Subjective Information and Market Efficiency in a Betting Market

Stephen Figlewski

New York University Graduate School of Business Administration

Much of the information available to participants in speculative markets is in the nature of expert opinion, analysis, professional advice, and so on. Markets discount widely held factual information very well; this paper studies market efficiency with respect to subjective information. We examine the "market" for bets on thoroughbred horse races to determine whether the published forecasts of professional handicappers are completely discounted. A multinomial logit probability model is used to measure the information content of the forecasts, and we find that they do contain considerable information but that the track odds generated by betting discount almost all of it. Within the population of bettors, those betting at the track appear to discount the handicapper information fully, but those betting through New York's off-track betting system do not.

I. Introduction

Statistical analysis of financial markets provides considerable evidence that they process publicly available information very well. The collective wisdom of the market rarely overlooks any widely held information which would allow an astute speculator to make higher-than-average profits. But while empirical tests almost always support the efficient-markets hypothesis that a speculative market fully discounts

I would like to thank the New York Racing Association and the New York Off-Track Betting Corporation for help in obtaining the data used in this study. During the research I profited greatly from conversations with Ray Kerrison of the *New York Post* and Henry Baker of the *Racing Form*. Thanks also to Steven Sheffrin, William Silber, and, especially, to Naeem Fayyaz, whose careful assistance in gathering the data was invaluable.

[*Journal of Political Economy*, 1979, vol. 87, no. 1]

all publicly available information in producing a market price, it is impossible to test specifically whether a market really discounts *all* information. The simplest tests, which involve looking for statistical irregularities in the price series itself that might be exploited by a mechanical trading rule, invariably support market efficiency. Other tests examine market behavior around specific events such as stock splits, dividend changes, changes in the Federal Reserve's discount rate, and so on, and have also tended to show that once news of an event is public knowledge it is completely discounted by the market.

However, much publicly available information is of a type for which market efficiency has never really been tested. This information is what we might call "subjective information" or "expert opinion." In addition to factual data, participants in speculative markets also have access to interpretation and opinion from market analysts, brokerage houses, other investors, and other more or less qualified sources. This subjective information is a major input into the decision-making process of virtually all market participants. We should like to know whether competitive markets are as efficient in discounting it as they seem to be in dealing with factual data. The question is especially interesting since the objective of most market analysis is to produce information that has not already been discounted. A market letter that consistently said the market was efficient and all securities were correctly priced would be of no interest to investors.

Unfortunately, testing market efficiency with respect to subjective information is problematical, since such information is seldom available in an explicit enough form that it can be easily quantified and analyzed statistically. Ideally, we would like to observe a market in which, first, a number of experts made precise forecasts about the outcomes of a set of investment strategies; next, these forecasts were widely disseminated among the participants in the market; and, finally, the market generated a set of prices which could be analyzed to see if they fully discounted the experts' information. Normal financial markets do not operate in this way, but this exact situation does occur in the "market" for bets on thoroughbred horse races. Before each race at a major track a number of professional handicappers study each horse, pick the three they consider most likely to come in first, second, and third, and then publish their opinions. The choices of the handicappers writing for major newspapers are widely available to the bettors well before the race. If the market for bets is efficient, it should be true that the track odds at race time fully reflect the information contained in these handicappers' forecasts.

Data on racetrack betting have been used before to study economic behavior in a risky situation with well-defined payoffs. Rosett (1965) compared the odds on simple bets with those on more complicated strategies involving compound bets and concluded that bettors are rational and sophisticated in evaluating the various betting possibilities they have available. Weitzman (1965) and Ali (1977) both analyzed the relation between the track odds produced by betting and the objective probability that a given horse will win, in order to estimate bettors' utility functions. So far, however, the possibility of using horse-race data to observe the impact of subjective information on a competitive betting market has been overlooked.

In this paper I analyze the forecasts of 14 professional handicappers and test whether the competitive market for bets is efficient with respect to this information. It is possible to show with a very high degree of confidence that the handicappers do possess information and that much of it, though perhaps not all, is discounted by the bettors. Further, efforts at out-of-sample prediction indicate that it may be very difficult to sift any additional useful information out of the forecasts even with relatively sophisticated statistical methods. Betting on thoroughbred horse races appears to be efficient with respect to the subjective information provided by professional handicappers. However, we do find that subsets of the bettor population may differ in their treatment of information. In particular, those betting through the New York off-track betting system seem to make less efficient use of the handicapper information than do on-track bettors.

Section II describes the data and the multinomial logit model I use to analyze the handicapper forecasts. Section III presents the empirical results, and Section IV contains concluding comments.

II. The Model

Betting on horse races is a custom dating back thousands of years. (One may wonder whether the comparatively recent development of trading in corporate equity will prove to be as durable an institution.) In the present-day United States, as in many other countries, horse racing is highly institutionalized, standardized, and regulated by governmental authorities, with the result that bettors are able to place wagers on what has become a closely controlled probabilistic event, at odds that are determined by impersonal market forces in a competitive market for bets. Betting on a horse race is by its nature very different from betting on a game of chance like roulette. Roulette is a "game against nature" in which the probabilities are known ahead of time. In Knight's (1965) terminology, it is a game with risk but no uncertainty. Horse racing, on the other hand, presents the bettor with both risk and uncertainty. While each horse in a race can be thought of as having a certain probability of winning, these probabilities are not and can not be known. This changes the complexion of the game considerably. A rational bettor places a bet on a horse not simply for the pleasure of taking a risk, but because he believes the odds being offered understate the true probability that the horse will win. He is betting that his estimate of the horse's chances is more accurate than the market's estimate. Unlike a casino gambler, a good horseplayer does not count only on luck. There is a reward to gathering information to improve one's probability estimates, and the bettors with superior information and ability to analyze it will be more successful than the rest.

We will be examining data from 189 thoroughbred races run at Belmont (New York) racetrack during June and July 1977. During the racing season there are nine races a day featuring between four and 12 bettable horses. In our sample, the average number of horses per race is between eight and nine. On occasion, more than 12 horses actually compete, but only if two or more of them run "coupled." This means that because they have the same owner or trainer they are treated by the track as a single entry. In the analysis, all data for coupled horses were merged, and the entry was treated as a single horse.

Races are made up about 2 days before they are actually run, and several factors insure that all of the horses in a race are of similar ability. Entrants are restricted by age, sex, and past record. For example, a "maiden" race is only for horses which have not won within a given period. Many races are "claiming" races in which the horses can be "claimed," that is, purchased for a specified price depending on the race. This prevents an owner from entering a very fast horse in a race with those of much lower quality. Although his horse might win the race easily, there is a good chance that the horse would be claimed for a price well below its true value. Further equalizing adjustments are made by altering the weight a horse must carry. Those with poorer past records are often allowed to carry a few pounds less. Apprentice jockeys are also permitted a weight advantage to compensate for their lack of experience. Nonetheless, after all preselection and adjustments, there still remain substantial differences between the horses' probabilities of winning in a given race.

Bettors have an enormous amount of information at their disposal to evaluate these probabilities. Detailed data about each horse are readily available in the *Daily Racing Form* and from other sources; these give full descriptions of each of its previous races, its bloodline, its performance at its latest workouts, and so on. The jockey's records and data about the track are also printed. Any relevant weather or track conditions, like rain or mud, will naturally be taken into account, as will the physical appearance of the horses just before the race.

There are also the published opinions of professional handicappers who may be able to bring more and better data under consideration and to apply their personal knowledge and professional expertise in evaluating it.

We will analyze the information contained in handicappers' selections taken from the published racing columns of three major newspapers, the *Daily Racing Form*, the *New York Post*, and the *New York Daily News*. Fourteen handicappers were available for the whole sample period. The handicappers make their picks at the time the race is made up, so they do not have access to certain relevant information, such as weather conditions, on the day of the race. Nor does a handicapper normally know the selections of other handicappers when he picks his favorites. The newspapers containing these choices are not available until the day of the race, and, in any case, in personal conversations several handicappers said they made no attempt to learn the opinions of other writers before a race.

Late scratches are also a problem. In some cases, almost all of the handicappers' choices are withdrawn before the race actually takes place. To maintain a reasonably homogeneous sample, races were eliminated when a significant number of the prerace favorites were scratched.[1] When a picked horse was scratched, the lower-ranked choices of that handicapper were moved up a position. To avoid a missing data problem for the third picks, we confined the analysis to the first and second choices only. The newspapers containing the handicappers' selections were always available to the bettors well in advance of the races.

There are two ways to bet legally on a race at Belmont: either at the track just before the race, or through New York's off-track betting system (OTB). Bets made both ways are pooled together, so the overall odds are the same.[2] However, there are differences in transactions costs borne by the two sets of bettors. On-track bettors must devote an afternoon to going physically to the track and must also pay transportation costs and an entrance fee. The OTB bettors save the time and other expenses, but, in order to reduce competition with the

[1] This was done in the following somewhat ad hoc, but efficient, manner to avoid having to gather too many data that would ultimately prove unusable. For each race, the choices of the five *Daily Racing Form* handicappers were checked to see which picked horses had been scratched. Two points were counted for each "win" pick scratched and one point for each "place." A race with a score of five or more was eliminated.

[2] The state of Connecticut also runs OTB betting on Belmont races, but does not pool its bets with those made in New York. Thus the odds generally differ somewhat between the Connecticut OTB and the New York betting market. The situation would give rise to arbitrage possibilities in the absence of transactions costs. However, in this case, given the significant cuts both New York and Connecticut take out of the betting pools, no profitable arbitrage exists.

racetracks, New York State takes an additional 5 percent (above its normal cut) from OTB winnings.

There is also an informational difference between betting at the track and betting through the OTB. On-track betting windows open about 20 minutes before the race, and, as bets flow in, the totalizer, or "tote" board, displays a running record of the current odds, updated every 30 seconds or so. Thus, an on-track bettor has a good (though not perfect) idea of the actual payoffs on his bets, while an OTB bettor must make do with the "morning line," a set of odds forecasts made by a track employee on the morning of the race.

In a fully competitive betting market, the odds on a given bet will be bid to a level that reflects the market's best estimate of the true probability of winning it. In this case, the market odds on horse i would be given by $(T-B_i)/B_i$, where B_i is the dollar amount bet on horse i, and $T = \sum_j B_j$ is the total betting pool. These odds reflect the bettors' aggregate opinion regarding horse i's probability of winning: that is, in the eyes of the market, horse i's probability is B_i/T. The posted odds differ from these market odds because of the amounts New York State and the track withdraw from the betting pool. First, a flat 17 percent is deducted from the total. Then the payoff on a $2.00 bet is calculated, and the resulting figure is rounded down to a multiple of 20 cents. This second adjustment is known as "breakage," and it makes it impossible to calculate the market's probability estimates directly from the posted odds. For example, before breakage two horses may be bet to pay $2.61 and $2.79, but after breakage both will be posted as paying $2.60. To calculate the market's true estimate of the probabilities, the actual dollar amounts before deductions were obtained from Belmont racetrack. (Throughout the paper, the terms "market odds" and "track odds" refer to these calculated probabilities, not to the posted odds.) Separate figures were available for on-track bettors and for those betting on the same races through the OTB, making it possible to study the information processing of two distinct bettor populations. (Off-track bettors are widely regarded as being less sophisticated.) A horse may be bet to "win," "place" (finish first or second), or "show" (finish first, second, or third). More exotic bets involving two or more horses like the "daily double," "perfecta," and "trifecta" are permitted for some races. In this paper we only consider the odds on bets to win.

The odds in an informationally efficient betting market fully reflect all widely held information. If some known factor, such as a horse's post position, improved its chances of winning over that implied in the odds, astute bettors would recognize this discrepancy and bet on the horse until the market odds moved into line. Free interplay of betting should bid the odds to a level where such information as the jockey, weather conditions, a horse's past performances, and also the published selections of professional handicappers is all accurately discounted. We can write this

$$P\ (\text{win} \mid \text{track odds}) = P\ (\text{win} \mid \text{track odds} + \text{all publicly held information}). \quad (1)$$

Given the track odds, it should not be possible to make more accurate probability forecasts by including further treatment of other publicly available knowledge. Conversely, one can test whether the market efficiently discounts specific information by seeing whether it can be combined with the track odds to produce significantly more accurate forecasts than can be made from the odds alone. If so, the market does not accurately discount the information. Thus our strategy in testing the efficiency of this market with respect to the subjective information contained in professional handicappers' forecasts will be to fit a probability model in which their independent contribution to forecasting accuracy can be measured.

For each race in the sample, we had the dollar amount bet on every horse, the names of the two horses each of 14 handicappers considered most likely to win and to run second, and the name of the horse which actually did win. This information was transformed into numerical data in the following way. For each horse a "track-odds" variable was created by dividing the amount bet on it by the total bet for the race. Twenty-eight zero-one dummy variables were created to summarize the handicappers' picks. For example, if handicapper 6 picked horse 3 to come in second, the variable "H6 PLACE" would be set to one for horse 3 and to zero for every other horse in that race. Finally, the outcome was expressed as a dummy: one if the horse won and zero otherwise.

How can we estimate how much information each handicapper gives us about a horse's true probability when he picks it to win? One possible approach would be simply to regress the outcome dummy on the track odds and the set of handicapper dummies. The handicappers' contributions would be measured by the estimated coefficients on their choices, and the fitted value for the dependent variable could be interpreted as the estimated probability of winning. However, it is well known that ordinary least squares is not the best choice of estimator for a probability model. Among other things, it makes no use of the fact that the fitted dependent variable represents a probability and must be between zero and one. The two estimation techniques commonly used in such problems are probit and logit. Logit is far easier computationally for large problems, so it was used.

In the multinomial logit format, each race is treated as an independent drawing from a multinomial distribution in which every horse, i, has its associated probability of winning, P_i. This probability is to be evaluated and explained by a set of independent variables, Z, in this case track odds and handicapper data. The conditional probability is assumed to be related to Z by the function

$$P_i = \frac{e^{Z_i \beta}}{\sum_{j=1}^{N} e^{Z_j \beta}}, \tag{2}$$

where Z_i is the vector of explanatory variables for horse i, β is a vector of coefficients, and N is the number of horses running. This has the advantage over the least-squares formulation that estimated values of P_i all lie between zero and one and the probabilities sum to one for a race.

If we denote the estimated probability associated with the horse that actually won the k^{th} race as P_w^k, then the joint probability of observing a whole set of outcomes in a sample of M independent races is the product of the individual win probabilities from each race,

$$L = \prod_{k=1}^{M} P_w^k. \tag{3}$$

The L is the likelihood function. Multinomial logit estimation consists of finding the parameters, β, in equation (2) such that the joint probability of the whole sample, L, is maximized. As a maximum likelihood technique, logit's large-sample theory is well developed and will allow us to make probability statements about the results. See Nerlove and Press (1973) or McFadden (1973) for a more complete treatment of multinomial logit estimation.

III. The Results

The full sample of 189 races was divided into two groups to permit in-sample estimation on 143 races and out-of-sample prediction on the remainder. The basic results did not appear to be sensitive to the manner in which the sample was split. Table 1 presents the results from the in-sample estimation. The first column represents the case in which neither the handicappers nor the bettors possess any information; β is set to zero for all of the variables, and every horse's predicted probability of winning is just $1/N$, where N is the number of horses in the race. Although this varies from race to race, the average probability of winning of a horse picked at random from the sample works out to 12.0 percent. In the next column the market odds are the only explanatory variable. The odds contribute significantly to explaining the probabilities of winning. The estimated coefficient is seven times

TABLE 1

IN-SAMPLE LOGIT ESTIMATION

	RUN					
VARIABLE	No Information	Track Odds Only	Handicappers Only		Track Odds and Handicappers	
Log likelihood	−298.8	−274.0	−264.5		−254.6	
Degrees of freedom	143	142	115		114	
Coefficients of track odds	. . .	5.198 (.73)	. . .		6.255 (1.41)	
Individual handicapper:			Win	Place	Win	Place
H1	. . .	. . .	.483 (.28)	.050 (.28)	.336 (.28)	−.039 (.29)
H2	. . .	. . .	−.211 (.31)	.210 (.26)	−.378 (.33)	.185 (.27)
H3	. . .	. . .	.194 (.28)	.119 (.27)	.066 (.28)	.054 (.27)
H4	. . .	. . .	.214 (.30)	−.359 (.30)	.038 (.31)	−.400 (.29)
H5	. . .	. . .	−.196 (.30)	.007 (.28)	−.647 (.33)	−.128 (.29)
H6	. . .	. . .	.291 (.26)	.020 (.27)	.205 (.26)	−.068 (.28)
H7	. . .	. . .	−.533 (.29)	.046 (.26)	−.560 (.30)	−.034 (.27)
H8	. . .	. . .	.073 (.26)	−.395 (.31)	−.049 (.26)	−.381 (.32)
H9	. . .	. . .	.364 (.26)	.315 (.25)	.407 (.27)	.334 (.26)
H10	. . .	. . .	.431 (.25)	.278 (.27)	.242 (.27)	.118 (.27)
H11	. . .	. . .	.226 (.29)	.629 (.25)	−.205 (.32)	.432 (.25)
H12	. . .	. . .	.060 (.35)	.632 (.25)	−.009 (.35)	.605 (.25)
H13	. . .	. . .	.029 (.29)	.243 (.25)	−.239 (.30)	.103 (.25)
H14	. . .	. . .	.458 (.24)	−.061 (.28)	.409 (.25)	−.145 (.29)

NOTE.—Standard errors are given in parentheses.

its standard error. We should expect this result, since the track favorite normally wins 30.0 percent of the time. In our sample the favorite won in 29.4 percent of the races.

The next run included information from the handicappers only. As we might expect, there is considerable correlation among the handicappers' picks, so that individually the coefficients are not highly significant. Only two of them are twice their standard errors. However, as a group they have considerable explanatory power. In the sample the horse with the highest predicted probability based only on the handicappers' information won 28.7 percent of the time. The statistical significance of these variables can be shown with a χ^2 test. In maximum-likelihood estimation the quantity 2 ($\mathscr{L}_u - \mathscr{L}_r$) is distributed as $\chi^2(k)$, where $\mathscr{L}_r$ is the maximum of the logarithm of the likelihood function estimated when there are k restrictions on the parameters, and $\mathscr{L}_u$ is the unrestricted maximum. In this case the unrestricted maximum is −264.5, and the maximum when all 28 coefficients are restricted to be zero is −298.8 from column 1. The quantity 2 ($\mathscr{L}_u - \mathscr{L}_r$) is equal to 68.6, which is highly significant; the .001 critical value for a χ^2 distribution with 28 degrees of freedom is 56.9. The handicappers as a group clearly possess relevant information.

Looking at the individual coefficients we see a relatively wide dispersion. Coefficients differ between handicappers by more than an order of magnitude, and six of them are even negative. A negative coefficient should not be interpreted here as implying that the particular handicapper is individually worse than a random prediction. In fact, some of these variables are more accurate on an individual basis than several entering with positive coefficients. The negative coefficient arises because the variable contributes no additional information beyond what is already included in the forecasts of the other handicappers.

In the last column we combine the track odds and the handicapper picks. The significance level of the odds coefficient drops, but it is still more than four times its standard error. The important question is whether adding the handicapper information significantly improves the fit over what was obtained using the odds alone. Twice the difference in the log likelihoods from columns 4 and 2 is 38.8, a figure which is significant at the 10 percent but not at the 5 percent confidence level. While the handicappers do possess considerable information, the bulk of it is discounted in the market odds produced by the bettors. We cannot reject the hypothesis that the market is fully efficient with respect to this information at the customary 5 percent confidence level.

Next the betting pools were disaggregated into on-track betting and OTB in order to study the information processing of two distinct bettor populations. There is reason to suspect that OTB bettors as a group may be quite different from on-track bettors. The OTB is just one of a number of state-run games of chance designed to appeal to the small bettor. To at least some OTB bettors, a wager on a horse race is probably little different from a lottery ticket or a chance in one

TABLE 2

COMPARISON OF ON-TRACK AND OTB BETTORS

	RUN				
VARIABLE	On-Track Odds Only	OTB Odds Only	On-Track Odds and Handi-cappers	OTB Odds and Handi-cappers	On-Track Odds and OTB Odds
Log likelihood	−272.1	−283.1	−253.1	−261.6	−272.0
Degrees of freedom	142	142	114	114	141
Coefficients:					
On-track odds	4.65	. . .	5.58	. . .	5.04
	(.62)		(1.18)		(1.07)
OTB odds	. . .	4.70	. . .	3.28	−.65
		(.82)		(1.35)	(1.44)

NOTE.—Standard errors are given in parentheses.

of a variety of legal or illegal "numbers" games. Such bettors will probably not devote much effort to actually analyzing the race. Also, since the state takes an additional 5 percent from OTB winnings, a serious horseplayer who is likely to place large bets will find it worthwhile to go to the track rather than bet via the OTB. The difference in predictive ability between these two populations shows up in the fact that OTB bettors as a group generally receive a smaller fraction of their bets back in winnings than do on-track bettors.

By disaggregating the total pools we can calculate the probabilities implied in each group's betting separately and compare their efficiency in treating handicapper data. The log likelihoods shown in the first two columns of table 2 confirm that the on-track bettors' probabilities have considerably more explanatory power than the OTB odds. This is borne out by running both together, as shown in the last column. When run along with the on-track odds, the coefficient on the OTB odds is negative and insignificant.

When the handicapper information is added to the on-track odds, the log likelihood increases by 19.0. Twice this, 38.0, is just significant at the 10 percent level for a χ^2 distribution with 28 degrees of freedom. As with the overall track odds, we are not able to reject the hypothesis that on-track bettors fully discount the handicappers' picks. For the OTB bettors, however, the situation is different. When the handicapper data are added to the OTB odds the log likelihood increases by 21.5, the coefficient on the OTB odds drops, and its standard error increases. Twice the difference in the log likelihoods is 43.0, a figure which is significant at the 5 percent level. The OTB bettors apparently do not discount the subjective handicapper infor-

TABLE 3

Out-of-Sample Estimation Results

Variable	Run: No Information	Track Odds Only	Handicappers Only	Track Odds and Handicappers
Log likelihood	−97.6	−70.3	−90.4	−68.0
Degrees of freedom	46	45	45	44
Coefficients:				
Track odds	. . .	9.56 (1.49)	. . .	9.35 (1.50)
Handicappers alone	. . .	. . .	.72 (.19)	. . .
Handicappers with odds	. . .	. . .	. . .	−.57 (.28)

Note.—Standard errors are given in parentheses.

mation as accurately as do on-track bettors. This can also be seen in the predicted probabilities. Correct treatment of the information contained in published handicapper selections would allow OTB bettors to raise their average win percentage from 27.3 percent to 30.8 percent.

Table 1 shows that the track odds discount much, perhaps all, of the handicappers' information. If the logit model we fitted has in fact captured something from the handicappers' picks which was not reflected in the odds, it should be possible to predict out of sample using the coefficients from table 1. Table 3 gives the results for out-of-sample tests on the other 46 races. Column 1 is the base "no information" run, and the second column shows the results for the track odds alone. For the run labeled "handicappers only," we created a single handicapper variable calculated as the sum of the handicapper dummies multiplied by their coefficients from column 3 of table 1. (These were the maximum-likelihood coefficients from the run with the handicapper data alone.) The estimated coefficient on this composite variable is nearly four times its standard error, so the logit model does indeed forecast out of sample.

However, when a combined handicapper variable is run along with the track odds in the last column, it only contributes negatively. Here the handicapper dummies are multiplied by the coefficients from column 4 in table 1. The estimated out-of-sample coefficient on the composite is negative and even significant. The fact that the best combination of handicapper data contributes nothing in out-of-sample prediction provides further evidence that the track odds do indeed fully discount the published choices of professional handicappers.

IV. Concluding Comments

It is a cliché that differences of opinion are what makes a horse race. To a large extent differences of opinion are also responsible for the trading activity in any speculative market. When people have different information or use different techniques to analyze information it is natural for them to differ about the probabilities of uncertain events. When a market in which they may speculate on their beliefs exists, the market price that emerges embodies an aggregate opinion which to some degree incorporates all of the information available to any of them. Over time we have accumulated a great deal of evidence that competitive speculative markets are informationally efficient with respect to widely held information. That is, free interaction of agents with different beliefs results in a market price which fully and accurately discounts the information they share in common.

Much of the information available in actual markets is subjective information, the result of analysis of factual data by professional information producers. Specialization in information production is normal when data are costly to gather and analyze, especially given the inherent economies of scale. But, for an individual market participant, evaluating subjective information is potentially quite difficult, since he will not normally know the precise factual data available to the expert or his techniques for analyzing it. Further, there will normally be a variety of conflicting opinions from different experts. This poses the question of whether a competitive market will be equally efficient in discounting widely available subjective information (thereby making it valueless).

The results of this paper indicate that the market for bets on thoroughbred horse races does quite well in discounting the subjective information contained in the published predictions of professional handicappers. They did produce valuable information, but it could not be used to improve the forecast accuracy of the market odds significantly. Apparently bettors as a group are able to place appropriate weight on the handicappers' choices in making their bets. However, there is evidence that, within the population of bettors, those who bet through the OTB appear to be less good at evaluating the subjective handicapper information than on-track bettors.

References

Ali, Mukhtar. "Probability and Utility Estimates for Racetrack Bettors." *J.P.E.* 85, no. 4 (August 1977): 803–15.

Knight, Frank H. *Risk, Uncertainty and Profit.* New York: Harper Torchbooks, 1965.

McFadden, D. "Conditional Logit Analysis of Qualitative Choice Behavior." In *Frontiers in Econometrics*, edited by Paul Zarembka. New York: Academic Press, 1973.

Nerlove, M., and Press, S. "Univariate and Multivariate Log-Linear and Logistic Models." RAND Report no. R-1306-EDA/NTH, 1973.

Rosett, Richard N. "Gambling and Rationality." *J.P.E.* 73, no. 6 (December 1965): 595–607.

Weitzman, Martin. "Utility Analysis and Group Behavior: An Empirical Study." *J.P.E.* 73, no. 1 (February 1965): 18–26.