Anatomy of an Experimental Political Stock Market

By ROBERT FORSYTHE, FORREST NELSON, GEORGE R. NEUMANN, AND JACK WRIGHT*

Results from the Iowa Political Stock Market are analyzed to ascertain how well markets work as aggregators of information. We find that the market worked extremely well, dominating opinion polls in forecasting the outcome of the 1988 presidential election, even though traders in the market exhibited substantial amounts of judgment biases. Our explanation is that judgment bias refers to average behavior, while in markets it is marginal traders who influence price. We present evidence that in this market a sufficient number of traders were free of judgment bias so that the market was able to work well. (*JEL* A10, C90, G10)

Markets allocate scarce resources to their most valued use. To most economists this statement is almost self-evident, subject to standard conditions about market failure; but to some social scientists, economists among them, the claim is viewed with suspicion, and in some cases, it is denied. Such denial is usually a claim that markets do not work well because participants do not behave as economic theory assumes. Particularly in markets that require probabilistic calculations and forecasts of future outcomes the information demands of efficient operation of markets may be unattainable. Research in psychology and political science, for example, is replete with examples in which individuals do not, or perhaps even cannot, perceive circumstances objectively. The existence of judgment bias, wish fulfillment, and other failures in rational information processing raises serious questions about how and why markets achieve any degree of efficiency at all.

Our prime concern in this paper is to analyze the operation of a market that worked extremely well, despite failures in the information-processing capabilities of some of its participants. We analyze data from the Iowa Presidential Stock Market (IPSM), which was designed and implemented in 1988 to yield predictions of the expected vote shares of the presidential candidates in that fall's election.[1] Prices in that market fully reflected the available information about the election in the sense that the market's prediction of President Bush's margin of victory was off by less than one percentage point.

There are several reasons why this market is interesting. First, its finite duration coupled with payoffs determined by an observable outcome allows us to ask how well the market worked. Did the prices generated in this market predict the outcome of the election? Also, there are alternative

*Department of Economics, University of Iowa, Iowa City, IA 52242; and, for the last author, Department of Political Science, University of Iowa, Iowa City, IA 52242. Support of the College of Business, University of Iowa; the Interdisciplinary Research Grant Program, University of Iowa; and National Science Foundation grants SES-90-10667, SES-90-22587, and SES-91-22253 is gratefully acknowledged. We have benefited from comments made by Doug Bernheim, Jim Hines, John Kennan, Alan Krueger, and seminar participants at Aarhus, Iowa, and Princeton. Four anonymous referees made several suggestions that were valuable. All remaining errors are ours. The data used in this paper can be obtained on diskette or from FTP by contacting any of the authors.

[1]The idea of using markets to predict presidential outcomes is not original. We are aware of related markets run at the University of Rochester, California Institute of Technology, Princeton University, the Wharton School of the University of Pennsylvania, and the Brookings Institution during the 1988 campaign. The IPSM, however, appears to be the first and only such market operated strictly as a research venture.

means of generating forecasts of election outcomes (namely, public opinion polls), and these serve as a comparative standard to judge market performance. Second, by placing this market in the untraditional area of political science we can study the effect of judgment biases on market performance. Existing political-science research provides strong evidence that ideology colors perceptions of reality and hence may affect individual trading behavior. Finally, the data collected in this market allow for the detailed analysis of the behavior of traders. In full-scale financial markets even the most finely detailed available data consist merely of anonymous prices and quantities of separate trades, making investigation of the behavior and influence of individual traders impossible. In this field experiment, however, every action taken is dated, identified by trader codes, and recorded in an extensive data trail, allowing for a microscopic examination of every event on the market. In addition, extensive surveys of traders were conducted to allow an analysis of the correlation between trading behavior and demographic and preference information.

Our findings lend support to the so-called Hayek hypothesis (Vernon L. Smith, 1982), which asserts that markets can work correctly even if the participants have very limited knowledge about their environment or about other participants. How markets can work under such circumstances was never clearly specified by Hayek, and the hypothesis is not easily tested. Further, existing data do not permit any deeper insight into how and why a market serves in this capacity. The virtue of the data from the IPSM is that it permits us to analyze the behavior of individual traders in the market and to show that the efficiency of the market depended not on the average trader, but on what we call the "marginal trader," a trader relatively free of judgment bias who consistently bought and sold at prices very close to the equilibrium price.

The remainder of this paper is organized as follows. In Section I we describe the organization and operation of the IPSM, and in Section II we examine how well the market did in predicting the 1988 election, both in absolute terms and relative to opinion polls. The third section of the paper describes what is meant by the term "judgment bias" and shows how it is related to the excess-demand function for an asset. We find substantial evidence of judgment bias, precisely in the manner that political scientists (e.g., Lee Sigelman and Carol K. Sigelman, 1984; Carole J. Uhlander and Bernard Grofman, 1986) have claimed, yet the market worked quite well. To explain this, we examine the behavior of marginal traders and study how their trading behavior affected prices. The final section contains our conclusions.

I. The Iowa Presidential Stock Market

A. *Market Description*

The IPSM was initiated in April 1988, and the market opened for trading on June 1.[2] Aspiring traders were sold portfolios of shares in candidates at $2.50 each, with each basic portfolio consisting of one share in each major candidate in the campaign. The slate of candidates included George Bush, Michael Dukakis, Jesse Jackson, and a candidate labeled "Rest-of-Field."[3] Shares were given value by the dividends paid after the election, with the dividend on each share determined as the candidate's fraction of the popular vote times $2.50. Since Rest-of-Field covered all third-party candidates who earned votes in the election, the vote shares summed to 1 across the four candidates, and the total dividend paid on a basic portfolio of one share in each candidate just matched the fee charged for that

[2]Specific details of the IPSM can be found in Forsythe et al. (1991a); here we summarize only the essential aspects.

[3]Though it was never invoked, the contingency for the admission of new candidates to the slate was an integral part of the market description and may have been a factor in the Rest-of-Field pricing decisions by traders. Had a new major candidate emerged, say from a brokered convention or the premature withdrawal of a major-party nominee, each Rest-of-Field share would have been split into a continuing Rest-of-Field share and a new new-named candidate share.

portfolio. This investment/payoff rule was adopted for the IPSM because it provides a direct translation of market prices into estimates of vote shares,

$$\text{expected vote share} = \text{price}/\$2.50$$

and thus offers a prediction of not only the election winner, but also the margin of victory.[4]

The IPSM operated as a computerized double-auction market. Upon payment of a portfolio investment fee and posting of a deposit to a cash account, each trader was assigned an individual computer account that enabled him to trade in the market. The market operated around the clock from June 1 until 9:00 A.M. on November 9, 1988, the day after the election. For each $2.50, a participant received a basic portfolio consisting of one share in each candidate on the slate. Since separate markets operated in each of the candidates, investors unbundled their purchases, trading each component separately. In addition to the portfolio investment, each new trader was required to keep a cash account. This account provided the liquidity for stock transactions; purchases of shares were charged to it, and sales of shares were credited to it. New traders were allowed to enter the market at any time, and existing traders could purchase new $2.50 portfolios (each consisting of one share in each candidate on the current slate) or make additions to their cash account at any time. In addition, conversions of cash account deposits into basic portfolios or basic portfolios into cash were allowed at any time. No withdrawals from cash accounts were allowed before the close of the market on November 9.

Participants traded in the double-auction market by issuing offers to buy (bids) or offers to sell (asks). There could be many bid and ask prices in the system at any time; they were maintained in bid and ask queues, ordered first by offer price and then by time of issuance. When an offer was entered into the bid or ask queue, it remained there until one of these events occurred: (a) it was withdrawn by the bidder, (b) it reached the top of the bid queue (bottom of the ask queue) and was found to be infeasible, as described below, or (c) it reached the top of the bid queue (bottom of the ask queue) and was subsequently matched with an opposing offer. The actual transactions were executed by the system when it found overlapping bid and ask prices in the respective queues. Note too that the system did not check the identities of the traders when overlapping bid and ask prices resulted in a trade, so that a trader could sell stock to himself.

The computerized market provided facilities for obtaining information about the trader's account and the market, as well as for issuing bids and offers. For individual traders, available account information included the number of shares held in each candidate, the balance in one's cash account, a list of outstanding offers, and a list of transactions. Market information available to an individual included the current high-bid, low-ask, and last transaction prices, as well as a record of the previous day's activity including opening and closing bid and ask prices, the last transaction price, the average transaction price, and the number of shares traded in each candidate. As on most stock markets, information on the depth of the bid and ask queues was not revealed. The computer system recorded all transactions (logins, logouts, requests for prices, bids, asks, and trades) in an audit trail for subsequent research. These data were not made available until the market closed in November.

Short sales and purchases on margin were disallowed by the system; offers to buy with insufficient funds in the cash account of the buyer or offers to sell when the seller's portfolio contained no shares in that candidate were ruled infeasible. Checks for feasibility were made only when an offer reached the top of its queue (high bid or low ask). If

[4]An all-or-nothing rule—a fixed payoff of, say, $2.50 to shares in the election winner with zero payoffs to shares in losers—is perhaps the most obvious example of the many possible alternative designs. This rule provides a prediction of only the probability of winning for each candidate, and in a runaway election, the expected convergence of prices to $2.50 and $0 with little or no subsequent price movement might well have stifled continuing interest in the market.

an offer failed the feasibility check it was withdrawn from the queue. Thus, the system accepted offers that were not feasible at the time of issuance. The intention was to enrich the set of strategies a trader might adopt while at the same time preventing the market itself from becoming a net creditor. Since only the high-bid and low-ask prices were revealed to traders, the system also avoided giving false information; traders were guaranteed that at least one share was available at quoted prices.

When a new feasible bid was entered with a price equal to or exceeding the current minimum price in the ask queue, the system recorded a trade at the ask price. Likewise, if a new ask was entered with a price equal to or less than the current maximum price in the bid queue, a trade was recorded at that bid price. Such trades were executed one share at a time, regardless of the number of shares bid or asked, with renewed checks for feasibility of both the bid and the ask after each one-share transaction. The recording of a transaction included notes in the transaction logs of the two traders involved, a credit to the cash account of the seller, a debit to the cash account of the buyer, and a transfer of the share of stock from the seller's portfolio to that of the buyer. The principles followed for the execution of trades were as follows: (a) offers to buy were processed "high-prices first"; (b) offers to sell were processed "low-prices first"; (c) in the case of ties (two offers at the same price), the earliest offer to arrive on the market was processed first; and (d) when an overlap between bid and ask prices was found, the trade was executed at the price of the older of these two offers.

In this market it is interesting to ask why people trade at all. One might view this market as a complete one, since the set of securities spans the possible outcomes of who might win. In this case, the no-trade theorem (Mark Rubinstein, 1975; Paul Milgrom and Nancy Stokey, 1982) ought to apply. Even in this instance, Hal Varian (1985, 1986) has shown that trade can occur due to differences in opinion.

An alternative view is that this market is incomplete: the payoff-relevant states are not just simply who wins but depend, in general, on a large number of factors. Even in this case, however, it is not apparent why traders participate. The rules of the market, particularly the investment/payoff rule, made the IPSM a zero-sum game in monetary rewards: all funds invested were returned in the form of dividends, and gains by one trader were exactly offset by losses of other traders. This structure raises the question of what motivated traders to participate in the IPSM. This zero-sum nature of the IPSM makes it resemble markets for futures contracts, except that an important motive for the existence of futures contracts is the transfer of risk from one party to another. In a large-scale presidential stock market there might conceivably be a transfer of risk: arms-makers might enter the market to buy shares in a pacifist candidate as a hedge against the loss of future income should that candidate win the election, for example. Such risk transfers could hardly be the motive for trading in the 1988 IPSM. That leaves five reasonable motives: a novelty factor (i.e., a trader treats this like another leisure-time activity, such as playing a computer game), confidence in one's knowledge about the election and its likely outcome relative to the knowledge of other traders (i.e., differences in information), confidence in one's ability to interpret news relative to other traders (i.e., differences in opinion), confidence in one's talents as a trader, and risk-seeking behavior. We expected these differing motivations to attract a diverse group of traders to the market. Our expectations were correct. Besides differing in their motives for participating in the market, traders varied considerably in their demographic characteristics, in their partisan and ideological preferences, and in their investments and earnings. We review some of these differences below.

B. *The Traders*

Participation in the IPSM was voluntary.[5] A total of 192 traders enrolled in the mar-

[5]Ideally, enrollments would have been unrestricted, but University of Iowa attorneys expressed concern

TABLE 1—POLITICAL CHARACTERISTICS OF TRADERS

Characteristic	Market survey (percentage)	CBS/*New York Times* survey (percentage)
Presidential preference among those planning to vote:		
Bush	50	46
Dukakis	41	40
Other	2	0
Undecided	7	14
Party identification:		
Republican	46	31
Democrat	32	36
Independent	22	33
Other	1	0
Registered to vote	80	67

Note: The CBS/*New York Times* poll was conducted during 21–23 September. See *The New York Times*, September 25, 1988.

ket; those traders held a total of 1,462 shares in each candidate. The total investment, including share portfolios and cash reserves, was $4,967, yielding an average individual investment of roughly $25. The smallest investment was $5 ($N = 7$), and the largest was $420 ($N = 1$).

During the third week of September, we asked existing traders to complete a brief mail survey to provide descriptive information about their political preferences and demographic characteristics. Subsequently, new traders were asked to complete the survey when they enrolled in the market. The full text of the survey form is included in Forsythe et al. (1991a). Using the data from traders registered as of September 27, Table 1 compares the traders' presidential preferences, partisan orientations, and frequency of voter registration with comparable figures from a CBS/*New York Times* national poll conducted during the same week. Because the CBS/*New York Times* data pertain only to likely voters, we utilize responses from those traders who stated that they were likely to vote in the election.

IPSM traders were slightly more supportive of George Bush, more Republican, and less independent in the partisan leanings. Except for the undecided voters, preferences of the IPSM traders were well within the margin of sampling error for the national poll. However, in direct contrast with the national tendency in partisanship, the traders professed a noticeably higher affinity for the Republican party. Moreover, the traders were clearly less independent than voters nationally. The traders also appeared to be more politically active than the average citizen, with 80 percent claiming to be registered to vote as compared with only 67 percent in the national sample. That IPSM traders tended to be relatively conservative and politically active is consistent with the demographic makeup of the group. In general, the traders were predominantly male, white, well-educated, and among the middle and upper income categories. Students comprised the largest category of traders, and the largest share of these students were business majors. One-third of the traders had at least a college degree, more than 90

over the possible violation of laws governing either or both gambling and stock exchanges. Cost aside, the time required to obtain necessary exemption from state and federal securities laws would have thwarted plans to conduct the market during the 1988 election year. Thus, we opted to operate under the provisions of Chapter 99B.12 of the Iowa State Code, which allows betting pools within employee groups. This meant restricting participation to members of the University of Iowa community.

TABLE 2—TRADERS' PREFERRED CANDIDATES

Time	Preferred candidate			
	Bush	Dukakis	Other	Undecided
Before debate 1	47.0 (47)	42.0 (42)	2.0 (2)	9.0 (9)
Debate 1	43.1 (53)	47.2 (58)	0.8 (1)	8.9 (11)
Debate 2	45.8 (65)	45.1 (64)	2.1 (3)	7.0 (10)
Debate 3	50.0 (68)	41.1 (56)	1.5 (2)	7.4 (10)
Election night	46.7 (56)	50.8 (61)	2.5 (3)	—

Notes: Table entries are the fractions of the responding traders who favor a given candidate. The actual frequencies are given in parentheses.

percent were white, 71 percent were male, and more than 70 percent placed themselves in the middle or upper end of the income distribution.

Besides the initial enrollment survey, we also conducted telephone surveys of traders immediately following each debate and another telephone poll on election night to ask traders which candidate they had voted for if they had voted or which candidate they now preferred if they had not voted. These results are summarized in Table 2.[6] Whether and how the political and demographic characteristics of our trader population affect investment strategies are unclear. Theoretically, of course, we expect that these factors should have little bearing on individuals' efforts to maximize the return on their investments. Yet, individuals' backgrounds and preferences may affect their selection of information and perceptions of a candidate's viability and in turn affect their trading behavior. We investigate this possible connection below.

Final payouts were made following the election. The largest profit earned in the market was \$13.54 on an investment of \$250.00, while the largest loss was \$22.48 on an investment of \$95.00. Since investment levels varied across traders and traders entered the market at different times, an annualized rate of return may be a more informative measure of market gains and losses.[7] By this measure, one trader realized an annualized gain of 65.8 percent, while the biggest loser suffered an annualized loss of 659.0 percent. One-fifth of the traders earned an annualized rate of return of more than 10 percent, while 36 of the 190 traders incurred annualized losses of more than 10 percent.

The weekly trading-volume history of the market is given in Table 3. A simple characterization of these data is that trading started slowly, but by June 20 the market activity was sufficient to produce meaningful (i.e., approximately continuous) changes in prices. The month of July, which included the Democratic convention (July 18–21), shows relatively thick markets also, but trading drops off precipitously in August. This period covers the end of the summer term at the University of Iowa, and it is a popular time for vacations by faculty, staff, and students. Since the market required computer access, the departure of traders from the area had an obvious and pronounced effect on market activity. Unfortunately, this period includes the Republican Convention (August 15–18), and so we have nothing to say about the effect of Bush's selection of Dan Quayle as the Republican vice presi-

[6]The entries in the first row of Table 2 are different than those reported in Table 1. Table 2 includes all respondents, while the previous table only included information on those who said they were planning to vote.

[7]This annualized rate is computed as $365 \times$ profit/(investment dollar days), where "investment dollar days" is the weighted sum of investments with the weight being the number of days remaining on the market at the time of the investment.

TABLE 3—WEEKLY MARKET ACTIVITY

Week	Number of shares traded	Cumulative number of trades
June 1–5	46	46
June 6–12	31	77
June 13–19	34	111
June 20–26	241	352
June 27–July 3	68	420
July 4–10	86	506
July 11–17	183	689
July 18–24	90	779
July 25–31	41	820
August 1–7	4	824
August 8–14	8	832
August 15–21	3	835
August 22–28	34	869
August 29–September 4	190	1,059
September 5–11	1,625	2,684
September 12–18	1,302	3,986
September 19–25	1,075	5,061
September 26–October 2	1,503	6,564
October 3–9	1,656	8,220
October 10–16	4,087	12,307
October 17–23	783	13,090
October 24–30	1,349	14,324
October 31–November 6	1,342	15,666
November 7–9	832	16,498

dential candidate. Thinness of the market during August also precludes identification of the source of information leading to an increase (five cents) in Bush's price by late August. These gaps in the market are a noticeable consequence of the restriction of participation to members of the University of Iowa community.

II. Results

The outcome of the election and the market's predicted vote shares are shown in panel A of Table 4. The first column records the popular vote for the candidates, while column 2 shows the share of the popular vote. The third column contains the predicted shares of the vote based on the market prices of November 7, the eve of the election. These predictions are quite accurate. The actual dividends paid, which are based on information that was not available until November 22, are listed in column 4 of panel A. In essence, the market undervalued Dukakis by a penny and overvalued the combined strength of all third-party candidates by a penny.

Panel B of Table 4 contains the results of the final opinion polls for the six major polling organizations: ABC/*Washington Post*, CBS/*New York Times*, CNN/*USA Today*, The Gallup Organization, Lou Harris and Associates, and NBC/*Wall Street Journal*. The second column reports the date that the poll was conducted, and the third and fourth columns contain the fraction of the sample favoring Bush and the fraction favoring Dukakis, respectively. Because there are many respondents who either refuse to answer the presidential-preference question or list themselves as undecided, and because the polls do not usually report results for third-party candidates, except in elections with a significant national third-party candidate, we cannot directly compare the IPSM to the polls. To make the comparison we follow the convention of examining the marginalized lead for a two-candidate race. In the case of the polls, this is given by (percentage favoring Bush − percentage favoring Dukakis)/(percentage favoring Bush + percentage favoring Dukakis), and for the stock market the lead is given by (price of Bush security − price of Dukakis security)/(price of Bush security + price of Dukakis security). This transformation assumes, in effect, that the undecideds in the polls either will not vote or will vote in the same proportions as those who express a preference. Overall, 1988 was a good year for polls, but the stock market did even better judging from the results in Table 4.[8]

A comparison of poll results and IPSM prices is useful not only to measure how

[8]It might be objected that we have biased the comparison in favor of the IPSM by using information up to November 7, which is closer to the election than the dates of any of the polls. The average of the leads computed from market closing prices over the period November 2–5 is 7.03, for November 2–4 it is 6.92, for November 3–6 it is 7.16, and for November 1–5 it is 7.02. In all cases these estimates are closer to the outcome than polls conducted on the same dates.

TABLE 4—ELECTION RESULTS AND FORECASTS

A. *Election Results*

Candidate	Total votes	Vote share	IPSM forecast vote share	IPSM dividend
Bush	48,138,478	53.2	53.2	$1.33
Dukakis	41,114,068	45.4	45.2	$1.14
Third party	1,219,240	1.4	2.0	$0.03

B. *Opinion-Poll Forecasts of Bush's Lead*

Organization	Date of poll	Bush (percentage)	Dukakis (percentage)	Marginalized lead
ABC/*Washington Post*	November 2–5	54	44	10.20
CBS/*New York Times*	November 2–4	48	40	9.09
CNN/*USA Today*	November 3–6	52	42	10.64
Gallup	November 3–6	53	42	11.58
Harris	November 2–5	50	46	4.17
NBC/*Wall Street Journal*	November 1–5	48	43	5.49
IPSM	November 7	53.2	45.2	8.13
Election	November 8	53.2	45.4	7.87

well markets perform relative to polls,[9] but also to explore the information content of polls. For example, since the 1948 upset of Harry Truman over Thomas Dewey, pollsters have claimed that polls are not forecasts and thus no comparison is possible. As Irving Crespi (1988) has noted, adherence to this argument would make even election-eve polls worthless. Even if some forecasting role for polls is acknowledged, it requires a behavioral model of poll response to interpret polls. For example, suppose in a two-candidate race that voter i perceives now the probability that he will vote for candidate A in November as P_i. In response to a question concerning how he would vote if the election were held today, one response model would be to reply $v_i = 1$ if $P_i \geq \frac{1}{2}$, and $v_i = 0$ if $P_i < \frac{1}{2}$. The fraction favoring candidate A in the poll is the sum of the v_i. However, the expected fraction voting for candidate A in November is the sum of the P_i, and these can be very different. The market price should reflect the sum of the P_i, and thus there is reason to suspect that markets might perform better than polls. Alternatively, suppose that the voters respond to polls by stating $v_i = 1$ with probability P_i, and $v_i = 0$ with probability $1 - P_i$. Now polls and markets would report the same estimate.

Regardless of what polls measure, if traders find them to be a valuable information source it is possible for the market outcome to be just a thinly veiled shadow of the polls: traders, being relatively unsophisticated in discerning the mood of the electorate, will make trades based on the latest

[9]While this seems an obvious comparison, a pollster's view is that a poll should be interpreted as a reading of how the electorate would vote if an election were held that day, and not as a forecast of an election sometime in the future. Indeed, the polling question usually asked is a variant of "If the election were held today who [*sic*] would you vote for?" Nevertheless, it is difficult to see why this invalidates the use of poll results as a forecast. Suppose, for example, that a poll asked two questions: (i) "If the election were held today who would you vote for?" and (ii) "In the election to be held on November 8th who do you think you will vote for?" The results of the second question would surely be interpreted as a forecast, and yet it is difficult to imagine the circumstances under which truthful revelation of preferences would produce a different answer to the two questions. Of course, new information may arrive which leads an individual to vote differently than he previously thought he would, but this just reflects the obvious fact that forecasts are not 100-percent accurate.

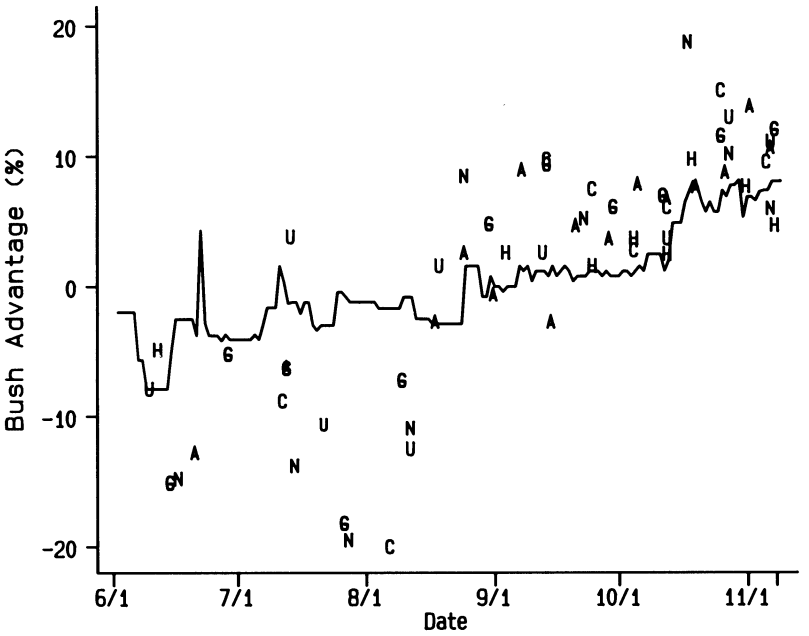

FIGURE 1. BUSH'S LEAD IN OPINION POLLS AND THE IPSM

Key: A = ABC/*Washington Post*, C = CBS/*New York Times*, G = Gallup, H = Harris, N = NBC/*Wall Street Journal*, and U = CNN/*USA Today*. The solid line depicts the IPSM results.

poll results.[10] If opinion polls contained genuine news, then one would expect their publication to affect the price of a candidate's stock. Alternatively, if polls do not convey news at all, but merely confirm what a knowledgeable observer could glean from other sources, then the advance knowledge of poll results should provide no additional profit opportunities. In this section, we provide graphical and regression-analysis evidence on the relation between opinion polls and the IPSM.

In Figure 1 we display the behavior of polls during the campaign in a comparison of the six major polls to the IPSM. As in Table 4, we use the marginalized lead for Bush. In the figure, the lead indicated by each poll is shown by a letter representing the polling organization. Since the IPSM was open for trading continuously, the lead implied by market prices is shown as a continuous series, though the points actually plotted are computed from the last-trade prices as of the end of the day. Inspection of the figure suggests that opinion polls are excessively volatile, certainly more so than can be attributed to sampling error alone. For example, after the Democratic convention Bush's lead falls to −17 percent, according to the average of the six polls, and yet in two weeks—after the Republican convention—it increases to +5 percent. Stated differently, if we were to believe these numbers it would imply that 22 percent of the American voters changed their support from Dukakis to Bush over a three-week period during which Bush chose Quayle as his running mate. In contrast, the lead im-

[10]Indeed, this claim was echoed by several of our colleagues.

plied by the market changed little over this period.

The view that the opinion-poll results should drive the market is a plausible one, a priori, since polls provide the most direct and most widely publicized evidence regarding the forthcoming election. However, the graphical evidence of Figure 1 raises doubts about such causality: given the large number of polls and their excessive volatility, both across polling organizations and across time, which poll should traders in the market believe? Moreover, examination of Figure 1 suggests that, while market prices and poll results followed the same general trend over the course of the campaign, movements in market prices often preceded movements in poll results, an effect contrary to the argument of causality leading from polls to prices; but such a visual test may not be very powerful.

The conventional assessment of causality, a Granger-Sims test, is not directly applicable to the problem at hand, since market prices comprise a continuous time series while poll results arrive at discrete and irregular intervals and from several sources. As an alternative, we present below the estimates of parameters in the following model:

$$(1)\quad M_t = E_t(M_t|\phi_{t-1}) + \beta[P_t - E_t(P_t|\phi_{t-1})]D_t + u_t$$

where M_t represents the market price on day t, E_t is the expectation operator which conditions on all information available just prior to time t (ϕ_{t-1}), P_t is the result of an opinion poll released on day t, D_t is an indicator variable which takes on the value 1 or 0 according to whether a new poll is released on day t, and u_t is white noise. The idea of this model is that polls convey information to the extent that their results are news (i.e., unanticipated), and only such unanticipated, information should influence market prices.

To operationalize this model, we measure M_t as the marginalized advantage for Bush as indicated by last-trade prices on day t, and expected market prices are assumed to be determined according to

$$(2)\quad E_t(M_t|\phi_{t-1}) = \alpha_0 + \alpha_1 M_{t-1} + \alpha_2 M_{t-2}.$$

P_t is treated as a vector of six elements, one for each of the major polling organizations,[11] and it too is measured as the marginalized advantage for Bush. The anticipated poll result is taken to be the previous poll result from the same polling organization, plus a response to news events common to the market and the poll since the last poll,

$$(3)\quad E_t(P_{i,t}|\phi_{t-1}) = \gamma_{i,0} + \gamma_{i,1}P_{i,t-s} + \gamma_{i,2}(M_{t-1} - M_{t-s}).$$

Least-squares estimates of the α's and β for this model appear in Table 5.[12] The first two columns contain results computed over the more active market period of September 1 through November 7, the day prior to the election. Entries in column (i) are for the constrained model with β fixed at zero, and those in column (ii) are for the unconstrained model. Values in parentheses below the regression coefficients are standard errors. As the p value in the last row indicates, the hypothesis that $\beta = 0$ (i.e., that the information in polls has no effect on market prices) would not be rejected at any reasonable significance level. No evidence of serial correlation can be found in the results for columns (i) and (ii), and at the 5-percent level one could not reject that $\alpha_2 = 0$. At least for this period, the evidence

[11] The numbers of polls released by each after the IPSM began operation on June 1 were as follows: ABC/*Washington Post*, 15; CBS/*New York Times*, 8; Gallup, 11; Harris, 9; NBC/*Wall Street Journal*, 9; CNN/*USA Today*, 9.

[12] Substituting (2) and (3) into (1) yields the estimating equation of interest. In Table 5 we report the estimates of α_i and β, which are of interest for the test of whether polls drive the market. The complete set of estimates is available from the authors upon request.

TABLE 5—INFLUENCE OF POLLS ON MARKET PRICES

Variable	September 1–November 7 (i)	September 1–November 7 (ii)	June 4–November 7 (iii)	June 4–November 7 (iv)	June 4–November 7 (v)
Intercept	0.1741	0.2958	0.0759	−0.0130	−0.0179
	(0.1452)	(0.1843)	(0.1003)	(0.1169)	(0.1174)
Price($t-1$)	0.7599	0.7650	0.7350	0.6967	0.7012
	(0.1190)	(0.1660)	(0.0785)	(0.0824)	(0.0828)
Price($t-2$)	0.2272	0.1517	0.2361	0.2493	0.2458
	(0.1195)	(0.1646)	(0.0795)	(0.0832)	(0.0835)
ABC/*WP*		0.0282		0.0376	0.0338
		(0.0632)		(0.0737)	(0.0741)
CBS/*NYT*		0.2299		−0.2009	−0.1935
		(0.5619)		(0.1387)	(0.1395)
Gallup		−0.5834		0.0098	0.0100
		(0.7427)		(0.0661)	(0.0663)
Harris		0.4180		0.2160	0.2090
		(0.3059)		(0.3234)	(0.3242)
NBC/*WSJ*		0.0916		0.1706	0.0795
		(0.0660)		(0.0565)	(0.1487)
CNN/*USA*		0.2223		0.0735	0.0459
		(0.2760)		(0.0775)	(0.0882)
Number of observations:	68	68	157	157	156
R^2:	0.923	0.943	0.890	0.909	0.909
p value for serial correlation:	0.850	0.451	0.836	0.849	0.882
p value for exclusion of poll variables:		0.568		0.052	0.789

Notes: Main table entries are regression coefficient estimates, and entries in parentheses are standard errors. The dependent variable is market price, transformed to reflect the vote-share advantage for Bush. Poll variables are computed as current result minus the expected result for the same poll, with the poll result measured as the vote-share advantage for Bush. The test for serial correlation is the LM test for four lags. Data used for column (v) exclude the August 25 observation.

mildly supports the hypothesis that the market followed a random walk.

The period from September 1 to the election is the period when the market had a relatively large number of active traders. As noted above, the market was largely inactive in early August because most traders were away from campus. However, this is the period when Bush's lead in the polls first appears, and thus focusing on the post-August period may bias the case against polls. Columns (iii) and (iv) contain results from the full market period,[13] with the null model in column (iii) and the model including poll information effects in column (iv). The overall test for the effect of polls is

[13]The first day with shares traded and thus with observable market prices in both Bush and Dukakis was June 2, and two observations are lost because of the lagged values of the dependent variable. Thus the sample period covers 157 market days from June 4 through November 7.

statistically significant over this full period, as the p value of 0.052 indicates. However, the most important poll, judging from the size of the coefficients and the t ratios for them, is the NBC/*Wall Street Journal* poll, and close inspection of the data reveals that all of the influence of this poll arises from a single poll release. The NBC/*Wall Street Journal* poll of August 25 suggested a Bush lead of 8 percent, up nearly 20 percent from the −11.4 lead reported in the August 12 poll.[14] August 25 was also the beginning of the fall semester at the University of Iowa, and the return of student and faculty traders led to greatly increased market activity, including the first trades in over a week and a one-day volume exceeding the total of all previous 24 days in the month. The coincidence of these two events might well represent a spurious correlation between market prices and the NBC/*Wall Street Journal* poll. When the August 25 observation is excluded, as it is in the results in column (v), the significance of polls disappears.

These results are robust to extensive "fiddling" with the specification of the model. For example, use of a single measure of poll results, computed either as the average of the six major polls or as the latest poll release irrespective of the polling organization, yields effects of polls that are slightly less significant. The same is true if the information content is computed using the most recent poll release from any polling organization, rather than the organization's own most recent poll. Monthly indicator variables, introduced as a crude measure of structural shifts, are highly significant, but the pattern of significance of polls is unaffected. When lagged poll releases are included as part of the specification of $E_t(M_t)$, they are insignificant and have no effect on the pattern or level of significance of the current poll information variables. Finally, inclusion of both current and lagged values of poll information variables, as a test of whether the impact of polls is fully realized on the release day alone, shows no significant effect.

Thus the data reject the view that polls drive the market. Evidently, traders were able to find out about the mood of the electorate without relying on opinion polls.[15] In this sense, polls are not "news" to traders who have an incentive to seek out information from other sources. This does not mean that the market does not react to news events. Indeed, events like the naming of Lloyd Bentsen as a vice presidential candidate, Jesse Jackson's announcement that he would not run on a third-party ticket, and the outcomes of the presidential debates had predictable effects on the stocks of the candidates. In each of these cases the market reacted quickly to the events, but not to the poll results about the event.[16]

III. Judgment Bias in Markets

In this section we examine whether the traders in this market exhibited any systematic trading biases. Based on survey data, political scientists often claim to find some biases in respondents' judgments. For example, respondents often engage in wishful thinking and respond more often than not that their candidate is likely to win (Uhlaner and Grofman, 1986). Since we conducted regular surveys on our traders, we can use these data to test whether they reflect similar biases and, to the extent that they do, we can also examine the extent to which these biases affect their trading behavior. As we report below, we find substantial evidence that these biases affect trading behavior *on average*. Since the market worked quite well, we proceed to identify a set of *marginal* traders and examine how their trading behavior affected prices.

[14]The only other poll released on August 25 was an ABC/*Washington Post* poll which reported a Bush margin of 2.1, up from a deficit of 3.2 according to the August 18 ABC/*Washington Post* poll.

[15]That newspaper reports of events might not constitute "news" is not novel. Studies of detailed industries, such as Norman Frey and John W. Labuszewski's 1981 examination of the copper industry, show that market participants typically have access to more informed sources than do reporters.

[16]See Forsythe et al. (1991b) for a description of news in the 1988 election.

TABLE 6—RELATION BETWEEN PREFERENCE AND EXPECTATION IN U.S. PRESIDENTIAL ELECTIONS

Year	Democrat/Republican	Percentage of respondents intending to vote Democratic who expect Democrat to win	Percentage of respondents intending to vote Republican who expect Democrat to win
1988	Dukakis/Bush	51.7	5.8
1984	Mondale/Reagan	28.8	1.0
1980	Carter/Reagan	87.0	19.6
1976	Carter/Ford	84.2	19.4
1972	McGovern/Nixon	24.7	0.4
1968	Humphrey/Nixon	62.5	4.6
1964	Johnson/Goldwater	98.6	69.5
1960	Kennedy/Nixon	78.4	15.8
1956	Stevenson/Eisenhower	54.6	2.4
1952	Stevenson/Eisenhower	81.4	14.1

Source: Donald Granberg and Edward Brent (1983), who use survey data collected by the Survey Research Center/Center for Political Studies of the University of Michigan. Entries for 1984 and 1988 were obtained from correspondence with Donald Granberg.

Political scientists and psychologists have documented two types of bias whereby an individual's preference for a candidate colors the voter's objective view of the election campaign. The first bias, termed the assimilation-contrast effect in psychology (M. Sherif and C. I. Hovland, 1961; Allen Parducci and Louise M. Marshall, 1962), states that an individual's preference for an outcome biases his or her interpretation of information about the likelihood of the outcome occurring. Thus, after viewing an event such as a debate, partisans are more likely to respond that their candidate's fortunes were favorably affected by the event. The second bias, termed the false-consensus effect by psychologists (Lee Ross et al., 1977; Clifford E. Brown, 1982) states that individuals tend to overestimate the extent to which their views are representative of the population. An example of this bias in elections is shown in Table 6. Even in a very close election like 1976, more than 80 percent of the supporters of each candidate thought their preferred candidate would win. Although the bias is truly large, it is not clear whether it is due to the infirmities of polling or whether the consequences of judgment bias are truly significant. In particular, the respondents had nothing at stake when they answered the question of who they thought would win the election.

Even when participants do have something at stake, we still find evidence of judgment bias. The nature of the first type of judgment bias observed in the IPSM is as follows. When we studied the effect of demographic variables on trading behavior, we observed that traders appeared to be maximizing income;[17] but when we examined political preferences, we found a different result. While trading behavior around the presidential debates was affected by beliefs about the outcome of the debate, judgment

[17]For example, while female participants exhibited a gender bias in preferences roughly similar to that detected in national polls (a preference advantage of 8–11 percent for Dukakis over Bush), their trading behavior was indistinguishable from that of males. Similar results were found for other demographic variables (e.g., income level, education, religious preference). It is noteworthy that participants in the IPSM were not a representative sample of either the U.S. or Iowa electorates: they are too well educated, too white, too well-off, and too liberal, even by Iowa standards.

TABLE 7—PORTFOLIO COMPOSITION BY PERCEIVED DEBATE OUTCOME: CHANGE IN FRACTION OF MAJOR-CANDIDATE SHARES HELD IN BUSH STOCK (Standard Errors in Parentheses)

	Candidate thought to have won			
	Bush	Dukakis	No one	All other traders
Debate 1	4.8	−4.0	0.0	1.0
	(2.3)	(2.2)	—	(1.0)
Debate 2	1.5	−1.6	0.0	−1.2
	(1.6)	(0.7)	—	(1.7)
Debate 3	2.2	−0.9	0.0	−2.8
	(3.4)	(1.0)	—	(1.8)

bias arose because traders who expressed preferences for a candidate were significantly more likely to believe that their candidate had won the debate. For example, 80.4 percent of the Dukakis supporters thought Dukakis did better than Bush in the third debate, while only 22.2 percent of the Bush supporters shared that view. While our observation of such a phenomenon is not unique, the distinctive result from the IPSM is that traders were willing to back their disparate opinions with cash.

Table 7 shows the changes in the fraction of Bush shares held (number of Bush shares held/[number of Bush shares + number of Dukakis shares]) categorized by which candidate was perceived to have won the debate. In the table it can be seen that those traders who believed that the Bush team won consistently increased their holdings in Bush stock, while those who believed Dukakis to have won the debate consistently reduced their holdings of Bush stock after the debate. Interestingly, those who thought that there was no winner in the debate (23 traders in debate 1, 10 traders in debate 2, and 16 traders in debate 3) acted consistently with those beliefs and did not trade within two days after the debate.

Evidence of the second type of bias can also be observed in the IPSM data. With this bias, traders who prefer Bush should anticipate a larger share of the popular vote for their candidate, and accordingly, they should demand more shares of Bush stock (and supply fewer shares) at each price than traders who prefer Dukakis. That is, the false-consensus effect in this context can be translated as shifts in the position of the excess-demand function for a good. Bush supporters, believing that more of the population share their beliefs, should buy more Bush stock than they sell, while Dukakis supporters should sell more Bush stock than they buy. By analogous reasoning, Dukakis supporters should buy more Dukakis stock than they sell, while Bush supporters should sell more Dukakis stock than they buy.[18]

These biases are evident in the data of Table 8. Here we have matched individual trading data to political preferences as revealed by the survey of trader preferences. We computed net purchases (purchases − sales) by day across all traders within the Bush and Dukakis constituencies. The second row of Table 8 contains the mean value of such net purchases across all days in the market with active traders from the respective constituency, while rows 3 and 4 provide counts of the number of days in which purchases by the constituency exceeded sales. Test statistics, including a t ratio for testing the difference between means for row 2 and a chi-square statistic for testing

[18]In a different context, Varian (1985, 1986) has argued that differences in opinion such as these can lead to trade. After observing news, Dukakis supporters will, in general, be optimistic about what they think Dukakis' shares are worth while Bush supporters will be pessimistic. These differences in opinion will cause the demand curve to rotate, and trade will occur.

TABLE 8—MEASURES OF QUANTITY DIFFERENCES BETWEEN BUSH AND DUKAKIS SUPPORTERS

Row	Measure	Transactions in Bush stock			Transactions in Dukakis stock		
		By Bush supporters	By Dukakis supporters	Test statistic	By Bush supporters	By Dukakis supporters	Test statistic
1	Number of days with active trade in either candidate	92	80		92	80	
2	Average of net purchases across active days	1.12 (1.11)	−1.59 (0.96)	1.85 [170]	−2.26 (0.90)	2.01 (0.86)	−3.43 [170]
3	Number of days in which purchases exceed sales	43	21		24	46	
4	Number of days in which sales exceed purchases	30	42	[1]	46	20	[1]
5	Average across active days of number of traders in this candidate	2.41 (0.20)	1.53 (0.15)		1.67 (0.17)	2.00 (0.19)	
6	Average proportion of traders with net purchases > 0	0.62 (0.038)	0.31 (0.048)	5.06 [146]	0.39 (0.044)	0.68 (0.045)	−4.61 [143]
7	Average proportion of traders with net purchases < 0	0.37 (0.038)	0.68 (0.048)	−5.06 [146]	0.60 (0.044)	0.31 (0.044)	4.66 [143]
8	Number of days with active *new* traders	51	45		51	45	
9	Average of net purchases across active days by new traders	1.63 (0.58)	−0.20 (0.49)	2.41 [94]	0.49 (0.56)	2.09 (0.46)	−2.21 [94]
10	Number of days with active *old* traders	82	66		82	66	
11	Average of net purchases across active days by old traders	0.24 (1.23)	−1.79 (1.12)	1.22 [146]	−2.84 (1.01)	1.01 (1.00)	−2.71 [146]

Note: Entries in parentheses are standard errors; entries in brackets are degrees of freedom.

independence with the count data of rows 3 and 4, suggest significant differences in trading behavior between the two constituencies. These differences could have been generated by a small number of traders with deep pockets, but that was not the case. Rows 6 and 7 of the table, which contain means of the proportion of traders in a constituency who were net buyers and net sellers, respectively, reveal that on average 62 percent of the Bush supporters bought more Bush stock than they sold while 68 percent of Dukakis supporters bought more Dukakis stock than they sold. Thus, the phenomenon appears to be pervasive. Again, the difference between support groups is statistically significant and in the direction suggested by judgment bias: Bush supporters revealed a higher valuation of Bush stock relative to Dukakis supporters by buying significantly more of it. Finally, the differences are not attributable merely to new and naive traders. Rows 9 and 11 contain the daily mean net purchases of stock by a constituent group as in row 2, including only trades undertaken by a trader during the first seven days after entry into the market (row 9) and only trades undertaken after the first seven days in the market (row 11). Again, all the differences between constituent groups are in the direction suggested by judgment bias, and all but one would be judged statistically significant.

IV. Why the Market Worked

In spite of these judgment biases, the market provided an accurate prediction of the outcome. How could this happen? To explain it, we have three competing hypotheses. The first is that we were simply lucky and that other markets will perform poorly. We find this unlikely because since the 1988 U.S. presidential election we have

conducted seven more election markets: the 1990 Iowa Senate race, the 1990 Illinois Senate race, three markets in Germany for the 1990 parliamentary election, and the parliamentary elections in Denmark (1990) and Turkey (1991). With the exception of the Illinois market,[19] these markets worked well in forecasting election results.[20] Thus, a pure-luck explanation is not possible.

A second possibility is that we had a representative sample of traders in terms of presidential preferences. If this were the case, traders who engage in market activities according to their (theoretically irrelevant) preferences could lead the market to a correct forecast if they were present in just the right numbers. Again the data argue against this interpretation. When we classify traders by political preference we find that, when first enrolled in the market, Bush supporters were 48.1 percent of all traders (compared to 49.6 percent for Dukakis), but that by election day only 46.7 percent of the traders supported Bush. Markets, of course, are not one-man, one-vote institutions. When we account for wealth holdings we find that Bush supporters held 55.2 percent of the total market value on enrollment day but that this had fallen to 49.7 percent by election day. In short, an offsetting-biases explanation fails because, measured by numbers or dollars, the traders were not sufficiently biased toward Bush. Indeed, over the campaign more traders switched to Dukakis.

A third hypothesis for the success of the market is the marginal-trader hypothesis. The motivation for this hypothesis is the simple argument that the results on judgment bias refer to the average trader, while prices are determined by the marginal trader. This hypothesis has a long tradition in economics, appearing in several guises as the argument demanded, but the hypothesis has rarely been made operational. For example, in markets for on-the-job safety (Walter Oi, 1974; W. Kip Viscusi, 1979, 1983) it is frequently asserted that allocations will be correctly made so long as the "marginal" worker is correctly informed. Since the preferences of the marginal worker usually are not known, identification of equilibrium prices with "correct" allocations is tautological. An alternative definition of marginal traders stems from the idea of arbitrageurs—a group of individuals who trade when market prices exceed their beliefs about, or forecasts of, future events. Under this definition the hypothesis is that there were some traders who did not suffer from judgment bias and that they profited by acting as arbitrageurs. In other words, these marginal traders would buy Dukakis stock from (biased) Bush supporters and sell it to (biased) Dukakis supporters. Theoretically, this would require arbitrageurs to have sufficient information about the other traders in the market. In particular, they would have to know the extent of the judgment bias exhibited by traders of each political preference, and they would have to know the fraction of traders who preferred each candidate. While it seems implausible that any trader in the IPSM had such information, perhaps some had a sufficient intuitive grasp to play the arbitrageur's role successfully.

Because the audit trail for this market contains information on the timing of all transactions we can construct a measure of trading activity that serves to identify marginal traders. In particular, the audit-trail information allows us to distinguish between "limit orders" (offers that are entered into the queues without resulting in immediate trades) and "market orders" (offers that amount to acceptance of an outstanding ask or bid and thus result in an immediate exchange). The traders to be identified as "marginal traders" are those who submit limit orders at prices close to the market price, while those not identified as marginal are traders who are inactive,

[19]The market for the 1990 Illinois Senate race was conducted concurrently with the Iowa Senate market, and enrollment in the market was limited to members of the University of Iowa community. Individuals could trade in either or both markets, but despite the large contingent of Illinois residents attending the University of Iowa, there was little interest in trading in the Illinois securities. Thus the poor performance of this one market is due mostly to lack of interest.

[20]We note, parenthetically, that one of these markets was run in Leipzig, in the former East Germany, an area which then had no experience with democracy or free markets for over 45 years.

TABLE 9—A COMPARISON OF MARGINAL TRADERS' CHARACTERISTICS

Trader characteristics	Marginal traders	Nonmarginal traders	p value
Number	22	170	
Party preference = Republican	38	39	0.46
Percentage favoring Bush:			
Initial survey	50	43	0.61
First debate	53	42	0.38
Second debate	59	44	0.25
Third debate	60	49	0.41
Male	100	68	0.01
Average total investment ($)	56	21	0.00
Median rate of return (percentage):			
Full period	9.6	0.0	
October 17–November 8	11.6	0.0	

Note: The p values in column 4 are obtained from Fisher's exact test for categorical variables (rows 2–7) and from the t test for continuous variables (row 8).

who submit only limit orders at prices far away from the market price, or who submit only market orders.[21]

Formally, we define an indicator for each trader on each day which takes the value 1 if either of two events occur: that trader had a bid or an ask in the queue at the end of the day at a price within two cents of the price of the last trade of the day, or the trader had a bid or ask in a queue which was accepted by another trader sometime during the day. We then identify a trader as marginal if and only if the indicator for that trader is 1 for at least three days during the relevant period. Because the greatest movement in market prices occurred after the second presidential debate, we focus on the period from October 18 to November 7. The precise definition of marginal trader is somewhat arbitrary: alternative definitions that use a few more or a few less days of market activity to classify a trader as a marginal trader lead to essentially the same results. Table 9 shows the characteristics of traders identified as "marginal" in this manner.[22]

Standard demographic factors differ little between marginal traders and other participants in the market. The fraction identifying the Republican party as their affiliation is roughly the same in each group. Similarly, education, income, age, and religion did not differ between the two groups. The one demographic factor that did stand out was sex: marginal traders were significantly more likely to be males than the other participants in the market. The reason for this difference is not obvious. Participation required access to a computer, but on-campus access to computers was available in computer clusters in the library, all computer laboratories, and in the dormitories. However, 34 percent of the males, and 23 percent of the females, reported that their primary access to the market was by a home or office personal computer.

[21]It is impossible to tell whether market orders represent marginal or nonmarginal trading behavior. For example, a market order resulting in a purchase at a price of $1.00 could have been submitted by someone willing to pay only $1.01, which would qualify it according to our definition as a "marginal trade." Alternatively, it might have been submitted by someone willing to pay as much as, say, $1.50, which would then disqualify it as a "marginal trade." Thus market orders are not used to identify marginal traders.

[22]One individual satisfied the technical definition of a marginal trader by selling all his Bush and Dukakis stock in order to buy Jackson stock (which was worthless at the time). His motive, offered in a postexperiment survey, was to "send a message to Jesse." We excluded this trader from the marginal-trader analysis which follows.

The marginal traders appear to have preferred Bush as a candidate more than the other traders, but the sample size is too small to pin this down with any precision. Interestingly, the marginal traders moved more to Bush after the vice presidential debate, a result that is consistent with the national pattern in polls.

The major difference between marginal and nonmarginal traders is the total investment made in the market. Marginal traders on average invested \$56, more than twice the level of nonmarginal traders ($t = 6.0$). They also traded more shares and were active in the market on more days. This is consistent with the hypothesis that judgment bias diminishes when more is at stake. More importantly, marginal traders earned higher returns. Over the entire market the median rate of return for nonmarginal traders was 0 percent, but for the group of marginal traders it was 9.6 percent. Some evidence of learning about trading is indicated in that returns for the later periods are even higher for marginal traders. Specifically, calculating the annualized rate of return based on portfolio and cash holdings immediately after the third presidential debate shows that the median return for nonmarginal traders was 0 percent, but for marginal traders it was 11.6 percent. In this case, it appears that the prospect of supracompetitive earnings led, or allowed, some individuals to arbitrage away the judgment biases of other traders. Interestingly, apart from the sex and market investment differences noted in Table 9, standard probit analyses of who is a marginal trader failed to reveal any significant factors among the data we had collected. Whatever leads individuals to operate without judgment bias does not appear to be related to conventional demographic or political preference measures, a result that probably is not surprising to Wall Street recruiters.

The pattern of returns suggests that our definition of marginal traders does pick out individuals whose judgment about political events was not clouded. To examine the issue of judgment bias directly we have repeated part of the analysis shown in Table 8 for marginal and nonmarginal traders separately. Table 10 reveals that the marginal traders show no indication of judgment bias in their transactions. The difference between the average net purchases of stock in the two major candidates is small and statistically insignificant for both Dukakis and Bush supporters, (values of 0.24 and -1.07, both with 104 degrees of freedom). In contrast, the remaining traders exhibit large and statistically significant evidence of judgment bias: nonmarginal Bush supporters purchase significantly more shares of Bush stock, while nonmarginal Dukakis supporters purchase significantly less. While we cannot be sure that we have identified all marginal traders, these trading patterns persist if the set of marginal traders is narrowed or broadened somewhat.[23] How did the marginal traders accomplish these gains? The trading data suggest that they did so by being able to recognize when "news" happened and when it did not. If nothing has happened to change a trader's view of the outcome of an election, she should be selling when the price rises, and buying when the price falls. That is, changes in prices and changes in quantities should be negatively correlated when there is no new information. Conversely, when news arrives, traders should be buying when the price is rising, and selling when the price is falling. In other words, when news occurs, changes in prices and changes in quantities should be positively correlated.

To measure how traders performed we need an operational definition of news. As defining what is news in a market is notoriously difficult, we followed two approaches. First we defined a day when news arrived as one when the absolute value of the change in price of Bush stock was greater than two standard deviations of the price series in practice. Our second method of identifying "news" days, was to examine the postelection assessments of reporters for what they viewed as significant events in the cam-

[23]Specifically, varying the number of trading days on the margin from 2 through 5 produced essentially the same results. Similarly, changing the two-cent rule to a one-cent rule also produced the same patterns.

TABLE 10—MEASURES OF QUANTITY DIFFERENCES BETWEEN BUSH AND DUKAKIS SUPPORTERS

Measure	Transactions in Bush stock			Transactions in Dukakis stock		
	By Bush supporters	By Dukakis supporters	Test statistic	By Bush supporters	By Dukakis supporters	Test statistic
Marginal traders:						
Number of days with active trade in either candidate	63	43		63	43	
Average of net purchases across active days	0.10 (1.80)	−0.49 (1.39)	0.24 [104]	−2.21 (1.43)	0.26 (1.83)	−1.07 [104]
Nonmarginal traders:						
Number of days with active trade in either candidate	82	71		82	71	
Average of net purchases across active days	1.18 (0.56)	−1.49 (0.63)	3.18 [151]	−0.84 (0.53)	2.11 (0.65)	−3.55 [151]

Note: Entries in parentheses are standard errors; entries in brackets are degrees of freedom.

paign.[24] From these we counted events as "news" if the price change on the associated days were in the lower or upper 5-percent tail of price changes.

The period from September 1 to election day comprises 68 days. During this period the number of days when prices and quantities in Bush stock were negatively and positively related for marginal traders, using $\text{abs}(\Delta p) > 2\sigma$ to define news days, was:

	$\Delta p \Delta q \leq 0$	$\Delta p \Delta q > 0$
No news	52	12
News	1	3.

Using news defined by reporters yields the same cross-tabulation results, although the categorization of what is news does differ. Using changes in prices to define news, we see that in 52 of the 64 days when there was no "news" about the campaign, marginal traders were not buying Bush stock when prices were rising, or not selling it when prices were falling. Conversely, when significant price changes occurred, the marginal traders were trading in the direction of the price change on three of the four days. Essentially the same story is seen when we use political commentators' assessments of news events. Thus, a significant feature of a marginal trader appears to be the ability to recognize what really is news.

V. Summary and Conclusions

The Hayek hypothesis is a bold claim about market behavior. It asserts that markets work even when participants know very little about the environment or about other participants. Even economists are rightfully surprised when the hypothesis appears to hold, as it does in the experimental markets surveyed by Smith (1982). This paper provides further evidence favorable to the hypothesis in a broader environment. Specifically, we find that even in a nontraditional market such as the presidential stock market analyzed here, and even with a relatively small number of traders, interaction through a market structure produces correct prices. Moreover, these interactions produce the correct prices even in an information-processing environment that should impede, if not prevent, markets from operating efficiently. Judged on a relative standard (i.e., compared to polls), the market worked well, and our results suggest that even the availability of polling information had little effect on market prices. Indeed, the information environment, including the information

[24] See Forsythe et al. (1992) for a description of the sources used in constructing "news" events from political scientists' and reporters' assessments. The reporter-defined news events were: (i) September 20, Dukakis proposes mandatory employer-provided health insurance; (ii) October 13, second presidential debate; (iii) October 17, Teamsters endorse Bush; (iv) October 31, Dukakis asserts he is a liberal in the Roosevelt sense.

provided by polls, provides insight into why the Hayek hypothesis works. Traders in any, perhaps all, markets have different talents, interests, and abilities; they may interpret data differently or be swayed by fads. However, as long as not all traders are so influenced there is room for markets to function efficiently. Measures of judgment bias produced by social scientists invariably are measures of *average* behavior. But, as the good old-time religion has it, market-clearing prices are set by marginal, not average, behavior, and it is for this reason that the Hayek hypothesis is robust.

REFERENCES

Brown, Clifford E., "A False Consensus Bias in 1980 Presidential Preferences," *Journal of Social Psychology*, October 1982, *118*, 137–8.

Crespi, Irving, *Pre-Election Polling: Sources of Accuracy and Error*, New York: Russell Sage Foundation, 1988.

Forsythe, Robert, Nelson, Forrest, Neumann, George R. and Wright, Jack, (1991a) "Forecasting the 1988 Presidential Election: A Field Experiment" in R. M. Isaac, ed., *Research in Experimental Economics*, Vol. 4, Greenwich, CT: JAI Press, 1991, pp. 1–44.

_____, (1991b) "Forecasting Elections: A Market Alternative to Polls," in Thomas R. Palfrey, ed., *Contemporary Laboratory Experiments in Political Economy*, Ann Arbor: University of Michigan Press, 1991, pp. 69–111.

_____, "The Dynamics of an Election Campaign: An Analysis of Events in the 1988 Presidential and 1990 Iowa Senate Elections," mimeo, University of Iowa, 1992.

Frey, Norman and Labuszewski, John W., "Newspaper Articles and Their Impact on Commodity Price Formation—Case Study: Copper", *Journal of Futures Markets*, Spring 1981, *1*, 89–91.

Granberg, Donald and Brent, Edward, "When Prophecy Bends: The Preference–Expectation Link in U.S. Presidential Elections," *Journal of Personality and Social Psychology*, September 1983, *45*, 477–91.

Milgrom, Paul and Stokey, Nancy, "Information, Trade and Common Knowledge," *Journal of Economic Theory*, February 1982, *26*, 17–27.

Oi, Walter Y., "On the Economics of Industrial Safety," *Law and Contemporary Problems*, Summer/Autumn 1974, *38*, 669–99.

Parducci, Allen and Marshall, Louise M., "Assimilation vs. Contrast in the Anchoring of Perceptual Judgments of Weight," *Journal of Experimental Psychology*, May 1962, *63*, 426–37.

Ross, Lee, Greene, David and House, Pamela, "The False Consensus Effect: An Egocentric Bias in Social Perception and Attribution Processes," *Journal of Experimental Social Psychology*, May 1977, *13*, 279–301.

Rubinstein, Mark, "Securities Market Efficiency in an Arrow-Debreu Economy," *American Economic Review*, December 1975, *65*, 812–24.

Sherif, M. and Hovland, C. I., *Social Judgment: Assimilation and Contrast Effects in Communication and Attitude Change*, New Haven, CT: Yale University Press, 1961.

Sigelman, Lee and Sigelman, Carol K., "Judgments of the Carter-Reagan Debate: The Eyes of the Beholders," *Public Opinion Quarterly*, Fall 1984, *48*, 624–8.

Smith, Vernon L., "Markets as Economizers of Information: Experimental Examination of the 'Hayek Hypothesis'," *Economic Inquiry*, April 1982, *20*, 165–79.

Uhlaner, Carole J. and Grofman, Bernard, "The Race May Be Close but My Horse Is Going To Win: Wish Fulfillment in the 1980 Presidential Election," *Political Behavior*, 1986, *8* (2), 101–29.

Varian, Hal R., "Divergence of Opinion in Complete Markets," *Journal of Finance*, March 1985, *40*, 309–17.

_____, "Differences of Opinion in Financial Markets," in C. C. Stone, ed., *Financial Risk: Theory, Evidence and Implications*, Norwell, MA: Kluwer, 1986.

Viscusi, W. Kip, "Job Hazards and Worker Quit Rates: An Analysis of Adaptive Work Behavior," *International Economic Review*, February 1979, *20*, 29–58.

_____, *Risk By Choice: Regulating Health and Safety in the Workplace*, Cambridge, MA: Harvard University Press, 1983.