecasts is 0.30, we find no substantial difference in forecasting accuracy between the HSX and IEM.

This may seem surprising given the results from IEM political markets where markets are superior to other sources of public information. While both HSX and the IEM are market mechanisms, the IEM uses real money and the majority of its movie market traders should be well informed (student traders were required to research the movies and submit forecasts before trading). These should be excellent conditions for improving on the forecasts from the HSX game. Here, again, our small number of observations could be a limitation.

There are a number of other possible explanations rooted in the differences between political polls and the information provided by the HSX.com web site. HSX.com, like the IEM, is an asset market, albeit with non-monetary payoffs. And, both markets are public to one another. Therefore, we might expect that the aggregation of information from traders at work in the IEM is also at work in the HSX.com game.

Another explanation is that the HSX.com game asks players the right question. HSX.com traders know that their ultimate payoffs are tied to future box office receipts. This differs from pre-election polls where questions are of the form, "If the election were held today, for which candidate A or B would you vote?" The poll asks respondents to provide an assessment of their own preferences at the current moment, not an assessment of other's actions at a future point in time.

Conclusions

In this paper, we extend field experiments of real money prediction markets to the problem of forecasting the success of a new product. Using the IEM, which has been so successful in the past in predicting the outcomes of political elections, we forecasted the four-week box office receipts for eight new movies being released in U.S. theaters.

These field experiments differ in many ways from traditional laboratory experiments. The outcome of the IEM movie markets is unknown to the researchers. In addition, each trader brings his or her own information to the market rather than having it provided by the experimenter. These conditions more closely approximate those associated with new product forecasting.

While our field experiments relax the assumptions of perfect or complete trader information, our results suggest that markets successfully aggregate traders' private information. This is illustrated by the high correlation between the survey-based forecasts (from the students' initial forecasting assignment) and market-based forecasts. In addition, we provide some evidence that the market reduces the variance of the forecasts relative to surveys.

While our markets appear to improve forecasts by reducing variance, we do not find evidence of the "crystal ball equilibrium" in the movie markets. Traders are not able to improve their forecasting performance relative to a publicly available signal from the Internet game HSX.com. While there could be an effect that our small sample size prevents us from detecting, it may also be the case that structural differences between polls and the HSX mechanism account for the difference. Given this very different result between this set of markets and those in the political forecasting realms, more research on how and why markets work is needed before they can replace traditional marketing research methods.

Acknowledgments

This research is funded in part by NSF Grant No. 9952362. We would like to thank Bob Forsythe for his comments on an earlier version of this paper.

Notes

1. Asset markets differ from traditional goods markets with respect to the role of participants. In asset markets, participants can be both buyers and sellers. Therefore, they are referred to as traders.
2. This is similar to the system used on the NASDAQ exchange.

3. Readers interested in share markets should read Berg et al. (2000) for a summary of those markets.
4. Or the exchange of ($1 ∗ share) for each of the contracts in a share market.
5. We report only those results from markets in which we collected forecasts from traders before trading began. We also conducted 4 markets without forecasts which are not included in this paper.
6. The implied median is $\$50$ million $+$ $\$20$ million $\times$ $(0.044/0.280) = \$53.1$ million, where 0.044/0.280 is the interpolation factor assuming each point outcome in the contract range is equally likely.
7. Whether this difference is due strictly to market forces is an open issue. To prevent contamination, we collected the survey forecasts before the market opened for trading. It is possible that this small difference in time could have provided more information to rule out the likelihood of extremely unusual outcomes.
8. This is analogous to the use of return in measuring stock performance in finance research.
9. A more diverse set of traders in open markets could have several different effects. They may have greater dispersion of information and, therefore, a greater propensity to trade. There may be more liquidity traders and, therefore, greater potential for information dissemination among all traders. There may be better-informed traders in this group. We cannot distinguish among these alternatives in our markets.

References

Berg J, Forsythe R, Nelson F, Rietz T. Results from a decade of election futures markets research. In: Plott C, Smith V, eds. *Handbook of Experimental Results*, Elsevier Press, 2000.

Berg J, Forsythe R, Rietz T. What makes markets predict well? Evidence from the Iowa electronic markets. In: Albers W et al., eds. *Understanding Strategic Interaction: Essays in Honor of Reinhard Selten*, New York: Springer, 1997:444–463.

Churchill GA, Iacobucci D. *Marketing Research: Methodological Foundations, 8th edn*. Harcourt Publishing: New York, 2001.

Eliashberg J, Sawhney MS. Modeling goes to Hollywood: Predicting individual differences in movie enjoyment. *Management Science* 1994;40(Sept.):1151–1173.

Forsythe R, Lundholm R. Information aggregation in an experimental market. *Econometrica* 1990;58(2):309–347.

Forsythe R, Nelson F, Neuman GR, Wright J. Anatomy of a political stock market. *American Economic Review* 1992;82:1142–1161.

Forsythe R, Rietz TA, Ross TW. Wishes, expectation and action: A survey of price formation in election stock markets. *Journal of Economic Behavior and Organization* 1999;39:83–110.

Grossman SJ. An introduction to the theory of rational expectations under asymmetric information. *Review of Economic Studies* 1981;48:541–559.

Gruca TS. The IEM movie box office market: Integrating marketing and finance using electronic markets. *Journal of Marketing Education* 2000;22(1):5–14.

Kou SG, Sobel ME. Forecasting the vote: An analytical comparison of election markets and public opinion polls. Working Paper, Columbia University, 2001.

Krider RE, Weinberg CB. Competitive dynamics and the introduction of new products: The motion picture timing game. *Journal of Marketing Research* 1998;35:1–15.

Lucas RE. Expectations and the neutrality of money. *Journal of Economic Theory* 1972;4:103–124.

Lundholm RJ. What affects the efficiency of the market? Some answers from the laboratory. *Accounting Review* 1991;66:468–515.

Neelamegham R, Chintagunta P. Forecasting the performance of new product launches in international markets. *Marketing Science* 1999;18:115–136.

Pennock DM, Lawrence S, Giles CL, Nielsen FA. The power of play: Efficiency and forecast accuracy of web market games. NEC Research Institute Technical Report, 2000-168, 2000.

Plott CR, Sunder S. Efficiency of experimental security markets with insider information: An application of rational expectations models. *Journal of Political Economy* 1982;90:663–698.

Plott CR, Sunder S. Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica* 1988;56:1085–1118.

Rosenstone SJ. *Forecasting Presidential Elections*. Yale University Press, 1983.

Sawhney MS, Eliashberg J. A parsimonious model for forecasting box-office revenues of motion pictures. *Marketing Science* 1996;15(Spring):113–131.

Urban GL, Hauser JR. *Design and Marketing of New Products, 2nd edn*. Prentice-Hall, 1993.

Zufryden F. Linking advertising to box office performance of new film releases—A marketing planning model. *Journal of Advertising Research* 1996;36(July–Aug.):29–43.

Thomas S. Gruca is the Lloyd J. and Thelma W. Palmer Research Fellow and Associate Professor of Marketing in the Tippie College of Business at the University of Iowa. His research interests include marketing and strategy. His current research projects involve electronic prediction markets, the effect of customer satisfaction on financial performance and internet pricing. He is a member of the editorial board of *Marketing Letters* and reviewer for a number of management science journals. He received his Ph.D. and M.B.A. from the University of Illinois at Champaign-Urbana.

Michael Cipriano is currently an Assistant Professor in the Department of Business Administration at Roanoke College in Salem, Virginia where he teaches accounting. His research interests involve issues in market efficiency and effectiveness using experimental and archival methods, as well as issues related to disclosure quality in financial reports. He is a Certified Public Accountant in several states. He earned an MBA from the University of Iowa where he also spent four years as a doctoral student in accounting.