

Internet-Based Virtual Stock Markets for Business Forecasting

Martin Spann • Bernd Skiera

*School of Business and Economics, Johann Wolfgang Goethe-University Frankfurt am Main,
Mertonstrasse 17, 60054 Frankfurt am Main, Germany
spann@spann.de • skiera@skiera.de*

The application of Internet-based virtual stock markets (VSMs) is an additional approach that can be used to predict short- and medium-term market developments. The basic concept involves bringing a group of participants together via the Internet and allowing them to trade shares of virtual stocks. These stocks represent a bet on the outcome of future market situations, and their value depends on the realization of these market situations. In this process, a VSM elicits and aggregates the assessments of its participants concerning future market developments. The aim of this article is to evaluate the potential use and the different design possibilities as well as the forecast accuracy and performance of VSMs compared to expert predictions for their application to business forecasting. After introducing the basic idea of using VSMs for business forecasting, we discuss the different design possibilities for such VSMs. In three real-world applications, we analyze the feasibility and forecast accuracy of VSMs, compare the performance of VSMs to expert predictions, and propose a new validity test for VSM forecasts. Finally, we draw conclusions and provide suggestions for future research.

(Internet; Forecasting; Virtual Stock Market; Design of Virtual Stock Markets; Movies; Virtual Markets)

1. Introduction

The prediction of future market developments is an ongoing challenge for business forecasting. Inaccurate or delayed predictions concerning future sales or consumer behavior can result in substantial costs for a company, and may weaken its position in negotiations. Many of these forecasting problems occur on a short- and medium-term basis. In particular, predictions concerning the success of new products create substantial difficulties, for example, the prelaunch uncertainty regarding the success of the Xbox, a game station launched by Microsoft at the end of 2001. In addition, predicting the demand for products with short product life cycles, such as movies, is problematic (e.g., Sawhney and Eliashberg 1996) as well as the prediction of sales in unstable market situations.

Apart from decision-support systems that usually make use of many different data sources, two

different approaches attempt to solve such prediction problems (Armstrong 2001). The first approach is to derive a prediction from existing data via univariate or multivariate statistical methods (e.g., extrapolation models or econometric models). The second approach is to collect new data by asking consumers or experts to evaluate different product alternatives or future market scenarios. This approach frequently leads to survey-type market research studies (e.g., Conjoint Analysis or Delphi studies).

The prediction of future market situations based on past data requires that past data contain information about the future development (Lütkepohl 1993). However, this might not be the case when new products are introduced or market conditions are generally unstable. Therefore, consumer and expert surveys have been applied in these situations. However, a representative sample of consumers might be hard

to draw and consumers tend to have difficulties in expressing their preferences for fairly new products. Experts, on the other hand, are difficult to identify and the researcher still faces the problem of how to weight the various expert opinions (see, e.g., Ashton and Ashton 1985, Batchelor and Dua 1995). In addition, all of these efforts require substantial time to complete, making it difficult to measure in real time the effects of particular information on the topic to be predicted.

The application of Internet-based virtual stock markets (VSMs) is an additional approach that can be used to predict short- and medium-term market developments. Its basic idea is to bring a group of participants together via the Internet and let them trade shares of virtual stocks. These stocks represent a bet on the outcome of future market situations, and their value depends on the realization of these market situations. Once the outcome of a specific market situation is known, each share of virtual stock receives a cash dividend (payoff) according to that specific market outcome (e.g., \$1 for each unit sold). Such types of VSMs were first applied in the form of a political stock market to predict the outcome of the Bush versus Dukakis 1988 U.S. Presidential Election with participation restricted to members of the University of Iowa community (Forsythe et al. 1992). The results of these studies are encouraging, as they demonstrate that the predictions derived from political stock markets outperform opinion polls in terms of forecast accuracy (for a recent overview, see Forsythe et al. 1999). Furthermore, the results of political stock markets show that VSMs can well perform even if their participants are not a representative sample of the electorate (or consumers on the market to be predicted). The reason is that a VSM elicits the participants' assessments of the election result (or the market outcome), and a rational participant should not trade according to his or her individual preferences, but according to his or her prediction of the market outcome due to the overall preferences of all voters (or all market participants) (Forsythe et al. 1992).

Given these positive results from the Iowa political forecasts, we would like to apply the concept of VSMs to solve short- and medium-term business forecasting

problems. These business forecasting problems differ from those taken into consideration by political stock markets (i.e., elections) in several aspects. First, the market situations involved in business forecasting are usually much more complex. Instead of only having two possible outcomes—that is, the winner is either candidate A or candidate B, or predicting vote shares of different political parties—a prediction may involve entirely different products whose sales figures need to be estimated. Second, the available information to solve business forecasting problems is usually poorly accessible, whereas political stock markets are conducted in situations with readily available information. Third, the application of VSMs for business forecasting problems requires a well-designed incentive structure, because the desired experts might be reluctant to participate and unwilling to invest their own time and money in VSMs. Fourth, predictions related to business forecasting may be required much more frequently to solve recurring problems, for example, sales figures often need to be predicted at least on a quarterly basis.

The use of the Internet to conduct such VSMs provides additional advantages. Participants can conveniently access the VSM from almost anywhere in the world at any time. Thereby, the Internet increases the pool of possible participants and allows, if necessary, the anonymity of traders on the VSM.

Should the use of VSMs lead to accurate forecasts, such a method would offer a number of advantages compared to traditional forecasting methods. First, it allows for an almost real-time reaction of stock prices to additional information and, hence, a quick prediction of the impact of that information on future market situations. Second, it does not burden the researcher with the task of weighting and aggregating different expert opinions as this is achieved by the trading mechanism implemented in the VSM. Participants, for example, weight their assessments by the volume and the price of the purchase or sale order they place or accept. Third, once established, a VSM can operate at rather moderate operating costs. Fourth, a VSM provides participants with an incentive to reveal their true assessments (Forsythe et al. 1999) if an adequate remuneration is properly linked to the participants' performance on the VSM. Hence,

whereas many consumer surveys remunerate consumers for their participation in a survey, a VSM usually remunerates participants for their successful participation. Wertebroch and Skiera (2002) show, for example, that consumers' willingness to pay significantly differs according to the incentive structure being provided. Finally, participants in a VSM might have more fun than their counterparts partaking in consumer or expert surveys.

Hence, the aim of this paper is to evaluate the potential use of VSMs for their application to business forecasting, discuss the different design possibilities, analyze the forecast accuracy, compare the performance of the VSM to expert predictions, and propose a new validity test for VSM forecasts. In §2, we describe the fundamental concepts of a VSM in more detail and outline the theoretical framework underlying why the trading mechanism embedded in a VSM should lead to accurate forecasts. Section 3 explores the various design possibilities of a VSM. Sections 4 and 5 describe and summarize the design and application of a VSM in three real-world settings, analyze the feasibility and forecast accuracy of VSMs, compare the performance of the VSM to expert predictions, and perform a new validity test for VSM forecasts. Section 6 contains conclusions and provides suggestions for future research.

2. Fundamental Concepts and Theoretical Foundations of Virtual Stock Markets

The fundamental concept of applying VSMs for business forecasting is to make future market situations expressible and tradable through virtual stocks. Thereby, the cash dividend (payoff) of such shares of virtual stocks depends on the market outcome; that is, the actual outcome of a specified event at time T

$$d_{i,T} = \phi(Z_{i,T}) \quad (i \in I), \quad (1)$$

where

$d_{i,T}$ = cash dividend of the stock modeling the outcome of the i th event at time T ,
 $\phi(\bullet)$ = transformation function,
 $Z_{i,T}$ = outcome of the i th event at time T ,

I = index set of events,

T = point or period in time that is relevant for the determination of the outcome of the event.

T is usually predetermined (e.g., the end of the election at political stock markets). However, it could also be a random variable, if the goal of the VSM is either to predict the point in time when a specific event occurs (e.g., when will an airline file for Chapter 11 or when will a hurricane hit the United States) or to predict an outcome whose specific time of occurrence is not known (e.g., which airline will file for Chapter 11 next or which state will be hit by a hurricane next).

The transformation function might have different forms and needs to be invertible. The one most commonly used in political stock markets is to pay a cash dividend of \$1 multiplied by the fraction of votes the particular candidate receives (Forsythe et al. 1999). An alternative transformation function might pay a cash dividend of \$1 if the candidate is elected and \$0 otherwise. Equation (1) also points out that we are actually not dealing with stocks, but with securities whose terminal values are contingent upon the outcome of an uncertain event. However, in accordance with the literature dealing with political stock markets (e.g., Forsythe et al. 1992, 1999), we use the denomination "stocks" to make the concept easier for the participants to understand.

The price of one share of a virtual stock should correspond to the VSM's aggregate expectations of the event outcome and, therefore, to the discounted expected cash dividend of a share of the stock. Discounting, however, might be negligible if the total duration of the VSM is rather short with the value of $T - t$ converging to zero:

$$\begin{aligned} \widehat{Z}_{i,T,t} &= \phi^{-1}(\widehat{d}_{i,T,t}) \\ &= \phi^{-1}(p_{i,T,t} \cdot (1 + \delta)^{T-t}) \quad (i \in I, t < T), \quad (2) \end{aligned}$$

where

$\widehat{Z}_{i,T,t}$ = expected value of the outcome of the i th event at the t th point in time T ,
 $\phi^{-1}(\bullet)$ = inverse of the transformation function,
 $\widehat{d}_{i,T,t}$ = expected value of cash dividend of the stock at the t th point in time for the i th event at time T ,

$p_{i,T,t}$ = price of the one share of stock at the t th point in time for the i th event at time T ,
 δ = discounting factor (assumed to be constant).

Participants of the VSM use their (individual) assessment of the market outcome to derive an (individual) expectation of the cash dividend of the related share of virtual stock. Accordingly, they compare their expected cash dividend with the VSM's aggregate expectation $\widehat{Z}_{i,T,t}$, which is a function of the stock price $p_{i,T,t}$. Thereby, they trade their individual assessment. For example, a participant of the VSM anticipates sales of 100 units during the next month for a specific product. The cash dividend of the related share of virtual stock would, consequently, be \$100 in the case that each unit of sales would correspond to \$1. In the case of a current price $p_{i,T,t}$ of \$95 (\$105); that is, $\widehat{Z}_{i,T,t} = 95$ (105) units, this stock would be undervalued (overvalued) according to the estimates of this participant. This individual could try to attain an expected profit of \$5 via buying (selling). If the potential gains in the virtual portfolio value create a sufficiently high incentive for the participants to perform well on the VSM, it is then their best strategy to make transactions according to their individual assessments. In this way, the participants reveal their true assessments of future market situations via their purchases and sales of stocks.

The theoretical foundations for the use of prices in VSMs for forecasting future market developments are the efficient market hypothesis and the Hayek hypothesis (Fama 1970, 1991; Hayek 1945). The efficient market hypothesis is a state description: A market is efficient if all available information is always fully reflected in the prices (Fama 1970). Therefore, if a VSM is efficient, the price of a specific stock reflects all information on the corresponding future market state and, thus, serves as a forecast. The rationale for why the trading mechanism of a VSM should aggregate and display individual assessments in stock prices; that is, be efficient, is based upon the so-called Hayek hypothesis. It states that the price mechanism on a competitive market is the most efficient instrument to aggregate the asymmetrically dispersed information of market participants (Hayek 1945, Smith 1982). Because individual assessments are tradable via

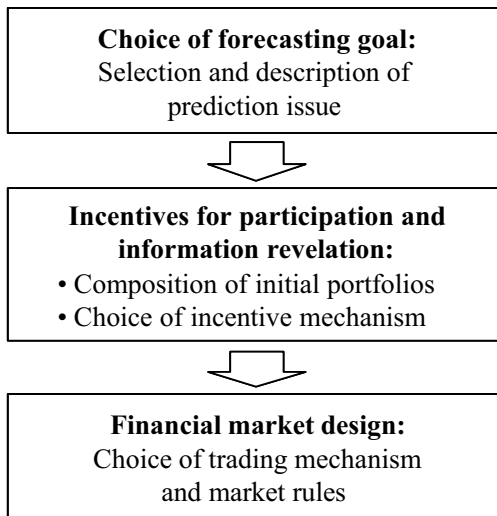
shares of virtual stocks in a VSM, VSMs create a market for the prediction of a future market situation in which the participants compete on the basis of their individual assessments. Thus, their aggregated information should be reflected in the stock prices. Extensive studies using empirical data and laboratory experiments support this informational efficiency of financial markets (see the overviews of Fama 1970, 1991; Sunder 1995). The rather good results of political stock markets strengthen this notion.

The ability of VSMs to reflect participants' information in the stock prices depends upon several prerequisites. There has to be a way to make future market situations tradable via shares of virtual stock. That means that $Z_{i,T}$ of (1) is unambiguously the basis for the cash dividend of the respective shares of stock; that is, the determination of the cash dividend is clear for participants. Additionally, the participants need to have some knowledge about the future market situation, otherwise, stock prices would be randomly set, and markets could not economize on information as stated by the Hayek Hypothesis. Furthermore, the incentive created by the payoff mechanism has to be high enough so experts decide to participate and reveal their true valuations.

3. Design of Virtual Stock Markets

Although political stock markets have been repeatedly applied (for a review, see Forsythe et al. 1999), they all basically used the same design, i.e., the one originally proposed by Forsythe et al. (1992). That design required the participants to invest a moderate amount of real money and used so-called unit portfolios to avoid any financial risk for the initiator of the VSM. Apparently, this specific design seems to be sufficient to handle the requirements of political stock markets. Business forecasting problems, however, have some unique characteristics (see §1) that might require other designs of a VSM. Therefore, we want to outline a broader range of design possibilities for VSMs by deriving conclusions from the results of different areas of research, in particular, political stock markets (Forsythe et al. 1992, 1999), experimental economics (Smith 1976, Friedman and Sunder 1994, Davis and Holt 1993), and financial market

Figure 1 Steps for Designing a Virtual Stock Market



design (e.g., Madhavan 1992, Harris and Hasbrouck 1996). Thereby, we separated the different issues to be taken into consideration into three categories (see Figure 1).

First, the forecasting goal of a VSM determines which future event or market state is to be predicted and is, subsequently, modeled as a specific stock to be traded. Second, participants need to have an incentive to reveal their true assessments. This requires the design of the initial portfolio and the final remuneration that together make up the incentive mechanism. The third category represents the design of the actual trading on the VSM itself; that is, the selection of market and trading rules. Table 1 provides an overview of the different design possibilities for VSMs according to these categories, which will be discussed in this section.

3.1. Choice of Forecasting Goal

The forecasting goal represents the forecasting problem to be predicted. Possible prediction issues corresponding to specified events $Z_{i,T}$ could be (i) the prediction of an absolute number, for example, sales in a particular period; (ii) the prediction of a relative number, for example, market share in a particular period; or (iii) the occurrence or nonoccurrence of a particular event, for example, the completion of a development project at a particular point in time.

The appropriate formulation for the payoff function in (1) in cases (i) and (ii) could be the following linear transformation:

$$d_{i,T} = \rho \cdot Z_{i,T} \quad \text{with } \rho > 0 \quad (i \in I), \quad (3)$$

where (e.g.)

ρ = transformation parameter, e.g.,

$$\begin{cases} 0.001 \cdot \text{unit sales} & \text{in case (i),} \\ 100 \cdot \text{market share} & \text{in case (ii).} \end{cases}$$

Case (iii) typically requires the formulation for the payoff function in (1) according to (4):

$$d_{i,T} = \begin{cases} v & \text{if particular event occurred,} \\ 0 & \text{otherwise,} \end{cases} \quad (i \in I), \quad (4)$$

where

v = cash dividend value, e.g., \$1.

An even more flexible approach could use a multi-part step function to account for different outcome levels in case that the event occurs, e.g., a payoff equal to \$1 if unit sales are above 100, a payoff equal to \$0.50 for unit sales above 50 and below 100, and a payoff of \$0 for unit sales below 50. The duration (if T in (1) is predetermined) and accessibility of the VSM must be determined when choosing the forecasting goal. A VSM could start as soon as participants are able to individually assess the prediction issue, and would have to finish before the outcome of the event is revealed. A VSM could, therefore, last several months or just a few hours. Participation in the VSM might be open to everyone at any time, or restricted to selected groups (e.g., in cases of confidential company data) during specific time periods. Our empirical results (shown later) and that from political stock markets (Forsythe et al. 1999) indicate that VSMs work for different payoff functions as long as they are comprehensible for participants, for different durations of VSMs and for unrestricted or restricted participation.

3.2. Incentives for Participation and Information Revelation

To create incentives for experts to participate and reveal their assessments of future market developments, a potential reward should be established

Table 1 Decisions Concerning the Design of a Virtual Stock Market

Step	Decisions
Choice of forecasting goal	<ul style="list-style-type: none"> Selection of the prediction issue Formulation for the payoff function Duration of the virtual stock market Open to the public or limited participation (specific groups or time periods)
Incentives for participation and information revelation	<ul style="list-style-type: none"> Composition of initial portfolios/endowment Investment of real money or endowment If endowment: Shares and/or virtual money Provision of loans and alternative investment opportunities (and choice of interest rates) Remuneration/incentive mechanism <ul style="list-style-type: none"> Monetary versus nonmonetary rewards Linear relationship between performance and reward versus nonlinear relationship (i.e., tournament) <ul style="list-style-type: none"> If linear relationship <ul style="list-style-type: none"> — Zero-sum game or not — If zero-sum game: “unit portfolios” or “jackpot” If tournament <ul style="list-style-type: none"> — Individual performance level or rank-order tournament — Predetermined prize ranks or not (e.g., “jackpot”) — Number and value of prize ranks — Time intervals Incentives not based on performance <ul style="list-style-type: none"> — Conditional: Random rewards among specific group — Unconditional: Random rewards among all participants
Financial market design	<ul style="list-style-type: none"> Trading mechanism: <ul style="list-style-type: none"> — Market maker (i.e., dealer) <ul style="list-style-type: none"> ◇ Choice of initial quotes ◇ Quote adjustment to order flow — Double auction (open or closed order book) — Combination of types Trading hours Long and/or short trading Order types: Limit and/or market; possible temporal restrictions Position limits and price limits Trading fees or no trading fees

according to the participants' performances; that is, their portfolio values *after* the cash dividends of the shares of virtual stock have been determined at the end of the VSM (Smith 1976, Friedman and Sunder 1994). Basically, the initiator of the VSM has the choice between two alternatives. The first alternative is to require participants to invest their own money and to link their monetary reward to their performance on the VSM. Thereby, the idea is that if participants' own money is at stake, this creates a strong

incentive for them to perform well. However, some experts might refuse to participate because they are unwilling to accept the risk of a financial loss. In addition, playing with real money might be considered to be illegal gambling. Furthermore, additional accounting and administration costs occur for managing the investments and the payoffs. Therefore, a second alternative is to waive a required investment, provide participants with an endowment of shares of virtual stocks and/or virtual money instead, and to

reward all or the best performing participants with monetary or nonmonetary rewards.

The investment of the participant's own money is usually required at political stock markets (e.g., Forsythe et al. 1999), whereas an endowment is common in microeconomic and financial market experiments (e.g., Smith 1976, Sunder 1995). However, if participants invest their own money and receive real cash payoffs as a reward for their final portfolio values, the initiator of a VSM faces the risk of losing an uncertain amount of money if the sum of the cash dividends of the stocks is unknown before the end of a VSM. To avoid this risk, a VSM has to be modeled as a zero-sum game. A zero-sum game can be achieved via two different approaches. First, so-called "unit portfolios" accomplish this task and have been frequently applied in political stock markets (Forsythe et al. 1999). In this case, the set of all stocks in the unit portfolio reflects all possible market outcomes so that (5) holds true and participants can only buy (and sell) such unit portfolios at the price of v :

$$\sum_{i \in I} d_{i,T} = v. \quad (5)$$

Therefore, the value of a unit portfolio and its cash-out value are always v irrespective of the specific cash dividends for the different stocks. As a result, the initiator faces no financial risk, yet is still able to offer each participant an incentive to perform well.

Such zero-sum games are easy to implement for political stock markets because either one of the candidates will win or the sum of the vote shares of all candidates or parties will be 100%, i.e., (5) holds true. Business forecasting, however, frequently requires predicting absolute numbers (e.g., sales figures) so that the sum of the cash dividends of stocks might considerably vary. Thus, real cash payoffs are risky if they are proportionally linked to each participant's final portfolio value. In this case, an initiator of a VSM might employ a second approach to create a zero-sum game by rewarding participants according to the participant's final portfolio value relative to the sum of all participants' final portfolio values. Thereby, the sum of the initial investments equals the sum of the

rewards, thus, treating the individual investments of all participants as a "jackpot"¹

$$R_{j,T} = \theta \cdot V_{j,T} \quad \text{with } \theta = \frac{\sum_{j \in J} F_{j,0}}{\sum_{j \in J} V_{j,T}} \quad (j \in J), \quad (6)$$

where

$R_{j,T}$ = reward of the j th participant at time T ,

$V_{j,T}$ = portfolio value in virtual currency units of the j th participant at time T ,

θ = transformation parameter of virtual currency units into real cash,

$F_{j,0}$ = initial financial investment of the j th participant,

J = index set of participants.

However, one disadvantage of this design might be that participants do not like or fully understand the reward function (6) because the parameter value θ is unknown until the end of the VSM.

An alternative to this linear transformation of final portfolio values into financial rewards is a tournament remuneration mechanism where the participant's remuneration depends upon the participant's performance relative to others; that is, his or her rank (see, e.g., Rosen 1986, Lazear and Rosen 1981). The initiator of the VSM can limit his or her financial risk by committing to a minimum number of prizes and by determining the total number of prize ranks and the reward associated with each prize rank after the number of participants and, thus, the sum of the initial financial investments are known.

Another alternative is the use of unit portfolios where stocks would correspond to intervals so that predictions are not made for absolute numbers, but for the presence of an absolute number within a specific interval. An example might be a unit portfolio consisting of three stocks with a payoff of \$1 for each stock according to (4) if (i) sales are below \$1 million, (ii) between \$1 million and \$2 million, or (iii) above \$2 million. Then, the price of each stock represents the probability for the absolute number to be within the corresponding interval. The major drawback of this approach is that the prediction cannot be made

¹ The authors thank the associate editor for drawing their attention to this specific design.

for an absolute number itself, and that the payoff function is rather inflexible because all possible outcomes need to be accounted for (i.e., (5) needs to be fulfilled).

An endowment avoids these problems that are associated with requiring participants to invest real money. Thereby, each participant usually receives the same endowment that can consist of virtual money and/or shares of virtual stocks. Such an endowment allows high flexibility in terms of the reward function, which might provide all or only some participants with monetary or nonmonetary rewards. Microeconomic and financial market experiments usually employ an endowment and provide real cash payouts proportional to final portfolio values (e.g., Smith 1976, Sunder 1995). However, this design entails risky payoffs for the initiator of a VSM, with total costs considerably depending upon the number of participants. Although this design is applicable in laboratory experiments with a small and limited number of participants, it is risky for VSMs with unrestricted participation. Thereby, the initiator's financial risk might be limited by using one of the methods outlined above or, as we have done, by using nonlinear remuneration mechanisms such as a rank-order tournament. That way, the VSM initiators have to bear the costs of the rewards, but do not face any financial risk.

The design of the incentive mechanism can also have implications according to behavioral decision research. Participants endowed with a certain portfolio may have a tendency to stick to the initial endowment ("status quo bias") or their willingness to accept greatly exceeds their willingness to pay ("endowment effect") (Kahneman et al. 1991). However, if the traded stocks can be redeemed for cash, the endowment effect may not appear (Kahneman et al. 1990). An endowment may also lead to an increase in risky trading behavior (Thaler and Johnson 1990). On the other hand, if participants invest their own money, loss aversion may lead to a fairly conservative trading behavior.

In addition to incentives based on performance (i.e., an increase in portfolio value), nonperformance-based incentives may be used to make the participation in a VSM more attractive. Such rewards can either be

randomly given away, for example, by drawing a lottery, or based on a lottery among participants who fulfill certain conditions (e.g., having at least a portfolio value above a certain level, being among the top 100 participants or trading at least 10 times) so that participants have the incentive to show at least some performance or some trading behavior even if they are not doing well on the VSM. In addition, if VSMs are repeatedly applied for a recurring prediction issue (e.g., for quarterly sales forecasts), it might be beneficial to link the reward to the performance across several VSMs to create an incentive for participants to perform well on all of these VSMs.

A final conclusion is difficult to draw. Our empirical studies and the political stock markets show that both, either the investment of participants' own money and real money cashouts according to final portfolio values or participants' endowment in combination with rank-order tournaments and nonmonetary prizes, lead to good results. If participants' own money is at stake, they may have a higher incentive to perform well. However, well-designed rank-order tournaments might attract experts that are not willing to invest their own money, and might also provide participants with sufficiently high incentives. In addition, it avoids the burden of administering all the individual investments and payoffs.

3.3. Financial Market Design

Finally, the financial market design has to be defined, which entails the definition of the trading mechanism and market rules. The two most prominent trading mechanisms include a market maker (dealer) trading mechanism comparable to NASDAQ, and a double auction trading mechanism, comparable to the one applied at the New York Stock Exchange (NYSE). The market maker (dealer) accepts every purchase or sale order from a participant and executes it according to the purchase and sale prices (quotes) determined by the market maker (Bernstein 1987). If a market-maker trading mechanism is chosen, the initial quotes and how the quotes adjust to the order flow have to be determined. This adjustment mechanism can be either automatic (e.g., a price increase [decrease] of $x\%$ for every y shares of the i th stock bought [sold]) or performed by human dealers. A double auction directly matches the purchase and sale orders of participants

(Gjerstad and Dickhaut 1998). Thereby, both buyer and seller determine the price of matching a purchase and a sale order. If a double auction trading mechanism is chosen, it must be determined whether the order book of the VSM should be open or closed (i.e., whether or not information on unfulfilled orders is disclosed). A market-maker type of trading mechanism has the advantage of providing a permanent trading partner for participants and, thus, a higher liquidity than a double auction type of trading mechanism. However, if real money investments are present at the VSM, the market-making activity may result in a suboptimal portfolio structure that poses a financial risk for the market maker.

In addition, several decisions regarding the market rules must be made. Specific limits on portfolio positions can be imposed at a VSM (e.g., a maximum of y shares of the i th stock or $x\%$ of total portfolio value invested in the i th stock) and maximum and minimum prices for limit orders and quotes. The purpose of such position limits is to restrict the possible influence of a single trader on the price of a specific stock. It is important to note that imposing a trading fee should harm forecast accuracy, because it inhibits rational participants from executing trades with an expected profit of less than the trading fee, thus, creating a bias. In this case, a rational participant who expects a certain cash dividend will only buy (sell) shares if the expected cost—that is, the share price plus trading fee (revenue, i.e., share price minus trading fee)—of a trade is less (more) than the expected cash dividend of the shares of this trade. A comparable argument applies to fixed interest rate deposits as alternative investments, because participants might prefer these to a risky trade with an anticipated low profit. The VSMs in our studies employed market-maker and double auction market mechanisms, and both mechanisms achieved a good forecast accuracy in most cases. Political stock markets usually applied a double auction market mechanism that also yielded good results. Empirically, both trading mechanisms have not been directly compared for VSMs with respect to their forecast accuracy. Therefore, we conclude that the initiator has the choice between those two trading mechanisms as long as participants do not trade with real money. In case of the latter, a

double auction market mechanism does not require the initiator of the VSM to participate in the trading of stocks, whereas a market-maker mechanism forces the initiator to buy and sell stocks at his or her own financial risk.

4. Empirical Application: Forecasting U.S. Movie Gross Box-Office Revenues

The aim of the three empirical studies described and summarized in this paper is to analyze different applications of a VSM for business forecasting. The following applications differ with respect to the number of participants, the accessibility, the trading mechanism of the VSM, the data sources, and the amount of information available.

The goal of the first study is to analyze the empirical application of a VSM for predicting the U.S. box-office success of movies prior to their release. The release of a movie entails the introduction of a new product that is confronted with a high uncertainty of the corresponding demand, a significant failure rate and high financial stakes for production and marketing (e.g., Sawhney and Eliashberg 1996, Eliashberg et al. 2000). Furthermore, the prerelease prediction of the box-office success of a movie is considered to be a difficult and important task for distributors and exhibitors who have to determine the number of movie screens, for example, and related promotional activities (Sawhney and Eliashberg 1996, Eliashberg et al. 2000). Prerelease information on movies is normally heterogeneously dispersed among the possible participants. Each participant can evaluate a movie, for example, by visiting the movie's website and downloading a trailer, asking friends and family members about their opinion and intention to see the movie, or read movie critiques. In addition to analyzing the forecast accuracy of the results and its comparison to expert predictions, we attempt to increase the knowledge about the possibility to expand the simple point estimation of a VSM to derive prediction intervals and propose a new validity test for VSM forecasts.

4.1. Design Description of the Virtual Stock Market

The Hollywood Stock Exchange (HSX) is a VSM where participants can buy and sell shares of virtual movie stocks (www.HSX.com). One goal of the HSX is to use the data collected from the exchange as a market research instrument. The HSX has had no distinct end of trading since stock trade started in 1996. New stocks are continuously issued (with so-called IPOs) and stocks are constantly cashed out with payoffs. In our analysis, we focus on the data of 152 movies released between January 1, 2000 and May 25, 2001, and their related virtual stocks ("MovieStocks"). The HSX has more than 725,000 registered participants

with an average of more than 15,000 unique participants actively trading every day. See Table 2 for details on the design of the study.

The HSX issues MovieStocks by making this type of stock available for trade for an initial purchase and sale price (HSX refers to this as an "IPO"). The IPO of MovieStocks is usually some time ahead of the actual release of the respective movie, often more than a year in advance. A MovieStock is cashed out and removed from the exchange four weeks after the respective movie is released on more than 650 screens at the box office. The payoff function (the cash-out price) is equal to the movie's domestic (U.S.) box-office revenue total during those first four weeks of release, paying one

Table 2 Design of the Empirical Study Hollywood Stock Exchange

Step	Decisions
Choice of forecasting goal	Forecasting of movie gross box-office revenues in the United States Payoff function: Gross box-office revenues in the first four weeks of release in the United States: (7) Price adjustment according to the box-office revenues of the opening weekend: (8) No downtime since start of HSX in 1996—new stocks are continuously issued and stocks are constantly cashed out Open to the public, participants can join at any time
Incentives for participation and information revelation	Composition of Initial Portfolios/Endowment Endowment of two million "Hollywood Dollars (H\$)" per participant Alternative investment opportunities: Unspent cash pays annual interest rate of 6% Remuneration/Incentive Mechanism Nonmonetary rewards during certain periods (e.g., May 2000 to January 2001) Rank-order tournament: During reward periods rewards for participants with highest increase of (virtual) portfolio value during specific interval Time intervals: Weekly, monthly, per season to date, year to date, 1996 to date Incentives not based on performance: Additional sweepstakes (e.g., quizzes) not directly related to VSM
Financial market design	Trading mechanism: Market maker — Choice of initial quotes: "IPO Prices" — Quote adjustment to order flow: "Virtual Specialist" Trading times: 24 hours a day, 7 days a week Long and short trading Order types: Limit and market; with temporal restrictions possible Position limits: Maximum of 50,000 shares per type of movie stock One percent trading fee on the total value of a trade

“Hollywood Dollar (H\$)” for every million of box-office revenue:

$$d_{m, T_2} = \frac{Z_{m, T_2}}{1,000,000} \quad (m \in M), \quad (7)$$

where

Z_{m, T_2} = domestic (U.S.) box-office total during the first four weeks of release of the m th movie (T_2 = first four weeks of release),

M = index set of movies.

On Friday, the day of the release of a movie, trading on the respective MovieStock is stopped. On Sunday evening of the opening weekend, the price of this MovieStock is adjusted and trading commences. The price is adjusted according to the Sunday evening estimation of the box-office receipts of the opening weekend multiplied by a specific factor (see (8)). On Sunday evening, the actual box-office receipts for Sunday are not yet available so that an estimate is used based on the results from Friday and Saturday. Its purpose is to adjust the price according to a simple prediction (i.e., opening weekend box-office receipts multiplied by a factor) for the cash-out price. The reason for this adjustment is to incorporate the information of the box-office revenues of the opening weekend in the quotes provided by the market maker (“Virtual Specialist”) at HSX. This factor of 2.9 for regular opening weekends (Friday to Sunday) is publicly known (the factor for extended opening weekends are adjusted differently; that is, 2.5 for a four-day holiday weekend and 2.3 for a five-day holiday weekend):

$$a_{m, T_1} = \frac{\tilde{Z}_{m, T_1}}{1,000,000} \cdot \alpha \quad (m \in M), \quad (8)$$

where

a_{m, T_1} = adjusted price of the m th movie after the opening weekend (T_1 = opening weekend),

\tilde{Z}_{m, T_1} = estimated opening weekend box-office receipts of the m th movie (T_1 = opening weekend),

α = adjustment factor (2.9 for a regular three-day opening weekend).

Until trading is stopped on the Friday of the opening weekend, the stock price of a MovieStock represents a prediction of the price adjustment according to the expected opening weekend box-office receipts

and is, thereby, a prediction of the opening weekend box-office receipts (see (9)). Contrary to our design recommendation in §3, a 1% trading fee on the total value of a trade is charged. There are currently no rewards provided on HSX. However, participants with the highest increase of (virtual) portfolio value during specific tournament intervals are ranked and displayed, most likely creating some intrinsic motivation for participants:

$$\hat{a}_{m, T_1, t} = p_{m, T_1, t} \Rightarrow \hat{\tilde{Z}}_{m, T_1, t} = \frac{p_{m, T_1, t}}{\alpha} \cdot 1,000,000 \quad (m \in M, t < T_1), \quad (9)$$

where

$\hat{a}_{m, T_1, t}$ = expected adjusted price at the t th point in time for the m th movie after the opening weekend (T_1 = opening weekend),

$p_{m, T_1, t}$ = price of a share at the t th point in time for the m th movie stock (T_1 = opening weekend),

$\hat{\tilde{Z}}_{m, T_1, t}$ = expectation at the t th point in time on the estimated opening weekend box-office receipts for the m th movie (T_1 = opening weekend).

4.2. Forecast Accuracy

Our analysis of the forecast accuracy of HSX is conducted according to the prerelease prediction for the opening weekend according to (9). We compare the closing price before trading is stopped on Friday, the first day of the opening weekend, with the price adjustment according to the expected opening weekend box-office receipts on Sunday evening. We calculate the Mean Absolute Percentage Error (MAPE) of deviation between the Friday closing prices and the Sunday evening price adjustment for all 152 movies in our sample to be 30.96% (see (10)). To compare this result, Sawhney and Eliashberg (1996) calculated a MAPE of 71.1% in their model for the prediction of total box-office revenues for 10 movies without box-office data; that is, before the movie release:

$$\text{MAPE} = \frac{\sum_{m \in M} (|p_{m, T_1, t} - a_{m, T_1}| / a_{m, T_1})}{|M|}, \quad (10)$$

where

MAPE = Mean Absolute Percentage Error of the deviation between the Friday closing prices and Sunday evening price adjustment of the M MovieStocks (t = Friday close),

$|M|$ = number of different types of MovieStocks in index set M .

In addition to the measurement of the forecast error, we evaluate whether market prices at a VSM are efficient or systematically biased, thus, whether all available information is reflected in the prices. Therefore, we propose a combined validity test. The first test analyzes whether the VSM forecasts are systematically biased, and the second test examines whether incorporation of additional information (e.g., attributes of the products to be predicted) can significantly reduce the sum of squared errors in a regression analysis. This joint validity test is based on a regression with the actual value of the market situation to be predicted as the dependent variable and the VSM forecast of this market situation and additional information, e.g., specific attributes of this market situation, as independent variables.

In case of the HSX, the following regression model (11) tests whether HSX provides an unbiased estimate of the Sunday evening price adjustment that cannot be improved by taking movie specific factors into account. Exogenous factors for the success of a movie on an opening weekend can be its genre and the intensity of its distribution, i.e., the number of movie screens on which a movie is shown during opening weekend:

$$a_{m, T_1} = \beta_0 + \beta_1 \cdot \hat{a}_{m, T_1, \text{HSX}} + \beta_2 \cdot \text{TS}_m + \beta_3 \cdot \text{DVActThr}_m + \beta_4 \cdot \text{DVDraRo}_m + \beta_5 \cdot \text{DVCom}_m + \mu_m \quad (m \in M), \quad (11)$$

where

$\hat{a}_{m, T_1, \text{HSX}}$ = HSX forecast of the Sunday evening price adjustment of the m th movie according to its Friday closing price (T_1 = opening weekend),

TS_m = number of movie screens of m th movie on opening weekend (in thousands),

DVActThr_m = dummy variable: "1" if movie is in genre "action" or "thriller,"

DVDraRo_m = dummy variable: "1" if movie is in genre "drama" or "romance,"

DVCom_m = dummy variable: "1" if movie is in genre "comedy,"

μ_m = residual of the m th movie.

The HSX forecasts are unbiased if the constant (β_0) of the regression is equal to 0 and the parameter for the HSX forecasts (β_1) is equal to 1. The joint hypothesis test ($H_1: \beta_0 = 0, \beta_1 = 1$) is accomplished by testing the restriction H_1 through application of restricted least squares (Greene 2002). Further, the HSX forecasts incorporate additional information if movie-specific factors cannot significantly reduce the sum of squared errors. Thus, the corresponding hypothesis ($H_2: \beta_2 = \beta_3 = \beta_4 = \beta_5 = 0$) can be analogously tested through application of restricted least squares.

According to the regression results of Table 3, the possible effects of movie genre and intensity of distribution are already incorporated in the HSX forecasts based on the Friday closing price, because the respective parameters are insignificant in the original model. Thereby, application of restricted least squares yields that the hypothesis H_2 cannot be rejected ($F = 0.604, p = 0.660$). However, the HSX forecasts apparently

Table 3 Regression Results for Movie-Specific Factors and Hollywood Stock Exchange Forecasts

Parameters	Original model ^a	Enhanced model 1 ^a	Enhanced model 2 ^{a,b}
Constant	-12.017 (0.059)	-3.342 (0.078)	-3.376 (0.181)
Friday closing price (HSX)	1.043 (0.000)	1.081 (0.000)	1.087 (0.000)
Movie screens in thousands	3.636 (0.201)		
DV_Action_Thriller	0.735 (0.836)		
DV_Drama_Romance	2.774 (0.425)		
DV_Comedy	2.978 (0.404)		
R^2	0.862	0.859	0.861
F -value	182.065 (0.000)	917.601 (0.000)	458.819 (0.000)
Number of observations	152	152	76

^a p -value in parentheses.

^b Estimation for sample of first 76 movies.

display a bias toward underestimating the opening weekend box-office revenues for “large” movies (i.e., movies that generate high box-office revenues) and overestimating those revenues for “small” movies if the respecified model (enhanced model 1) is estimated. This bias is displayed by the negative constant β_0 , indicating that a downward adjustment increases the model’s explanatory power (which has the largest relative effect in case of small movies) while, on the other hand, the value of the parameter β_1 greater than one reflects that an upward adjustment is effective for large movies. However, application of restricted least squares yields that the hypothesis H_1 cannot be rejected for the enhanced model 1 at the 5% level ($F = 2.607$, $p = 0.077$). Further, we estimated the respecified model for an estimation sample of the first 76 movies (enhanced model 2) to allow for out-of-sample predictions when applying the model to enhance forecasts (see §4.3). Thereby, the estimation of the enhanced model 2 yields a negative constant β_0 and a value of the parameter β_1 greater than one, too, but the hypothesis H_1 cannot be rejected ($F = 1.485$, $p = 0.233$).

Based on the regression model (13), the application of standard econometric procedures allows prediction intervals for the Sunday evening price adjustment for a specific movie to be created (Gujarati 1995):

$$\Pr[\hat{a}_{m, T_1, OLS} - t_{\alpha/2} \cdot \text{se}(\hat{a}_{m, T_1, OLS}) \leq a_{m, T_1} \leq \hat{a}_{m, T_1, OLS} + t_{\alpha/2} \cdot \text{se}(\hat{a}_{m, T_1, OLS})] = 1 - \alpha \quad (m \in M), \quad (12)$$

$$\hat{a}_{m, T_1, OLS} = \beta_0 + \beta_1 \cdot \hat{a}_{m, T_1, HSX} \quad (m \in M), \quad (13)$$

where

$\hat{a}_{m, T_1, OLS}$ = estimate of Sunday evening price adjustment of the m th movie according to estimation of regression model (13) (T_1 = opening weekend),

$\Pr[\bullet]$ = probability,

$t_{\alpha/2}$ = critical value of t -distribution at $\alpha/2$ level of significance,

$\text{se}(\bullet)$ = standard error.

Thus, based on the Friday closing price, the estimated price adjustment according to the regression model (13) will lie with a probability of $(1 - \alpha)$ in the prediction interval of (12). Equation (12) and the

assumption of a normal distribution for the Sunday evening price adjustment of a particular movie allows a prior for movies’ opening weekend box-office success to be constructed (e.g., O’Hagan 1998, Kass and Wasserman 1996) by taking into account that the total amount of probability outside an interval of two standard deviations of a normal distribution is only 0.0456:

$$a_{m, T_1} \sim N(\hat{a}_{m, T_1, OLS}, \frac{1}{2} \cdot t_{0.05/2} \cdot \text{se}(\hat{a}_{m, T_1, OLS})) \quad (m \in M). \quad (14)$$

Another interesting aspect to take into consideration is that the 1% trading fee may have harmed the forecast accuracy. The same holds true for the secure alternative investment, an interest rate of 6% on unspent cash. The HSX forecast accuracy might, therefore, have been even better without the trading fee and secure alternative investment.

4.3. Performance Compared to Expert Judgments

To evaluate the performance of HSX, we compare the results of the HSX predictions to two renowned expert predictions for movie box-office revenues available on the Internet. Further, we add the results of the calculation of enhanced predictions of HSX according to models 1 and 2 in Table 3 to our comparisons. The first expert prediction is Box Office Mojo (BOM) by Brandon Gray who writes a widely read weekly column and provides forecasts on his website (www.boxofficemojo.com) for each weekend dating from mid-1998 to the present. Because BOM does not provide forecasts for several weekly periods, a total of only 140 movies could be used for a comparison of predictions between Brandon Gray and HSX. The second expert opinion comes from Box Office Report (BOR), which has published forecasts for each weekend as of March 2, 2001 on its website (www.boxofficereport.com). The HSX, BOM, and BOR have 24 movies in common, which can be used for a comparison. Table 4 displays the results of a direct comparison of all different prediction methods. It is interesting to note that although the overall MAPE for BOM is the lowest, HSX achieves twice as many hits. A “hit” is defined as a case where a specific forecasting instrument (HSX or an expert) has

Table 4 Comparison Between the Predictions of Hollywood Stock Exchange, Box Office Report, and Box Office Mojo*

Instrument	Hits**	MAPE (%)	HSX % improvement (<i>p</i> -value) ^a
HSX	14	40.62	—
HSX enhanced model 1 ^b	14	36.48	11.35% (0.080)
HSX enhanced model 2 ^c	14	36.81	10.35% (0.100)
Box Office Report (BOR)	3/2/2 ^d	53.40	-23.93% (0.049)
Box Office Mojo (BOM)	7/8/8 ^d	35.30	15.07% (0.485)

^a Percentage of improvement of alternative instrument over HSX = [MAPE HSX – MAPE instrument]/MAPE instrument (two-tailed paired *t*-test for difference).

^b Model estimated on movies 1–152.

^c Model estimated on movies 1–76.

^d Number of hits in comparison to HSX/HSX enhanced model 1/HSX enhanced model 2.

* Comparison of predictions for the 24 movies whose opening weekend box-office revenues were predicted by HSX, BOR, and BOM between March 2, 2001, and May 25, 2001 (two weeks without predictions from BOM).

** Number of movies. Method with lowest absolute percentage error for a specific movie.

the lowest absolute percentage error of all considered forecasting instruments for the same movie. The enhanced predictions of HSX according to models 1 and 2 lead to an increase in forecast accuracy relative to the ones directly derived from the Friday closing prices at HSX.

The comparison of HSX and BOM (see Table 5) shows that BOM performed slightly better than HSX (71 versus 69 hits and an insignificant difference of the MAPE). However, the enhanced model 1 of HSX significantly increases forecast accuracy of HSX. This

increase, however, does not necessarily need to occur because the validity test analyzes the effect on the sum of *squared* errors, but the MAPE represents the mean *absolute* percentage error.

5. Summary of Additional Empirical Applications

Given space limitations, we provide a brief summary of two additional VSM applications below. Please see the online Appendix (mansci.pubs.informs.org/ecompanion.html) for details on these applications.

Table 5 Comparison Between the Predictions of Hollywood Stock Exchange and Box Office Mojo*

Instrument	Sample	Hits**	MAPE (%)	HSX % improvement (<i>p</i> -value) ^a
HSX	140 movies	69	31.11	
HSX enhanced model 1 ^b	140 movies	75	28.40	9.54% (0.001)
Box Office Mojo (BOM)	140 movies	71/65 ^d	28.05	10.91% (0.152)
HSX	64 movies ^f	36	34.76	
HSX enhanced model 2 ^c	64 movies ^f	36	32.71	6.27% (0.077)
Box Office Mojo (BOM)	64 movies ^f	28/28 ^e	32.88	5.72% (0.601)

^a Percentage of improvement of alternative instrument over HSX = [MAPE HSX – MAPE instrument]/MAPE instrument (two-tailed paired *t*-test for difference).

^b Model estimated on movies 1–152.

^c Model estimated on movies 1–76.

^d Number of hits in comparison to HSX/HSX enhanced model 1.

^e Number of hits in comparison to HSX/HSX enhanced model 2.

^f Movies 77–152: movie predictions of HSX and BOM for 64 movies.

* Comparison of predictions for 140 movies whose opening weekend box-office revenues were predicted by HSX and BOM between January 1, 2000 and May 25, 2001 (four weeks without predictions from BOM).

** Number of movies. Method with lowest absolute percentage error for a specific movie.

5.1. Forecasting Movie Visitors and Pop Music Single-Chart Positions in Germany

In a second empirical study, called "Chart- & Movie-eXXchange" (CMXX) (www.CMXX.com), we analyzed the forecast accuracy and the performance of a VSM in an additional application for the prediction of the box-office success (number of visitors) of movies in Germany and for the chart position of pop music singles in Germany. We conducted the CMXX seven times for a total number of 81 movies and 11 pop music singles, using our own VSM software that employs a double auction trading mechanism. Thereby, we yielded similar results to the analysis of the HSX.

5.2. Forecasting the Usage of Different Services of a Large German Mobile Phone Operator

In another study, we used our own VSM software at a large German mobile phone operator to forecast the usage of some of its different mobile phone services in a specific month (further referred to as "MPSX" for "Mobile Phone Services eXchange"). MPSX differs from the other VSMs analyzed in this paper mainly because access was restricted to 20 selected employees in the marketing and planning departments of this company. Therefore, the goal of this study is to analyze the forecast accuracy and the performance of a VSM with restricted accessibility and a small number of participants. We conducted MPSX between May 22, 2001 and June 22, 2001. During this period, the 20 selected participants could trade five different stocks,

each one representing the expected usage of a specific mobile phone service in June 2001.

Table 6 compares the forecast accuracy of MPSX with the accuracy of forecasts based upon the arithmetic mean, the geometric mean, the linear trend, and the exponential trend of the previous five, respectively, three months of usage prior to June 2001 that were available. The forecast accuracy of the VSM MPSX is better than all of the four extrapolation models. However, due to the small sample size, none of the differences is statistically significant.

6. Conclusions and Future Research

Our empirical studies show rather encouraging results for the applicability of VSMs for business forecasting purposes. A further application of a VSM for the prediction of the German premier soccer league results also yielded a good forecast accuracy of the VSM in comparison to expert judgments and forecasts derived from betting odds (Spann and Skiera 2003). Apparently, such VSMs can be easily conducted via the Internet in terms of organizational and technical aspects. In addition, VSMs seem to work well under different incentive structures and even with a limited number of participants.

Encouraged by the results of our empirical studies, we propose a review of the forecast accuracy of VSMs for business forecasting purposes in future

Table 6 Comparison Between the Mobile Phone Services Exchange Predictions and Extrapolations

Mobile phone service	Arithmetic mean (APE) (%)	Geometric mean (APE) (%)	Linear trend (APE) (%)	Exponential trend (APE) (%)	MPSX (APE) (%)
New subscribers ^a	25.22	23.79	8.67	7.86	3.05
Unique WAP users ^a	6.33	6.51	3.12	3.39	6.35
WAP user proportion ^a	26.65	25.83	2.37	0.72	5.76
Unique GPRS connections ^b	0.20	0.11	8.60	8.70	8.60
Data volume per connection ^b	18.60	19.14	34.65	32.91	23.66
MAPE	15.40	15.08	11.48	10.72	9.48
MPSX % improvement	38.44	37.14	17.42	11.57	
(<i>p</i> -value) ^c	(0.415)	(0.418)	(0.512)	(0.660)	

^a Based upon monthly usage data for January 2001 to May 2001.

^b Based upon monthly usage data for March 2001 to May 2001.

^c Percentage of improvement of MPSX over extrapolations = [MAPE Extrapolation – MAPE MPSX]/MAPE Extrapolation (two-tailed paired *t*-test for difference).

studies. A further interesting usage could come from the area of manager appraisal. Because the forecasting of market developments is an important task for most companies, they also require this ability from their managers. Therefore, a company might link rewards to the forecast accuracy of their executives (Ashton and Ashton 1985). Consequently, manager appraisal and remuneration could be based on a manager's performance on a VSM for business forecasting problems. Thereby, managers with a good forecasting ability and market understanding could be identified and promoted to positions that require this capability. Applied with consumers as participants, for example, as in the case of HSX, VSMs can be a tool to detect knowledgeable consumers or "trend scouts," which, for example, could be consulted by companies for purposes such as testing new products. Otherwise, if experts or knowledgeable consumers are already identified, their willingness to participate in a VSM (e.g., fraction of invited experts who actually trade) can reflect their confidence in the forecasting task and, thus, be used as a predictor of the VSM's performance.

A rich field for future research is the analysis of different design alternatives for VSMs. Future research could analyze the different possibilities of a VSM design in more detail, for example, concerning the appropriate design of the incentives, the virtual stocks, the initial portfolio structure, the inclusion of possible financial losses, and the effects of alternative durations of a VSM. Additionally, the analysis of the real-time effects of particular events on the topic to be predicted (i.e., on the corresponding stock prices) can be of interest.

One possible shortcoming of VSMs is that although informational efficiency may hold true in the long term, market imperfections (e.g., "bubbles") may exist in the short term. As the stocks traded at VSMs are actually securities whose terminal values are contingent upon the outcome of an event in the short- and medium-term future, these short-term market imperfections appear less of a problem at the end of VSMs, i.e., when using the closing prices of VSMs to derive forecasts. A more detailed analysis of such market imperfections and effects recorded in behavioral decision research can be the focus of future experimental research of VSMs.

Although representative samples are not required to obtain good results, future studies might want to analyze whether improvements in forecast accuracy could be achieved if participants form a representative sample. Further, as the informational structure is an important prerequisite of VSMs, participants could be actively informed before trading on a VSM. Thereby, the specific information provided can be experimentally controlled and its effects on the forecast accuracy in separated VSMs analyzed. Furthermore, a profound analysis of the effects of providing groups of participants with more information or better tools to exploit information, for example, via the use of decision-support systems, on the resulting performance of individual participants would certainly contribute to our overall understanding of the usefulness of management science tools under different circumstances.

An attractive analysis would consider the possibilities of combining a VSM with other new data-gathering methods. It may be possible to create some sort of hybrid system. For example, a VSM could be combined with a focus group using the same set of participants. Such an approach would differ from the approach of combining (i.e., aggregating) the *results* of different forecasting techniques (e.g., Blattberg and Hoch 1990, Armstrong 2001). Two possible designs of such an approach would be the following: In the first design, a group of participants could initially trade on a VSM. After this trading, individual interviews could be conducted, followed by another round of trading on the VSM. That way, the participants might be sensitized by the VSM in so far as being able to quantify the outcome and estimate the sensitivity of the results. Another design would be to perform initial interviews, then allow trading on the VSM, and finally, finish with another round of interviews. A possible advantage might be that trading on the VSM would involve participants so much that they might think more carefully about future developments. The analysis of such a combination of methods could yield interesting results concerning the interaction between the two forecasting techniques and the possible elicitation of further information. All in all, we feel that the use of VSMs opens the doors to an interesting

area of research and could be the focus of additional studies in the future.

Acknowledgments

The authors gratefully acknowledge the provision of data from the Hollywood Stock Exchange and many helpful comments from Brian Dearth and Melissa Perrot of HSX Research, Sönke Albers, Sonja Gensler, Dominique Hanssens, Christian Laux, Frank Schmid, Gerry Tellis, Markus Zuber, and the participants of presentations at Penn State University, Stanford University, UCLA, and the Second INFORMS Marketing Science and the Internet Conference at the University of Southern California. In addition, the two editors, associate editor, and two referees provided us with many helpful suggestions.

References

- Armstrong, J. S., ed. 2001. *Principles of Forecasting*. Kluwer Academic Publishers, Dordrecht, The Netherlands.
- Ashton, A. H., R. H. Ashton. 1985. Aggregating subjective forecasts: Some empirical results. *Management Sci.* **31** 1499–1508.
- Batchelor, R., P. Dua. 1995. Forecaster diversity and the benefits of combining forecasts. *Management Sci.* **41** 68–75.
- Bernstein, P. L. 1987. Liquidity, stock markets and market makers. *Financial Management* **16** 54–62.
- Blattberg, R. C., S. Hoch. 1990. Database models and managerial intuition: 50% model and 50% managers. *Management Sci.* **36** 887–899.
- Davis, D. D., C. A. Holt. 1993. *Experimental Economics*. Princeton University Press, Princeton, NJ.
- Eliashberg, J., J.-J. Jonker, M. S. Sawhney, B. Wierenga. 2000. Moviemod: An implementable decision-support system for prerelease market evaluation of motion pictures. *Marketing Sci.* **19** 226–243.
- Fama, E. F. 1970. Efficient capital markets: A review of theory and empirical work. *J. Finance* **25** 383–417.
- . 1991. Efficient capital markets: II. *J. Finance* **46** 1575–1617.
- Forsythe, R., T. A. Rietz, T. W. Ross. 1999. Wishes, expectations and actions: A survey on price formation in election stock markets. *J. Econom. Behavior Organ.* **39** 83–110.
- , F. Nelson, G. R. Neumann, J. Wright. 1992. Anatomy of an experimental political stock market. *Amer. Econom. Rev.* **82** 1142–1161.
- Friedman, D., S. Sunder. 1994. *Experimental Methods. A Primer for Economists*. Cambridge University Press, Cambridge, U.K.
- Gjerstad, S., J. Dickhaut. 1998. Price formation in double auctions. *Games Econom. Behavior* **22** 1–29.
- Greene, W. H. 2002. *Econometric Analysis*. Prentice-Hall, Upper Saddle River, NJ.
- Gujarati, D. N. 1995. *Basic Econometrics*. McGraw-Hill, Singapore.
- Harris, L., J. Hasbrouck. 1996. Market vs. limit orders: The super-DOT evidence on order submission strategy. *J. Financial Quant. Anal.* **31** 213–231.
- Kahneman, D., J. L. Knetsch, R. H. Thaler. 1990. Experimental tests of the endowment effect and the coase theorem. *J. Political Econom.* **98** 1325–1348.
- , ———, ———. 1991. Anomalies: The endowment effect, loss aversion, and status quo bias. *J. Econom. Perspectives* **5** 193–206.
- Kass, R. E., L. Wasserman. 1996. The selection of prior distributions by formal rules. *J. Amer. Statist. Assoc.* **91** 1343–1370.
- Lazear, E. P., S. Rosen. 1981. Rank-order tournaments as optimum labor contracts. *J. Political Econom.* **89** 841–864.
- Lütkepohl, H. 1993. *Introduction to Multiple Times Series Analysis*. Springer, Berlin, Germany.
- Madhavan, A. 1992. Trading mechanisms in securities markets. *J. Finance* **47** 607–641.
- O'Hagan, A. 1998. Eliciting expert beliefs in substantial practical applications. *Statistician* **41** 21–35.
- Rosen, S. 1986. Prizes and incentives in elimination tournaments. *Amer. Econom. Rev.* **76** 701–715.
- Sawhney, M. S., J. Eliashberg. 1996. A parsimonious model for forecasting gross box-office revenues of motion pictures. *Marketing Sci.* **15** 113–131.
- Smith, V. L. 1976. Experimental economics: Induced value theory. *Amer. Econom. Rev.* **66** 274–279.
- . 1982. Microeconomic systems as an experimental science. *Amer. Econom. Rev.* **72** 923–955.
- Spann, M., B. Skiera. 2003. Predicting soccer results: Virtual stock market vs. betting odds and experts. Working paper, Goethe University, Frankfurt am Main, Germany.
- Sunder, S. 1995. Experimental asset markets: A survey. J. H. Kagel, A. E. Roth, eds. *Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ, 445–500.
- Thaler, R. H., E. J. Johnson. 1990. Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Management Sci.* **36** 643–660.
- von Hayek, F. A. 1945. The use of knowledge in society. *Amer. Econom. Rev.* **35** 519–530.
- Wertebroch, K., B. Skiera. 2002. Measuring consumer willingness to pay at the point of purchase. *J. Marketing Res.* **39** 228–241.

Accepted by Arthur M. Geoffrion and Ramayya Krishnan; received August 2001. This paper was with the authors 18 months for 5 revisions.