The Economics of Wagering Markets

RAYMOND D. SAUER[1]

1. *Introduction*

THIS PAPER ANALYZES the economics of wagering markets, specifically markets in which participants take a financial position on the outcome of a sporting event such as a horse race or a football game. Although these markets are a tiny feature of most economies, they present significant opportunities for economic analysis. These stem from the fact that wagering markets are especially simple financial markets, in which the scope of the pricing problem is reduced. As a result, wagering markets can provide a clear view of pricing issues which are more complicated elsewhere.

Two basic questions form the root of the research analyzed in this paper. First, given that expected returns in the aggregate are negative, why do people trade in these markets? Second, are the prices, i.e. odds and point spreads, we observe consistent with equilibrium models of agents who make informed, optimizing choices? These questions do not have simple answers, but in attacking them the literature has deepened our understanding of the forces at work in wagering markets and elsewhere.

Early work on wagering markets followed post-war advances in expected utility theory. These studies had the

modest aim of using data on racetrack betting to examine risk-taking behavior in a simple, repeated context. The assumption, explicit in William McGlothin (1956, p. 605) and implicit elsewhere, was that "there is sufficient comparability among individuals to make an investigation of group risk-taking behavior meaningful." This approach, in which aggregate data is used to characterize the preferences of a representative agent, was a successful first step in developing and explaining basic stylized facts. However, many results from recent research cannot be explained within this framework.

The efficient markets hypothesis has proven to be a more flexible empirical tool. Since these are gambling (as opposed to capital) markets, efficiency implies that the expected rate of return to bettors has an upper bound of zero; i.e., it denies the existence of profitable wagering opportunities. Although some papers are concerned with this issue in a narrow sense, many analyses exploit features unique to wagering markets to study questions of general interest in economics and finance. The optimality of learning in a stochastic environment, the nature of the error term in pricing models, and the identification of informed trading are examples where the concept of market efficiency has been creatively employed.

To preview the subsequent discussion, there is a substantial amount of

evidence that prices set in wagering markets are, to a first approximation, efficient forecasts of outcomes. Yet there are important (and interesting) exceptions to this rule. These exceptions often highlight the effects of differences in information possessed by different agents. Ultimately, it is clear that prices formed in the wagering market aggregate scarce information from diverse sources. An asset-pricing framework which assumes agents with identical information and identical preferences is incapable of explaining these findings. Models focusing on diverse information, heterogeneous agents, and transaction costs are required to better understand these markets.

Section 2 introduces the institutions, concepts, and notation that form the basis for the subsequent discussion of empirical work. It is particularly important to be clear about what is meant by the term "efficient" in various contexts. This section uses a simple model of the wagering market as an example to clarify this issue. This example, which emphasizes differential motivations and information of market participants, helps to explain evidence that is inconsistent with a generic concept of efficiency.

Section 3 considers the motivation for gambling. It is important to at least ask the question "why do people trade in these markets" even if the answers are not completely satisfying. The representative agent model of a wagering market based on expected utility maximization with local risk preference is discussed in detail here.

The bulk of the discussion considers the empirical evidence on the nature of prices in these markets. Section 4 reviews the literature on racetrack betting, with special attention given to the favorite-long shot bias. Section 5 focuses on the point spread betting market for North American football and basketball games. Section 6 attempts to put these findings in perspective.

2. Odds, Institutions, and Efficiency

2.1 Odds and Institutions

Betting on athletic events dates to ancient times. The indigenous population of North America, for example, bet on foot races and ball games (Stewart Culin 1921). The Circus Maximus in Rome was one of many racetracks scattered across the Roman Empire (John Humphrey 1986). This arena drew crowds of 260,000 people, who were described en route to the stadium as "already in a fury of anxiety about their bets."[2]

The laws regulating wagering markets are the result of a perpetual tug of war between gambling interests and reformers.[3] In the United States, pari-mutuel markets have been the principal means of wagering on horse races, primarily because state prohibitions on bookmaking were passed early in the 20th century. Currently, bookmaking on sporting events in the U.S. is legal only in Nevada. This market has grown rapidly, in part due to reductions in the Federal excise tax on wagering from 10 percent to 0.25 percent. Between 1972 and 1995, the annual wagering on sports (exclusive of horse races) in Nevada grew more than two hundredfold, from $1 million to over $2 billion.[4]

---

[2] Tertulian, an opponent of gambling, quoted in H. A. Harris, 1972, p. 229.

[3] See Reuven Brenner and Gabriele Brenner (1990, Ch. 5); also Wray Vamplew (1976) on bookmaking in the U.K., and William Robertson (1964, pp. 195–200 and 299–301) on racetrack betting in the U.S. John Findlay (1986) and Roger Longrigg (1972) provide informative histories of gambling in the U.S., and horse racing throughout the world, respectively.

[4] These figures are from the Monthly Revenue Report issued by the Nevada Gaming Control Board. The illegal market is believed to be far greater in size; *Gaming and Wagering Business Magazine* estimated it to be $20.7 billion in 1988.

Three forms of wagering are prominent in the literature covered here: pari-mutuel odds, odds offered by bookmakers, and point spreads offered by bookmakers. The pari-mutuel system is exclusively used by racetracks in North America, France, Hong Kong, and Japan, and coexists with a bookmaking market in Australia and Great Britain. Nevada bookmakers take bets on races at major tracks, and offer odds or point spreads on team sports such as baseball, basketball, and football. The legal bookmaking market is less restricted, more extensive, and more liquid in Great Britain and Australia.

In pari-mutuel betting, a predetermined percentage is taken out of the betting pool to cover the market maker's costs, and the remainder is returned to winning bettors in proportion to their individual stake. Consider a pari-mutuel betting pool on the winner of a horse race. Let W_i be the total dollars bet on horse i to win, and $W = \Sigma_{i=1,n} W_i$ denote the total wagered on all n horses. Also, denote horse i's share of the win pool as $w_i = W_i/W$. The takeout rate is τ, and the fraction $Q = (1 - \tau)$ of the wagering pool is returned to winning bettors. The gross return to a winning \$1 bet on horse i is thus:[5]

$$R_i = QW/W_i = Q/w_i \qquad (1)$$

Let $O_i = R_i - 1$ denote the odds against horse i, which measure the net return per dollar wagered. In the pari-mutuel system, the actual payoffs to winning bettors are determined once and for all at the close of the betting period.

Bookmakers offer their patrons a set of payoffs conditional on the outcome of a given event. The payoffs offered may change during the betting period, but the payoff to each bet is determined at the time the bet is placed. The return conditional on winning is thus known at the time of the wager, in contrast to the pari-mutuel system, where heavy betting late in the period can reduce returns below acceptable levels.[6] Individuals who make bets large enough to affect the odds naturally prefer to bet with a bookmaker.

Point spread betting on football games is the staple of the Las Vegas sports betting market. In a point spread wager the payoff depends on the difference in points scored by the two opposing teams. Point spreads (PS) are typically reported as the number of points by which one team is favored to beat another. Define the actual difference in points, DP, as the points scored by the favored team less those scored by the underdog. Bets on the favorite pay off when $DP - PS > 0$, bets on the underdog pay off when $DP - PS < 0$, and all bets are refunded when $DP = PS$. The "eleven for ten rule" characterizes standard terms in the Las Vegas market, where $\tau = .1$ and successful bets return net winnings of \$1 to every \$$(1 + \tau)$ wagered.

It is easy to see how the point spread represents a price in this market. Let p represent the probability that wagers on

---

[5] This is a simplification. A typical pari-mutuel system truncates the payoff to the nearest 5 cents on the dollar, with a minimum payoff of 5 cents on the dollar. Let tr(x) truncate x to the closest integer. The payoff determined by the system just described is $R_i = \max(1.05, .05 \cdot \text{tr}(20QW/W_i))$. It is extremely rare for the minimum to be binding in win betting, although it does arise periodically in show betting (where bets on the top three finishers receive payoffs) in races with extreme favorites.

[6] There is an important exception to this rule. In Britain's betting shops it is common to make wagers at the "starting price," which is the officially determined odds prevailing in the bookmaking market at post time of the race. The unknown nature of the return is thus similar to that in the pari-mutuel system. Note that many of these bets are made before the market on the race opens; hence the odds are not yet known.

the favorite pay off; i.e. $p = \text{prob}(DP - PS) > 0$. The expected cost of an attempt to gain \$1 by betting on the favorite is the amount wagered times the probability of losing, or $\$(1 + \tau)(1 - p)$.[7] Since p falls as PS increases, the expected cost of a wager on the favored team, i.e. its price, is increasing in PS.

2.2 *Efficiency Defined*

Although there is no single definition of market efficiency, there are several concepts that are widely utilized in various empirical settings. This section introduces three concepts that will be referred to in the paper. The most restrictive of the three requires that expected returns be equal across wagering prospects.

Constant Returns. The constant expected returns model is the basis for many investigations of the efficiency of racetrack betting. If it is assumed that the takeout rate $\tau = 0$, and that all agents are fully informed, have identical risk-neutral preferences, and maximize wealth, an immediate implication of optimizing behavior is that the expected return from betting on all horses equals 1. This is easily seen. Let p_i be the probability that horse i wins, $i = 1, ..., n$. Using (1), the expected return to betting \$1 on horse i is given by $ER_i = p_i Q / w_i$.

Consider a value-weighted portfolio with weights v_i equal to the share of the win pool for each horse, $v_i = w_i$. The return to a value-weighted portfolio is important since it represents the return to the representative agent. $Q = 1$ when $\tau = 0$, thus the expected (and actual) return to the value-weighted portfolio is $\Sigma_i v_i p_i / w_i = \Sigma_i p_i = 1$.

Now suppose $ER_k > 1$. This implies that the expected return to a value-weighted portfolio of all horses excluding k is less than 1. Since all agents are equally informed and have identical preferences, all agents observe that their wealth is increased by betting on horse k. Additional betting on horse k increases w_k, reduces ER_k, and increases the expected return for all other horses. Hence $ER_k > 1$ cannot represent an equilibrium. This argument extends to all n horses. Expected returns of 1 for all n horses is thus required for equilibrium this setup. When $\tau > 0$, equilibrium assuming constant expected returns requires that the return to any value weighted portfolio of all horses is $Q = (1 - \tau) < 1$.

Subjective probabilities refer to the betting market's estimate of each horse's chance of winning the race. This term has its root in the constant returns model. Let the subjective probability that horse i wins the race be s_i, and satisfy the equation $Q s_i / w_i = Q$. That is, people bet such that they expect the return from each horse to equal Q. This assumption is the basis for using betting shares from the win pool, w_i to measure s_i. Further, assume that $s_i = p_i$. If the constant returns model is true, then $w_i = p_i$ —horse i's share of the win pool measures its true probability of winning.

Empirical analyses which examine the rate of return to wagers across different classes of prospects (categories of horses) implicitly test a version of the constant returns model where the hypothesized return is $Q < 1$. The theoretical basis for such tests is weak, however, since the motivation for wagering in the presence of a positive takeout rate must be addressed. A preference for gambling must be introduced, and this preference must be "just right" to generate constant returns when $\tau > 0$. Local risk-loving in a utility of wealth framework will not generate constant

---

[7] We can assume prob(DP=PS)=0 with no loss of generality, since bets are refunded when DP=PS.

returns in equilibrium, as shown by Richard Quandt (1986) and discussed in Section 3.

The Absence of a Profit Opportunity and Efficient Point Spreads. Many papers calculate the profitability of wagering rules. These tests are motivated by the idea that gambling markets can't simultaneously yield profits for both market makers and the aggregate betting public. Since profits for market makers must be non-negative, the existence of profitable wagering rules is, at a minimum, inconsistent with informed behavior by bettors. A simple test of the absence of profit opportunities concept of efficiency examines returns to following a particular wagering rule. Rejection of the hypothesis that $ER_k \leq 1$ would provide statistical evidence against the absence of profit opportunities concept of efficiency.

In point spread betting, the absence of a profit opportunity restricts the probability of making a successful point spread wager. With zero takeout, payoffs are at odds of 1 to 1, and this concept of efficiency implies (i) that the probability of a successful wager is .5, and (ii) that PS is the median of the DP distribution.

Takeout weakens the implication. Again, let p be the probability that bets on the favorite are winners, with odds set at 1 to $(1 + \tau)$. The absence of a profit opportunity requires that the expected return to a bet on the favorite be no greater than the cost of a wager, i.e., that $p(1 + (1 + \tau)) \leq (1 + \tau)$. The same requirement must hold for wagers on the underdog, which restricts $(1 - p)(1 + (1 + \tau)) \leq (1 + \tau)$. Combining these conditions, the absence of a profit opportunity requires that[8]

$$1/(2 + \tau) \leq p \leq (1 + \tau)/(2 + \tau) \qquad (2)$$

[8] This implication obviously holds for any two-sided condition.

which is less restrictive than the median condition when $\tau > 0$.

Equilibrium Pricing Functions. The equilibrium of an explicit model can be viewed as characterization of efficient prices. Equilibrium pricing functions depend importantly on the information structure and behavioral assertions imposed on market participants. Examples include Quandt (1986) and Hyun Song Shin (1992), two leading models of the favorite-long shot bias.

Pricing functions derived from explicit models are important because they fill the gap between the two generic concepts of efficiency discussed above. The absence of profit opportunities is too weak to explain some persistent features of the data which are inconsistent with constant returns. The favorite-long shot bias is one such example.

If one insists that the efficiency concept be restricted to models in which all agents are fully informed and share common preferences, then equilibrium pricing functions derived under alternative assumptions would not qualify as efficient. I believe a broader definition of efficiency is warranted. Suppose we consider a model in which uninformed agents have a taste for betting on horses, and a group of informed bettors maximizes a concave function of wealth. Let the model generate an equilibrium in which both sets of agents are at least as well off pursuing this activity than the next best alternative, with no profit opportunity for new entrants. Basic models of exchange characterize such outcomes as efficient. However, the constant returns concept of financial market efficiency would exclude this pricing function if, for instance, the expected return to an informed bettor differed across categories of horses. What is important is that we have testable models which predict observed phenomena. For this reason, it is useful

to adopt a broad definition of efficiency which includes equilibrium pricing functions from well-posed models of the wagering market.

2.3 *Efficiency in Markets with Heterogeneous Agents: An Example*

Wagering markets are fascinating in part because cheering for competitors is linked to loyalty, passion, and other "irrational" concepts. These traits surely affect the behavior of some market participants. But since prices (odds and point spreads in this context) are determined at the margin, it is not a given that these traits will affect market prices. It is the marginal agent, about whom little is known from psychological studies, that is crucial to pricing.

The following example illustrates some important points related to market efficiency, and is intended to provide a unifying basis for interpreting much of the empirical work that follows.

A Horse Race with "Irrational" Betting. Consider a horse race contested by three horses of equal ability but different color. Index each horse by h_i and denote the probability that h_i wins by $p_i = 1/3, i = 1, 2, 3$. Further assume that each of N uninformed bettors wagers \$1 on the horse of their favorite color. Color preferences are distributed evenly so that \$$N/3$ is bet on each horse. For simplicity, assume there is no takeout so that all money bet is distributed to the winning bettors.

In this example, each bet is a fair one since the expected return to a \$1 bet on h_i is $ER_i = p_i N/(N/3) = 1$. Suppose an economist using the constant returns concept examines data generated by a series of these races. The economist would find this market to be efficient,[9]

despite the fact that the basis for decision making has nothing to do with the outcome of the race.

Now change the scenario. Suppose that in each race, horse k is randomly chosen to be more likely to win. Specifically, let $\text{prob}(h_k = h_i) = 1/3, i = 1, 2, 3$; and $p_k = p = 2/3, p_j = 1/6, j \neq k$. The uninformed bettors place wagers as before. The efficiency conclusion is the same if the economist, like the bettors, does not know which horse has been chosen more likely to win.

Informed Betting and Efficiency. Now introduce a single informed agent with knowledge of the correct set of probabilities and of the betting by others, and assume his objective is to maximize the expected profit from betting. Let s_i^u represent the subjective probabilities (i.e. betting shares) of the uninformed bettors. Rufus Isaacs (1953) studied this problem and showed that the optimal wager is always positive when $Qp_i/s_i^u > 1$, which is clearly the case for h_k.[10] Denoting the amount wagered on h_k by x, expected profits are

$$E\Pi = p\,\frac{(N + x)}{(N/3 + x)}\,x - x$$

where $p = 2/3$ by assumption. In this case, the expected profit maximizing wager is $x* = N/3$. If the informed agent bets optimally, the total amount wagered is now $4N/3$, with $2N/3$ bet on h_k and $N/3$ bet on each of the other horses. The expected return to betting \$1 on h_k is now $ER_k = (2/3)(4N/3)/(2N/3) = 4/3 > 1$. Our economist will discover this eventually (h_k is now always favored and thus easily identified), and correctly pronounce this market to be inefficient.

As the number of informed agents is increased, the betting on h_k increases and expected returns to bets on all

---

[9] Which it is, since there is no incentive to alter choices.

[10] See L.C. MacLean, W.T. Ziemba, and G. Blazenko (1992 for an analysis incorporating risk aversion).

horses converge towards 1 again. William Hurley and Lawrence McDonough (1995) derive the non-cooperative Nash solution for the case of $M \geq 1$ informed agents. Since informed agents will only consider betting on h_k in this example, the problem for informed agent m is to maximize

$$E\Pi_m = p\frac{(N + \Sigma x)}{(N/3 + \Sigma x)}x_m - x_m$$

where Σx represents the betting of all M informed agents.

Solving this problem assuming informed agents are identical yields a solution for $x_m = x* = x*(N, M, p, \sigma_k)$.[11] Total betting by informed agents $Mx*$ is increasing in M. As a result, the expected return to betting h_k, $ER_k = p(N + Mx*)/(N/3 + Mx*)$ asymptotically approaches 1.

Table 1 presents the optimal wager $x*$ and total wagering $Mx*$ by informed agents as M increases. The rightmost columns of Table 1 present the expected return to betting h_k and the aggregate subjective probability of all agents. The calculations are made for the case described above, with N set to 100. This enables the wagering magnitudes to be interpreted as percentages relative to the volume of uninformed betting.

As indicated by the case in which $M = 1$, the difference between 1 and ER_i can be substantial when the number of informed agents is small.[12] Nevertheless, it does not take many informed agents to drive expected returns close to 1. It is worth noting also that the uninformed agents are treated somewhat implausibly in this example, since it is profitable for them to switch their bets from their preferred color to the betting favorite. Such behavior increases the tendency for expected returns to converge to 1.

Discussion. This example illustrates several points that are sometimes obscured in the debate on financial market efficiency. First, all agents need not be informed for a market to be efficient. In this example, a conclusion of market efficiency is reached with either zero (which is fortuitous) or a small number of informed agents. Second, a finding of efficiency does not imply the absence of irrational behavior. By irrational, I mean simply that some participants base their wagers on irrelevant data, as in this example. This line of thinking has important implications. An example is how we interpret survey evidence on expectations and behavior. Survey evidence documenting biased expectations does not necessarily imply that market prices will be similarly biased.

Third, it is the marginal agent that is crucial to efficient pricing. Indeed, once the information asymmetry is introduced, competition among marginal agents is required for efficiency. Representative agent models in which all bettors have access to identical public information are quite restrictive, and will

---

[11] The solution for x_m proceeds as before, using $\partial x_j/\partial x_m = 0, j = 1, 2, ..., M$.

The resulting quadratic equation

$$x^2 + x(2N/3)\left(\frac{1}{M} - \frac{(M-1)}{M^2} \cdot \frac{p}{1-p}\right)$$
$$+ \frac{(N/3)^2}{M^2} - \frac{p}{1-p} \cdot \frac{2N^2/9}{M^2} = 0$$

yields the solution $x*(N, M, p, s_k)$:

$$x* = (N/3)\left((M-1)z - \frac{1}{M}\right) + \frac{1}{2M}$$

$$\sqrt{(4N^2/9)\left(\frac{1}{M} - z(M-1)\right)^2 - 4\left(\frac{N^2/9}{M^2} - z(N/3)(2N/3)\right)}$$

where $z = p/(1-p)M^2$.

[12] This requires of course that there be a substantial difference between the true probabilities and subjective probabilities of the uninformed agents, as in this example.

TABLE 1
INFORMED AGENTS, OPTIMAL WAGER SIZE AND EXPECTED RETURNS

Number of Informed	Optimal Wager	Total Informed	Expected Return	Probability Subjective
M	x°	Mx°	pR_k	R_k^{-1}
0	—	0.0	2.000	0.333
1	33.3	33.3	1.333	0.500
2	28.9	57.7	1.155	0.577
3	23.3	69.9	1.097	0.608
4	19.2	76.8	1.070	0.623
5	16.2	81.1	1.055	0.632
6	14.0	84.1	1.045	0.638
7	12.3	86.3	1.038	0.642
8	11.0	87.9	1.033	0.645
9	9.9	89.2	1.029	0.648
10	9.0	90.3	1.026	0.650
15	6.2	93.4	1.017	0.655
20	4.8	95.1	1.013	0.658
25	3.8	96.0	1.010	0.660
50	2.0	98.0	1.005	0.663

Notes: The calculations assume $p_k = 2/3$ and $s_k = 1/3$. N is set to 100 so that the wagers listed can be interpreted as percentages relative to the size of the uninformed betting pool.

be rejected when private information is important and limited to small groups. Horse race betting is just such a case. Many results in this literature which depart from the constant expected returns standard are related to these factors.

Finally, it is well known but worth reiterating that efficiency is not a stand-alone concept. Efficient prices embody properties that are implied by a given model, and are therefore dependent on the behavioral assertions, constraints, and information structure that characterize the model. Hence, the source of error when efficiency is rejected is by no means immediately obvious. My own view is that the generic efficient markets hypothesis is a very useful benchmark. Its generality is at once a great strength, since it can be widely applied, and a great weakness, since it will be rejected in settings where idiosyncratic conditions are important. But rejections of efficiency don't just highlight limitations of the basic model; they must be

studied carefully, for it is these cases which add the most to our understanding of the forces that create market prices.

3. *Models of Gambling Behavior and Gambling Markets*

3.1 *Utility-of-Wealth Models of Gambling*

Models of gambling based on expected utility date back to Daniel Bernoulli's famous solution to the St. Petersburg paradox. Bernoulli posited that individuals value a gamble using a probability-weighted utility function instead of the standard mathematical expectation (see Kenneth Arrow 1952, pp. 420–21; and Mark Machina 1987, p. 122–23).[13] Since the solution requires

[13] Arrow (1952) provides a particularly comprehensive discussion of axiomatic foundations for and methodological objections to expected utility theory. Machina (1987) provides a helpful introduction to non-expected utility models, in which

marginal utility to be decreasing in wealth, it is incapable of explaining the acceptance of fair gambles.[14] That people accept some fair (and unfair) gambles implies that either they obtain pleasure from gambling or that marginal utility is not universally declining.

The classic paper of Milton Friedman and Leonard Savage (1948) assumes the latter. The Friedman–Savage explanation is based on the assumption of a convex (increasing marginal utility) segment in the middle range of an otherwise concave utility function. Hence, individuals in the first concave segment are predicted to purchase low probability, high payoff gambles that reach well into the convex segment, while simultaneously insuring against wealth-decreasing risks. The concave segment at higher levels of wealth limits the gains from gambling and hence the size of the prize for which people are willing to gamble. Several facts about lotteries—that prizes are typically not "winner take all" but multiple in nature, and that poor people tend to play the lottery and rich ones do not—are consistent with Friedman and Savage.[15]

Jack Hirshleifer (1966) pointed out a critical flaw in this approach. While the purchase of low probability, high payoff gambles by the poor fits the data, the application to middle-income individuals implies that they are "plungers of an extreme sort."[16] Strictly taken, the model implies that middle-income individuals would gamble their way into either the poor or rich sections of the income distribution.

The approach of Harry Markowitz (1952) avoids this pitfall. The Markowitz model places the convex segment of utility at current wealth, and thereby treats gambling as exploitation of local risk preference. The local nature of this model has the advantage of allowing all segments of the income distribution to make rational gambles. Nevertheless, like Friedman and Savage, Markowitz relies exclusively on the curvature of a utility of wealth function to explain gambling behavior.[17]

3.2 Local Risk Preference in a Representative Agent Model of the Betting Market

Early models of racetrack betting—notably Martin Weitzman (1965), Mukhtar Ali (1977), and Quandt (1986)—employ the local risk preference assumption. Weitzman's model is the first in which prices observed in the betting market are given a coherent equilibrium interpretation. Weitzman's data are a sample of over 12,000 races run at

---

the objective function is non-linear in the probabilities. These latter models have not been adapted to empirical studies of gambling, with the exception of John Conlisk (1993), discussed below.

[14] Declining marginal utility implies that the utility of current wealth exceeds the expected utility of any fair gamble. Hence, it can explain why people are unwilling to pay more than a small amount for a gamble with infinite expected value, i.e., the St. Petersburg game. This game tosses a fair coin n times and pays the gambler 2^n dollars, where n is the first time a head appears.

[15] The data unequivocally show that the percent of income allocated to lottery games declines sharply with income. Actual expenditure on lotteries varies in diverse ways across income and socio-economic classes. See Charles Clotfelter and Phillip Cook (1991, pp.95–104).

[16] For individuals in the middle range, the expected utility of a small gamble spanning a tiny segment of wealth will be dominated by that of a large gamble which reaches far into the convex segment of the utility function. These individuals are thus predicted to make much larger gambles (of higher probability) than those with lower wealth.

[17] Positing a wiggle in an otherwise smoothly declining marginal utility function strikes many as being equally ad hoc as positing a taste for gambling. Of course there are ways to generate such functions. Applebaum and Katz (1981) create a Friedman–Savage function based on an asset market argument, in which low wealth is a barrier to earning a high rate of return. An increasing rate of return in the middle range of wealth generates convexity in a utility function which is concave in the lower and upper wealth regions.

New York tracks over a 10–year period. Let R_i be the return per \$ conditional on winning and $\hat{p}_i$ the sample estimate of the probability of winning associated with that return. Weitzman found that a rectangular hyperbola and a modified hyperbolic form estimated by weighted least squares achieved a remarkably good fit to the $(R_i, \hat{p}_i)$ pairs (R^2 exceeded .98 in both cases). The latter function was consistent with local risk preference.[18]

Weitzman's model of the wagering market posits a representative agent, Mr. Avmart, who maximizes expected utility given a local preference for risk. Assume that Avmart bets \$b and let $m_i = R_i * b$. The points $(m_i, \hat{p}_i)$ on the hyperbolic curve represent an equilibrium relation between money prizes and the probability of winning them, and the representative agent is necessarily indifferent between points on this curve. Were Avmart not indifferent to points on this curve, more betting would take place on preferred combinations reducing their return (and raising others) until the agent was indifferent between the resulting combinations. The resulting curve then would be what we observe. Similarly, given a momentary deviation of an $(m, \hat{p})$ combination off the curve, "Avmarts will line up at the ticket window to drive the probability versus return curve back to the proper shape (Weitzman 1965, p. 24)." Weitzman and Ali (1977) employ the observed relation between probability and return to recover the (local) shape of the representative agent's utility function.

Quandt (1986) took the argument one step further. Weitzman's data, along with other data before and since, exhibited a phenomenon now referred to as "the favorite-long shot bias."[19] In short, this bias exists when the expected return to betting favorites exceeds that of betting long shots. Quandt showed that, assuming local risk preference in utility, a necessary condition for equilibrium in the wagering market is that favorites yield higher expected returns than long shots.

The favorite-long shot bias can be characterized by the condition $p_i/w_i > p_j/w_j$ when $w_i > w_j.$ Assume that agents know the true probabilities, and recall that $p_i = w_i, i = 1, ..., n,$ implies $ER_i = Q$ for each horse. Let $p_i = w_i$ hold for every horse. This cannot be an equilibrium if Avmart exhibits local risk preference. Since the longest shot is the riskiest (a preferred attribute) and its expected return is the same as all other horses, Avmart would prefer to bet on the longest shot when $p_i = w_i$ for all horses. This activity increases w_i for long shots and reduces it for others. Quandt shows that $p_i/w_i > p_j/w_j$ implies $p_i > p_j$ when Avmart exhibits local risk preference; that is, higher expected returns accrue to higher probability horses.

These papers represent the high-water mark of the representative agent model of racetrack betting. They show how aggregate returns in the betting market can be viewed as an equilibrium outcome of optimizing trades by a representative agent with local risk preference. Subsequent empirical work reveals the limitations of this model, however. First, not all markets exhibit the favorite-long shot bias. Second, there is substantial evidence of differential returns earned by informed

---

[18] The rectangular hyperbola $\rho = A/R$ of course generates a constant expected return of A when subjective probabilities equal their empirical counterparts. Weitzman's estimate of A was .8545, trivially higher than Q (.85) at New York tracks during this period. Although the modified hyperbolic form fit the data better, Weitzman notes that constant expected returns "is a good first-order approximation" to the data.

[19] These empirical studies and other attempts at explaining this phenomenon are discussed in Section 4.

market participants. These findings, discussed in Section 4, indicate that there is more to racetrack betting than is captured by this model. To fully understand these features requires more sophisticated modeling of the behavior of various market participants.

3.3 *Critiques and Alternatives to Utility-of-Wealth Models of Gambling*

Paul Samuelson (1952), Hirshleifer (1966), and others have argued that gambling is not wealth-oriented in the Friedman–Savage sense.[20] Hirshleifer's observation that most gambling involves repetitious wagers with small stakes illustrates the point. Most bets that people make (say $10 to win on a horse at odds of 3–1) are simply incapable of generating wealth changes of any consequence. But if gambling is motivated by something other than convex utility, what alternatives can shed light on the economics of gambling markets?

Richard Thaler and William Ziemba (1988) argue that behavioral propositions such as mental accounting and prospect theory are consistent with some empirical features of wagering markets. This is true, and future papers on wagering markets will certainly be motivated by the behavioral approach. The potential lack of generality in this approach is troublesome however. One can think of numerous stories to explain things not predicted by an economic model, perhaps one or more stories for every anomaly. Extensions of the basic economic model may ultimately be a more productive line of inquiry.

One promising extension is the model of John Conlisk (1993), which focuses on the pleasures of gambling.[21] In doing so, he is forced to confront the notion that a direct "association of utilities with gambles can lead to a theory . . . without scientific content," an issue of obvious importance. Conlisk appends what he calls "a tiny utility of gambling" to an otherwise standard expected utility model. He considers fair gambles in which amount G is won with probability p, and amount L is lost with probability $1–p$. Fairness implies that a given gamble can be summarized by the pair (G, p), where $pG = (1 − p)L$.

Gambles are assumed to add to utility in the following way. Let $U(W)$ represent a standard utility of wealth function, which is bounded and exhibits decreasing absolute risk aversion. The preference function is

$$E(G, p, W) = pU(W + G)$$
$$+ (1 − p)U[W − pG/(1 − p)] + \varepsilon V(G, p)$$

which is an expected utility function augmented with an additional utility of gambling, $\varepsilon V(G, p)$, where ε is a nonnegative scale parameter. $V(G, p)$ has the properties $V(0, p) = 0$, $V_1(G, p) > 0$, $V_{11}(G, p) < 0$, $V(G, 0) = 0$, and $V_2(G, p) > 0$ for $G > 0$.

These restrictions on the taste for gambling enable Conlisk to derive testable implications from the model. Assuming that ε is sufficiently small, he shows that (1) there is a limit to the size of an acceptable gambling prospect; i.e. there is a sense in which small gambles are preferred to large ones; (2) there is a uniquely preferred gamble size; and (3) both of these magnitudes increase

---

[20] In discussing the issue, Samuelson (1952, p. 671) writes "Warning: what constitutes a prize is a tricky concept. When I go to a casino, I go not alone for the dollar prizes but also for the pleasures of gaming—for the soft lights and the sweet music. In such cases the X's (prizes) should be complicated vectors embodying *all* these elements."

[21] Vernon Smith (1971) also tinkered with an expected utility model modified to include a direct utility of gambling. Smith's primary focus, however, was on how different payoff structures (odds vs. point spreads) affected expected utility in the absence of a taste for gambling. This topic is taken up in Section 3.4.

with wealth.[22] In particular, the model satisfies one essential requirement of Friedman and Savage—it predicts the acceptance of small gambles and the purchase of insurance when risks are large. Intuitively, the basis for these implications is that for small gambles, the utility-of-gambling effect is first-order small, whereas the risk aversion effect is second-order small. Finally, note that individuals of various wealth levels will accept fair gambles in this framework, an implication consistent with Markowitz' normalization of the convex segment of utility.

On the empirical side, Conlisk makes the case that experimental evidence of "excess risk-seeking" in prior studies is consistent with the implications of his model of gambling. Nevertheless, in its current form, the restrictions required for empirical testing are not present. In effect, with the curvature properties of two additive functions loosely specified, the variety of admissible choices is too extensive. Conlisk (p. 262) reports, however, that Alexander Pollatsek and Amos Tversky (1970) "present an axiomatic development that, specialized to the prospect (G, p)" can yield a specific functional form for $V(G, p)$. Restrictions on $V(G, p)$ would give the model additional punch.

Conlisk's model represents a significant advance over the ad hoc specifications of convexity in earlier models. Although additional elements are required for an empirically tractable model of the wagering market, this is a serious treatment of the motivation for gambling and its implications, and deserves further exploration.

3.4 *The Structure of Gambling Markets: Why Is There No Odds Market for Football Games?*

Why do Las Vegas bookmakers offer point spread betting at fixed odds of

10–11, but choose not to make an odds market on the winner of the game? The layman's answer is that the point spread wager is more exciting, particularly when a lopsided contest is scheduled. Gilbert Bassett, Jr. (1981) provides an elegant alternative explanation in which bettors have no innate preference for one type of wager over the other. Bettors are assumed to wager an amount which is a nondecreasing function of their perceived expected return. The bookmaker is assumed to set odds or spread such that a profit is realized independent of the game's outcome. To make a market in this setting obviously requires that bettors have disparate probability beliefs. These beliefs are modeled by distributions over the point difference which differ only by a location parameter.

The critical element in the Bassett model is the bookmaker's takeout. All bettors are charged a fixed fee, which, coupled with the odds, generates the price of a contingent claim to $1. In the absence of this fee, betting volume is the same in each betting game. With a positive fee, point spreads generate more betting when the distribution of point differences is symmetric and unimodal. The bookmaker's takeout drives a wedge between the two sides of the market. In Bassett's model this wedge is minimized when the odds are equalized, as in the point spread game.

The model of Bill Woodland and Linda Woodland (1991) emphasizes the role of risk aversion in determining the preferred betting game. Woodland and Woodland identify two problems with the Bassett model. First, the fixed takeout assumed by Bassett raises the price of a contingent claim for the long odds bettor by a much greater percentage than the short odds bettor. This effect can be very pronounced at long odds, for example, when the fee is 5 cents and

---

[22] If ε exceeds a critical level, then every fair prospect will be accepted.

the price of obtaining $1 by betting on the long shot is, say, 10 cents or less. Second, Woodland and Woodland are skeptical of the implication in Bassett that odds betting on football games would be just as likely to be observed as point spread betting in the absence of takeout. Something other than takeout (or perhaps in addition to it) may be at work.

Woodland and Woodland show that the optimal wager size of a risk-averse, U(W) maximizing individual is greater when the odds are close to even, as in the point spread game, than when they are disparate, as in betting on the winner.[23] Their model's assumption of risk-averse, utility of wealth maximizing agents provides the basis for arguing that spread betting will be preferred to odds at equilibrium values. Further, they argue that the standard assumption of local risk preference is inconsistent with point spread betting: since odds betting on the winner is inherently riskier, local risk preference would imply that betting on the winner dominates betting against the spread. This is a useful observation, but it does not imply that concavity in U(W) is the culprit in driving out an odds market for football games.

While these are interesting models, neither has solved the riddle. There may be better attacks on the question than comparisons of wagering volume, since the odds market in football probably fails for reasons that are independent of the success of point spreads. Note that several forms of wagering are offered on all of the sports discussed in this paper. Thus, the existence of one

form of wagering does not preclude the existence of another form. Consider also the bookmaking market on horse races in the U.K. In lopsided races with an extreme favorite, many bookmakers do not offer odds against the favorite, while simultaneously offering odds against the long shot horses in various forms. This suggests that it is not profitable to make a market on extreme favorites, in contrast to the long shots. This situation parallels that of a heavy favorite in a noncompetitive football game. If there is a solid economic (as opposed to a taste-based) explanation for the dominance of point spreads over odds, this explanation should explain why bookmakers make a market on long shots but not heavy favorites in lopsided races.

4. *Empirical Analyses of Racetrack Betting*

4.1 *The Betting Market and the Probability of Winning a Race*

Psychologists Richard Griffith (1949) and McGlothlin (1956) were the first to use the racetrack betting market as a vehicle for studying behavior under uncertainty. These studies used the win pool shares, $w_i = W_i/W$, to evaluate the ability of market participants to discern small differences in the probabilities of outcomes.

Consider an *n* horse race, where $p_i, (i = 1, 2, …, n)$ is the probability that horse *i* wins. The net share of the win pool after takeout is given by

$$\pi_i = W_i/QW = w_i/Q = R_i^{-1} \qquad (1')$$

which is the inverse of equation (1). π_i can be viewed as the price of obtaining a claim to $1 in the event that horse *i* wins the race. As before, R_i is the gross return to $1 bet on horse *i*, and the odds $O_i = R_i - 1$ measure the net return.

Griffith (1949) grouped horses from

---

[23] It is assumed that the perceived net return from the wager is positive. Of course this cannot be true in the aggregate, so this model fails to satisfy basic notions of rationality. Woodland and Woodland do not address the question of why individuals with their assumed objective would persist in this activity when takeout assures that on average they lose money.

1386 races into 11 categories according to their odds. Horses with $O_i \leq 1$ comprised the lowest-odds group, and horses with $O_i > 15$ comprised the longest odds group. Let H_k be the number of horses in class k, and N_k be the number of winning horses in this group. The expected number of winners in group k is $N_k^e = p_k H_k$. Griffith compares H_k with N_k/w_k for a sample of 1386 races. Clearly, $H_k \approx N_k/w_k$ only to the extent that $w_k \approx p_k$. Griffith (pp. 292–93) finds the "near congruence of" N_k/w_k and H_k "remarkable." This congruence implies that $w_k \approx p_k$, which justifies the use of win pool shares as estimates of subjective probability.

Arthur Hoerl and Herbert Fallin's (1974) study of the win pool is especially clear on this point. Hoerl and Fallin categorized races by the number of contestants, which ranged from 5 to 12 horses. They ranked horses within each race by their track odds, and compared the average subjective probability with the observed proportion of victories for each ranking. These comparisons are presented in Table 2. The subjective and objective probabilities move together quite closely; both decline in lock-step fashion as one moves from the highest ranked to the lowest ranked horse in a race. Table 3 presents Hoerl and Fallin's calculation of the average finishing position for the horses in these races. This number declines monotonically with the rank of the horse in all eight race categories. Like Griffith, Hoerl and Fallin (1974, p. 230) concluded that the betting market "demonstrates that individuals with incentive can on the average successfully discriminate small differences" in the likelihood of outcomes.

4.2 The Favorite–Long Shot Bias

Evidence. But Griffith also found that for horses with the lowest odds (favor-

ites), N_k/w_k was slightly higher than H_k, implying that $w_k < p_k$ for these classes. For horses with higher odds (long shots), Griffith observed the opposite, implying that $w_k > p_k$ in this range. Hence, the win pool shares were lower than win frequencies for favorites, and slightly higher than win frequencies for long shots. This tendency is also clear in Hoerl and Fallin's data, which can be discerned by looking carefully at the most and least favored horses in Table 2.[24]

McGlothlin (1956) framed the question in terms of expected returns, and carried out parametric statistical tests. McGlothlin classified horses from his sample of 9248 races into nine groups according to the odds of each horse. In the absence of takeout, the expected profit of a \$1 wager on horses in group k is $E_k = p_k/w_k - 1$.[25] McGlothlin used the ratio N_k/H_k to measure p_k, and calculated w_k using the published odds of each horse.

The benchmark value of E_k is zero, which is realized when $p_k = w_k$. McGlothlin calculated E_k to be .08 for the two groups of horses with the lowest odds (3–1 and less), and E_k to be roughly -.10 for the three groups of horses with the highest odds (8–1 and higher). These figures were 2–4 times their standard errors. Thus, expected returns on favorites exceeded those of long shots in McGlothlin's sample of races, a finding which has been repeated in many samples since.

Wayne Snyder (1978) surveyed six

[24] Although the departures from equality are small, the ordering of the differences between objective and subjective probabilities in the tails is systematic. Statistical tests on the differences as a group provided no evidence against the null hypothesis, although Hoerl and Fallin noted the tendency to over-predict in the right tail.

[25] Since $Q < 1$, the true expected profit of a \$1 wager is $E_k' = p_k/\pi_k - 1$. Note that $E_k' - E_k = \tau$ when $p_k = w_k$, so McGlothlin's E_k is essentially a normalization of E_k' to 0.

TABLE 2
COMPARISON OF SUBJECTIVE PROBABILITIES AND ACTUAL WINNING FREQUENCIES
BY ODDS RANK OF HORSE

No. of Entries	No. of Races		Odds Rank of Horse											
			1	2	3	4	5	6	7	8	9	10	11	12
5	69	Subj. prob.	.42	.25	.17	.11	.06							
		Obs. freq.*	.41	.30	.20	.07	.03							
6	181	Subj. prob.	.36	.23	.17	.12	.08	.04						
		Obs. freq.	.43	.21	.20	.11	.03	.02						
7	312	Subj. prob.	.33	.22	.16	.12	.09	.06	.03					
		Obs. freq.	.34	.21	.16	.12	.08	.08	.02					
8	352	Subj. prob.	.31	.20	.15	.12	.09	.06	.04	.03				
		Obs. freq.	.33	.25	.13	.09	.07	.06	.04	.02				
9	283	Subj. prob.	.30	.20	.05	.11	.09	.06	.05	.03	.02			
		Obs. freq.	.35	.15	.17	.13	.08	.06	.02	.01	.02			
10	241	Subj. prob.	.29	.19	.14	.11	.08	.06	.05	.03	.02	.02		
		Obs. freq.	.31	.17	.16	.10	.07	.07	.06	.04	.02	.01		
11	154	Subj. prob.	.27	.18	.14	.11	.08	.07	.05	.04	.03	.02	.01	
		Obs. freq.	.27	.18	.19	.08	.05	.05	.05	.05	.04	.04	.01	
12	233	Subj. prob.	.26	.17	.13	.10	.08	.07	.05	.04	.03	.02	.02	.01
		Obs. freq.	.28	.14	.17	.12	.10	.06	.02	.05	.03	.03	.01	.00

Source: Hoerl and Fallin (1974); data are from all 1,825 races run at Aqueduct and Belmont Park (NY) in 1970.
* Observed frequency.

TABLE 3
MEAN ORDER OF FINISH BY ODDS RANK OF HORSE

No. of Entries	No. of Races	Odds Rank of Horse											
		1	2	3	4	5	6	7	8	9	10	11	12
5	69	2.1	2.4	2.9	3.4	4.1							
6	181	2.2	2.9	3.2	3.6	4.2	4.9						
7	312	2.8	3.2	3.7	4.0	4.3	4.6	5.4					
8	352	2.8	3.2	3.9	4.2	4.7	5.1	5.7	6.4				
9	283	3.1	3.6	4.1	4.6	5.1	5.3	6.0	6.4	7.1			
10	241	3.1	4.0	4.3	5.1	5.3	5.6	6.2	6.5	7.0	7.9		
11	154	3.8	4.0	4.7	5.2	5.7	5.8	6.3	6.9	7.2	7.8	8.5	
12	233	3.9	4.6	5.1	5.4	6.0	6.2	6.7	7.2	7.6	7.7	8.7	9.1

Source: Hoerl and Fallin (1974); data are from all 1,825 races run at Aqueduct and Belmont Park (NY) in 1970.

studies (including Griffith, McGlothlin, and Weitzman) encompassing 50,000 races in North America, each documenting the favorite-long shot bias. Snyder aggregated the rates of return for various odds categories, with the takeout added back, as in McGlothlin. The pre-takeout rate of return varies
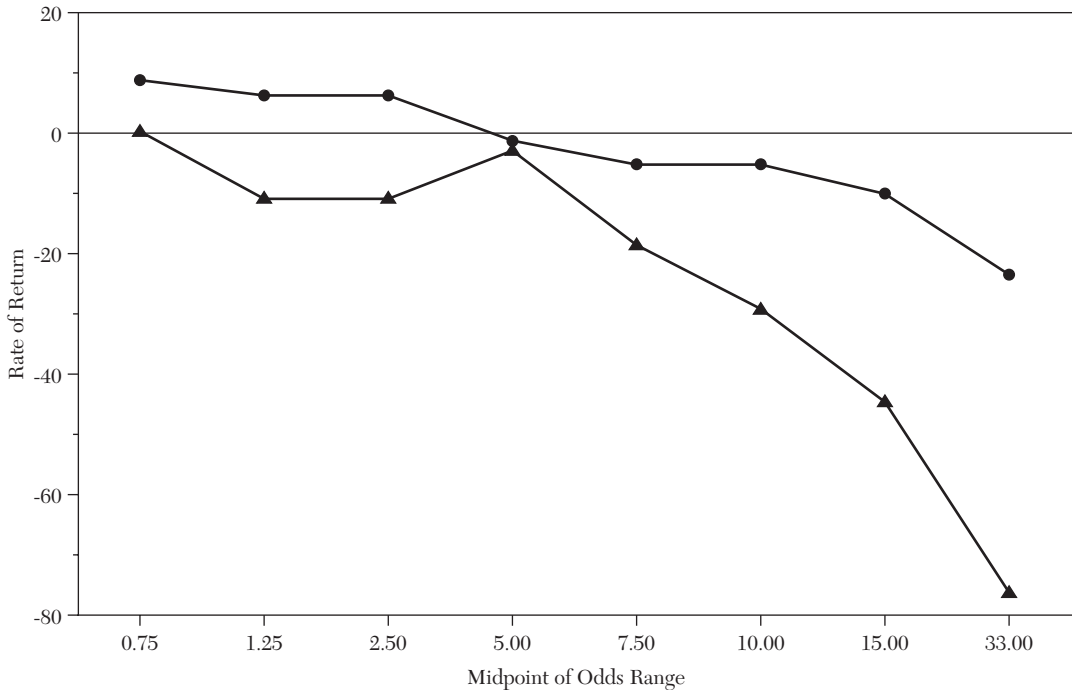
Figure 1. The Favorite-Longshot Bias in the U.S. Pari-mutuel (Snyder) and U.K. Bookmaking (Dowie) Markets
Note: Snyder's are pari-mutuel returns with the takeout added back; hence the norm is zero. Dowie's are pre-tax returns at SP odds. Hence the norm for Dowie is less than zero to account for the costs of bookmaking.

from 9.1 percent for odds-on horses ($O_i < 1$), to –23.7 percent for horses with the highest odds (33–1 and up).

Jack Dowie (1976) found the same pattern of returns in the British bookmaking market, hence this finding is not unique to pari-mutuel markets. Indeed, the bias is more pronounced in Dowie's data, which encompass all 2,777 races run in Britain during 1973. Figure 1 displays the rates of return from Snyder (1978) and a similar construction using Dowie's data. The returns to extreme long shots are markedly lower in the U.K. market. In addition, the returns to low odds horses present a puzzle. Breaking down the data more finely, Dowie's figures indicate that *bookmakers* lost money when taking bets on extreme favorites, horses with $O_i \leq 0.5$. There were 107 such horses, with a be-

fore-tax rate of return of .085 at final odds, which is roughly the bookmaker's loss for these bets.[26]

Ali (1977) analyzed betting in a sample of 20,247 harness races in a fashion similar to Hoerl and Fallin. In contrast to Hoerl and Fallin's data, Ali's data were strongly inconsistent with the null hypothesis that subjective and objective probabilities were equal. Table 4 contains estimates of objective

[26] An explanation consistent with evidence in Section 4.4 is that these horses were bet heavily and their odds were reduced during the betting cycle. (Note that this implies the bookies' losses were even greater than 8.5% for these horses). Offsetting these losses, perhaps, are profits for bookmakers on horses who were thought to be heavy favorites but whose odds drifted out in the betting. In both cases, incomplete adjustment of the odds occurs during the market period, resulting in improved but imperfect estimates of the probability of winning.

TABLE 4
THE DIFFERENCE BETWEEN SUBJECTIVE AND OBJECTIVE PROBABILITIES IN THREE SAMPLES

Rank	Ali (1977)		Asch, Malkiel, & Quandt (1982)		Busche & Hall (1988)	
	$\hat{p}_k$	$w_k - \hat{p}_k$	$\hat{p}_k$	$w_k - \hat{p}_k$	$\hat{p}_k$	$w_k - \hat{p}_k$
1	0.358	−0.035	0.361	−0.036	0.276	0.008
2	0.205	0.003	0.218	−0.013	0.190	−0.003
3	0.153	−0.001	0.170	−0.025	0.151	−0.009
4	0.105	0.007	0.115	−0.011	0.099	0.012
5	0.076	0.006	0.071	0.001	0.084	0.003
6	0.055	0.005	0.050	−0.002	0.063	0.003
7	0.034	0.008	0.030	0.004	0.048	0.001
8	0.021	0.007	0.017	0.008	0.047	−0.010
9			0.006	0.012	0.034	−0.006
10					0.021	0.000
11					0.023	−0.005
12					0.014	−0.001
13					0.010	0.001
14					0.013	−0.005

Notes: The probabilities need not sum to one because different numbers of horses participated in each race. $\hat{p}_k$ is the estimated objective probability for horses in each group, and $w_k - \hat{p}_k$ is the difference between the subjective and objective probabilities. A positive value for $w_k - \hat{p}_k$ indicates that horses in this group were overbet relative to a standard in which expected returns were equal across all classes of horses. The original sources reported tests of statistical significance for the variate $w_k - \hat{p}_k$. For the samples of Ali and Asch, Malkiel and Quandt, favorites (rank 1) were significantly underbet and extreme longshots were significantly underbet; there is no such pattern in the Busche and Hall data. Ali's sample is from 20,247 harness races run at 3 New York tracks from 1970–74. Asch, Malkiel, and Quandt's sample is from 729 races run at Atlantic City during the 1978 season. Busche and Hall's sample is from 2,653 races in Hong Kong from 1981 to 1986.

probabilities, and the differences between subjective and estimated objective prob-abilities, $w_k - \hat{p}_k$, for three studies: Ali (1977); Peter Asch, Burton Malkiel and Richard Quandt (1982); and Kelly Busche and Christopher Hall (1988). Ali's findings are in the first two columns of Table 4, where the bias is clearly evident.[27] Asch, Malkiel and Quandt (1982) found the same pattern in betting data from Atlantic City's race course.

The data from Hong Kong are different. Busche and Hall (1988) analyzed a sample of 2,653 races run in Hong Kong between 1981 and 1986. They found no evidence of biased returns in the Hong Kong pari-mutuel market. $w_k - \hat{p}_k$ is negative (−.035 and −.036) for the top ranked horse in the North American studies, indicating the favorite was relatively underbet.[28] In Hong Kong, this difference is positive. Whereas the subjective probability exceeds the objective by about .01 for North American long shots, there is no pattern in the Hong Kong long shots.

A Psychological Explanation. There are several potential reasons for the existence of the favorite-long shot bias. An obvious explanation is that bettors may underestimate the chances of

[27] The t-ratios reported by Ali testing the hypothesis that $w_k = \hat{p}_k$, are −10.3 for the favorite, insignificant for the second and third ranked horses, and above 3.0 for the 4th–8th ranked horses.

[28] It was also statistically significant, about 10 times its standard error mean in Ali, and twice as large in Asch, Malkiel, and Quandt.

favorites and overestimate those of the long shots. A common bias reported in the psychological literature is that people systematically overestimate the probabilities of small probability events. In laboratory experiments, subjects asked to estimate the probability of death from hazardous events underestimate common hazards, like a heart attack, and overestimate unlikely hazards, like being struck by lightning (Paul Slovic, Baruch Fischoff, and Sarah Lichtenstein 1982). Griffith viewed his findings on pari-mutuel betting as evidence that the biases found in surveys and experiments had a real-world counterpart in the favorite–long shot bias.

The favorite–long shot bias may thus be another manifestation of such misperceptions. But this argument is not persuasive. The questions asked in these experiments tend to be one-shot in nature. Other experiments, for example Gerald Dwyer et al. (1993), show that probability estimates made by individuals improve markedly as they acquire more experience with their environment. Similarly, biases reported by survey respondents are smaller in environments with repeated exposure to risk (W. Kip Viscusi and Charles O'Connor 1984).[29] The racetrack itself is a highly repetitious environment. Hence, it seems presumptuous to assign the cause of the favorite–long shot bias to systematic misperceptions.[30]

Variations in Local Risk Preference. As discussed above, Weitzman, Ali, and Quandt viewed the favorite–long shot bias as evidence of local risk preference in utility. The Hong Kong evidence shows that the bias is not universal, which casts doubt on the generality of this explanation. Ali (1977) estimated the degree of local risk preference exhibited by the representative agent and found that it varied in two important ways. First, the implied preference for risk was greater in the small market track than in the large market.[31] If one assumes there is a fixed cost of arbitrage, this result can be explained by the model in Section 2.3. The larger is the pool (N), the number of informed bettors who can cover their fixed costs by betting (n) increases. As can be seen in Table 1, increasing n causes the bias $w_k - p_k$ to shrink towards zero.

Second, Ali found that bettors exhibit greater risk preference in the last race of the day relative to earlier races. The last race offers a chance for losing bettors to "get out"; i.e. to recoup their losses with a final wager. Low probability/high odds horses provide that opportunity, and are thus over bet relative to a constant expected returns standard. Thaler and Ziemba (1988) interpret this as evidence of "mental accounting." Although one can speculate on rational reasons why bettors might want to avoid going home losers, this piece of Ali's evidence is certainly inconsistent with the standard economic framework.

Heterogeneous Bettors in Pari-mutuel Markets. Models with heterogeneous agents have been shown to yield outcomes consistent with observed returns in several dimensions. Hurley and McDonough's (1995) model provides a concise explanation of the bias. Suppose

---

[29] One hypothesis is that survey respondents overreact to newly identified risks, but adjust these expectations with more experience. A new hypothesis explored in Daniel Benjamin and William Dougan (1997) is that subgroups of the population report estimates specific to their own cohort. Correcting for cohort effects, the reported biases in the original experiments lose much of their significance.

[30] Note, however, that one psychological explanation for gambling is that those who persist in the activity in the face of unfavorable odds (and mounting losses) exhibit biased evaluation of outcomes (Thomas Gilovich, 1983).

[31] The parameter characterizing risk preference was .17 at the small market track and .06 at the large market track.

there are two groups of people betting on a two-horse race, with $p_F > .5$ for one of the two horses. The uninformed group bets equal amounts on each horse. The other group is informed; they know the true probabilities of winning and maximize the expected profit from betting. This setup is quite appealing in both its descriptive and analytical simplicity; it captures the essence of the racetrack market where a group of knowledgeable bettors seeks to profit on the action of a group that is poorly informed. In the model with no takeout rate and an unlimited number of informed bettors, the market equilibrium is perfect and win pool shares equal the objective probabilities of winning.

When the number of informed bettors is small, the influence of the uninformed remains in the odds in equilibrium. With an unlimited number of informed bettors the bias emerges when the takeout rate is positive. In this case, the equilibrium win pool share for the favorite is Qp_F. In this solution, informed bettors never bet on the long shot, and bets on the favorite take place until the expected profits are reduced to zero.[32]

Hurley and McDonough (1996b) also explore a model in which all bettors are boundedly rational. Here, bettors' probability estimates are drawn from an arbitrary distribution centered on the true probabilities. The simplest version is a two-horse race with zero takeout. Each bettor assumes that the current odds will prevail in the end, and wagers a fixed amount on the horse with the

greatest expected value. Hurley and McDonough show that the betting under these assumptions yields odds which converge to a unique equilibrium. The equilibrium win pool share of the favorite, w_F^*, solves the equation $w_F + G(w_F) = 1$, where $G(w_F)$ is the distribution function of the bettors probability beliefs. Let $p_F > .5$ be a solution. Since $G(p_F) = .5$ by assumption, $p_F + G(p_F) > 1$; hence only solutions where $w_F^* < p_F$ are possible when $p_F > .5$.[33]

The setup of this model has its appeal, but its implications are not so pleasant. First, in the absence of takeout there is a profit opportunity at equilibrium odds. Second, the favorite–long shot bias diminishes and ultimately reverses as takeout is increased. This is the opposite result of the earlier model. In sum, the models explored by Hurley and McDonough are useful and important attempts to capture what is empirically observed at North American racetracks. But more work of this sort is needed to get the magnitudes right.

Information Asymmetry in the Bookmaking Market. Hyun Song Shin (1991, 1992) modeled the bookmaking market with an emphasis on the information asymmetry between the bookmaker and the bettor. There is always a chance that the bettor knows a great deal more about the outcome of the race than the bookmaker. Shin's papers show that asymmetric information creates a favorite-long shot bias through an optimal pricing response by the bookmaker. Shin (1992) analyzes the problem for an *n*-horse race in a competitive equilibrium environment, and is our focus here.

The model allows for two types of bettors: a single insider, and a large group of uninformed outsiders. The

---

[32] Hurley and McDonough (1995) ran some experiments with this model, which did not support the implication that the takeout was responsible for the bias. In experiments with and without takeout, a bias of equal magnitude appeared in the two situations. The number of bettors was small however (17 and 18). More extensive experiments of this nature, in particular, experiments in which the number of bettors varies, seem worthwhile.

[33] For example, let probability beliefs be characterized by the uniform distribution [.55,.75] with $p_F = .65$. Then $G(w_F) = (w_F - .55)/.2$, and $w_F^* = .625 < p_F$.

bookmaker contracts with a single bettor of unknown identity. The bookmaker knows two sets of probabilities: the probability that the bettor is an insider, and the ex ante probabilities of each horse winning the race. The bookmaker sets prices with the knowledge that there is a probability, z, that he will contract with the insider, who has superior knowledge of the outcome of the race.

The probability beliefs of the bettors are extreme. The insider has been given a window to the future, and knows the outcome of the race with certainty. The outsiders don't know the outcome of the race, but each believes with certainty that a particular horse will win. Hence, both types will make a bet of the maximum amount allowed on their horse—let it be $1—regardless of the prices offered by the bookmaker. The outsiders are distributed in proportion to each horse's winning probability. Thus, if $p_i = \alpha$, then α is the proportion of outsiders that would bet on horse i. While extreme, these assumptions enable a clear view of the problem and yield an explicit solution.[34]

The problem for the bookmaker is to determine the expected profit-maximizing set of prices $\pi:(\pi_1, \ldots \pi_n)$. Hence the payoff when the bettor selects the winner is $1/\pi_i$.

The following considers the simplest case where z is constant across all horses. Conditional on the bettor being an outsider, expected profits for the bookmaker are $V(\pi|\text{outsider}) = 1 - \Sigma_i p_i^2/\pi_i$. If the bettor is an insider, expected profits are $V(\pi|\text{insider}) = 1 - \Sigma_i p_i/\pi_i$. Hence, the expected profit function for the bookmaker is

$$V(\pi) = 1 - \sum_i \left[z\, p_i + (1-z)p_i^2\right]/\pi_i \quad (3)$$

The bookmaker is assumed to operate in a competitive environment and earn zero profit. The simplest characterization of optimizing prices thus restricts the sum of prices to be the minimum sum with nonnegative expected profit for the bookmaker. Hence the objective is[35]

$$\min\{\pi_1 + \ldots + \pi_n | V(\pi_1, \ldots \pi_n) \geq 0\}$$

Letting $A_i = \sqrt{[zp_i + (1-z)p_i^2]}$, the solution is given by

$$\pi_i^* = A_i \cdot \sum_s A_s \quad (4)$$

Recall that π_i is the price of a contingent claim to $1 in the event horse i wins. Hence the deviation in the sum of the prices from 1, $D = \Sigma_i \pi_i^* - 1$, can be thought of as the market spread.

Using (4) and the definition of A_i it is simple to show that $\pi_i = p_i$ when $z = 0$ (there is no chance of an insider). Prices equal probabilities, $D = 0$, and there is no bias in the odds.

The favorite–long shot bias emerges when $z > 0$.[36] For example, consider a 2-horse race, where $p_1 = .25, p_2 = .75$, and z is constant at .05. Using (4), we find that $\pi_1^* = .2746$ and $\pi_2^* = .7746$, with $D = 1.0492$. The spread of 4.92 percent amounts to compensation for losses imposed on the bookmaker when the insider trades. The normalized prices π_i^*/D are .2617 and .7384, respectively. The normalized prices (the bookmaking market's counterpart to w_i) thus overstate the probability that the long shot wins, and understate the probability that the favorite wins.

Shin (1993) explores a testable implication of this model. Shin shows

---

[34] Shin (1992) shows that a more general specification of outsider expectations yields similar qualitative results, but without an explicit solution.

[35] This is a modification of Shin's setup which is used to simplify the discussion.

[36] The market is feasible only if z is "small," since large values of z result in strictly negative profits for all sets of the bookmakers prices. The model thus suggests a motive for restricting insider trading, since a market may not exist when z is large.

TABLE 5
DIFFERENCE BETWEEN PARI-MUTUEL AND BOOKMAKING PAYOFFS FOR WINNING 10P BETS

Odds Range of Winner	Number of Observations	Pari-mutuel Return	Return at Bookmakers' Odds	Difference	% Difference
Odds < 10	1271	52.7	48.4	4.3	8.9
Odds < 15	1353	63.9	53.7	10.2	19.0
Odds < 20	1408	74.8	59.1	15.7	26.6
All	1435	81.6	63.4	18.2	28.7

Source: Gabriel and Marsden (1991).
Notes: Data are grouped by the final odds in the bookmaking market. All differences are significant at the .01 level.

that the market spread can be approximated by the linear regression $D = z(n-1) + h(p)$. The main focus is on the coefficient of the $(n\text{-}1)$ term.[37] The estimated coefficient measures z, which must be positive (and small) for the model to be supported. For a sample of 136 races run in the U.K. during the first week of July 1991, Shin obtains a statistically significant estimate of .02 of the z coefficient.

Leighton Vaughan Williams and David Paton (1997) find that the market spread varies in predictable ways with proxies for insider trading. First, in higher quality races where the role of private information is relatively unimportant ($z \approx 0$), Vaughan Williams and Paton find no relation between the spread and n. Second, they find that the spread is greater in races where there are large changes in the odds, a feature consistent with the presence of inside information.

The evidence from Paul Gabriel and James Marsden (1990, 1991) provides additional support for Shin's model. Gabriel and Marsden compare the returns to winning bets in the pari-mutuel pools to those obtained in the parallel bookmaking market. Gabriel and

Marsden (1991) find that pari-mutuel returns exceed the returns from bookmakers (at final odds), and on this basis conclude that the betting market is not efficient. This result calls for explanation.[38] Nevertheless, notice how the returns differ as the odds increase, as displayed in Table 5. For winners at bookmakers' odds of 10–1 or less, the difference in the return is 8.9 percent; this difference grows to 29 percent as long shot winners are added to the group. This implies that the bias is more pronounced in the U.K. bookmaking market than in the U.K. tote market. This may help to explain the difference in Figure 1 between the returns in the U.S. pari-mutuel and U.K. bookmaker's market. Shin's model focuses on a bookmaker who adjusts prices as a

[37] $h(p)$ is an index related to the dispersion of p_i, but is not of interest here.

[38] Note that the pari-mutuel odds are by definition average odds, and that the final odds offered in the bookmaking market are more akin to marginal odds. Evidence discussed in the next section suggests that the volume of late betting (the marginal odds) is more informed than early betting, and that adjustment of pari-mutuel odds during the betting period is incomplete. Note that Gabriel and Marsden do not measure average returns to bets placed in the bookmaking market. If heavy betting is responsible for falling odds, the average return to bettors in the bookmaking market will be understated by their figures. There remains a puzzle however; the hypothesis just stated does not explain why bettors in off-course betting shops place wagers that are settled at final odds rather than in the pari-mutuel market.

consequence of the non-zero probability of confronting an insider that backs a long shot. This threat induces him to trim the odds on long shots relative to favorites. Pari-mutuel pools, by construction, are not subject to this source of the bias.

4.3 *Evidence of Informed Trading*

Stephen Figlewski (1979) used a multi-nomial logit model to determine if the pari-mutuel odds incorporated the opinion of experts. Let Z_i be a vector of information on horse i. p_i is assumed to be related to Z by the function

$$p_i = e^{\beta z_i} / \sum_{j=1}^{n} e^{\beta Z_j}$$

where n is the number of contestants in the race and β is a parameter estimated by maximum likelihood. Let Z be composed of two components, $Z = [Z_1 \, Z_2]$, with the associated parameter vector $\beta = [\beta_1 \, \beta_2]'$. For each horse i, the components are $Z_{1i} = w_i$, and Z_{2i}, which contains forecasts of the finishing position of horse i by 14 expert handicappers.[39] Denote the estimated probability of the horse that won the k^{th} race by P_w^k, and the likelihood function for the sample of M independent races is

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

The properties of this model imply that when $\beta = 0$, $p_i = 1/n$. The significance of $[Z_1 \, Z_2]$ is assessed using likelihood ratio tests. As Figlewski states, "the important question is whether adding handicapper information significantly improves the fit over what was obtained using the odds alone." The effect is marginal; the null hypothesis that the handicappers add no information ($\beta_2 = 0$) is rejected at the 10 percent but not the 5 percent level. When Figlewski uses the parameter estimates in an out-of-sample test, he finds that the "handicapper data contributes nothing." The "track odds fully discount the published choices of professional handicappers" (p. 86–87).

Figlewski also ran a test based on data where the on-track and off-track pools were disaggregated. Figlewski conjectured that the off-track bettors were less informed than their on-track counterparts. In this test, the handicapper information added significant improvement to the predictions based solely on the off-track pool, while producing no improvement in the predictions based solely on the on-track pool, supporting Figlewski's conjecture.

Asch, Malkiel, and Quandt (1982) suggest that informed bettors are more likely to wait until the end of the betting period in pari-mutuel markets before making their bets, a view which is widely held. There are two reasons for this. The first is a consequence of the fact that all payoffs are determined by the final odds in pari-mutuel markets. Early in the betting the posted odds are more volatile—since the pools are small at that point, a modest wager can drop the odds on a horse from 7–1 to 3–1. Informed bettors can obtain a lower variance of their estimate of the expected return late in the period, when the pool is bigger as a result of the accumulated betting. The second reason is strategic. A bettor (or group of bettors, such as a racing stable) with inside information obviously prefers to keep the information private. Betting large sums can affect the odds and alert tote-board watchers—the racetrack's version

---

[39] The handicappers' forecasts of the top two finishers were incorporated in the following way. Suppose that handicapper 1 picked horse i to win, handicapper 2 picked horse i to finish second, and that the other 12 handicappers do not pick horse i to finish in the top two. Then, if handicapper 1's picks (1st choice, followed by the second) are followed by handicapper 2's, and so on, Z_{2i} would be [1 0 0 1 0 0 0 0 0 0 0 0 ....].

TABLE 6A

RATIO OF FINAL ODDS AND MARGINAL ODDS TO
MORNING LINE ODDS BY FINISH POSITION FOR 729
RACES RUN AT ATLANTIC CITY (NJ) IN 1978

Horses Finishing	O_{FINAL}/O_{ML}	O_{LATE}/O_{ML}
First	0.96	0.82
Second	1.16	1.06
Third	1.22	1.17
Also Rans	1.59	1.63

Source: Asch, Malkiel, and Quandt (1982, p. 193).

Notes: Figures in the columns are ratios of average odds. O_{FINAL} are the final odds, O_{ML} are the morning line odds (projected by the track's handicapper), and O_{LATE} are the marginal odds produced by bettors in the last third of the betting period.

of a technical analyst—that something is up. Betting late into a larger pool thus reduces the likelihood of creating bandwagon effects in the odds.

Asch et al. use data from the 765 races run at Atlantic City Race Course in 1978. They find that the winning horse is "bet down," that is, its final odds, O_{FINAL}, are lower than the morning line estimate of the odds produced by the track's expert handicapper, O_{ML}. These results are presented in Table 6.A. O_{FINAL}/O_{ML} is 0.96 for horse the ultimately wins the race. The ratio for all other horses (including the second and third place finishers) ranges from 1.06 to 1.63.

Furthermore, late money appears to be more informed than early money. Asch, Malkiel, and Quandt calculated the marginal odds based exclusively on wagers made in the last eight minutes of the betting. Using marginal odds, O_{FINAL}/O_{ML} is 0.82 for the winner, and remains above 1 for horses that don't win. As Asch, Malkiel, and Quandt (1982, p. 306) put it, "winning horses are especially preferred by the late bettors."

N. F. R. Crafts (1985) studied the bookmaking market in the U.K. along these lines. In this market, the odds analogous to the North American morning line are issued as the FP, or forecast price, by *Sporting Life*, a daily trade publication. Representatives of *Sporting Life* observe the betting, noting in particular the very large bets and the odds at which they are transacted. A description of the betting is then printed in a subsequent edition of the paper. At the end of each betting period, these representatives determine the starting prices for the horses, SP, which are the odds offered in the on-course bookmaking market at the end of the betting period. Crafts notes that the practice of paying off at odds offered at the time of each transaction enhances the value of inside information, since bandwagon effects from betting large sums do not affect the payoff.[40]

Craft's sample covers 16,769 horses that ran between September 1982 and January 1983. Denote the subjective probabilities at FP and SP odds as s_{FP} and s_{SP}. Horses for which either (*i*) $1.5 \leq s_{SP}/s_{FP} < 2$, or (*ii*) $s_{SP}/s_{FP} \geq 2$ are considered to have been "heavily backed;" i.e. the wagers in the market

[40] It is important to understand how bookmakers change odds during the betting period. Standard practice is to begin with offers of low odds on all horses, and to push the odds out until they start to attract betting. This addresses, to a degree, an adverse selection problem that the bookmaker would face if he initially offered odds equal to his forecast of optimal prices (whereupon relatively informed bettors would wager only on his mistakes). Given the typical odds progression, the critical question for informed agents is how long to wait before accepting the offered odds. In the case of a tightly knit betting coup, the group will let the odds drift out until the final minute, and wager all of their money at once (at various locations) before the odds can adjust. Note that as the relevant information becomes more widely held, the incentive to wait is offset by the knowledge that others are likely to accept the bookmakers odds, and that this betting may reduce the odds available on your horse.

TABLE 6B
RATES OF RETURN AT FP AND SP ODDS FOR HORSES CHARACTERIZED BY ODDS MOVEMENTS
FOR 16,769 HORSES IN THE U.K., 1982–83

Odds Movement	Number of Horses	Rate of Return at FP	Rate of Return at SP
Heavily Backed:			
$p_{SP}/p_{FP} \geq 2$	397	1.41	−0.09
$1.5 \leq p_{SP}/p_{FP} < 2$	712	0.64	−0.01
Drifting Out			
$1.5 \leq p_{FP}/p_{SP} < 2$	858	−0.63	−0.38
$p_{FP}/p_{SP} \geq 2$	317	−0.64	−0.13

Source: Crafts (1985, p. 298)
Notes: p_{SP} is the implied subjective probability at SP odds, p_{FP} the same at FP odds. Horses whose odds decline in the betting will have p_{SP}/p_{FP} ratios which exceed 1.0; vice versa for horses whose odds increase during the betting. Rates of return do not include the 4% (10%) tax on course (off course) in effect at the time.

have pushed the odds down well below the forecast level. Table 6.B presents Crafts' results. There are 712 horses satisfying the former and 397 horses satisfying the latter, more pronounced condition. Bets on these horses at SP odds are not profitable, earning pre-tax returns of −.01 and −.09, respectively. Bets made at FP odds would have yielded returns of .64 and 1.41, the latter being phenomenally profitable. On the flip side, the 1,175 horses whose odds drifted out were very poor bets: horses with $s_{FP}/s_{SP} \geq 2$ yielded returns of −.64 at *FP* and −.13 at *SP* odds; horses with $1.5 \leq s_{FP}/s_{SP} < 2$ yielded returns of −0.63 at *FP* and −0.38 at *SP* odds.

Using pari-mutuel data from races run in Chicago, Robert Losey and John Talbott, Jr. (1980) obtained a similar result. Losey and Talbott's simulation placed 579 bets on all horses for which the Daily Racing Form's expert handicapper placed a morning line of 3–1 or less, but whose final pari-mutuel odds exceeded this estimate. The returns were −28.4 percent, which exceeds by a large margin the 17 percent loss (takeout rate + breakage) expected if the rates of return were equal across horses.

What does one make of these rates of return? It is clear that trading in these markets creates measures of the probability of winning which are significantly better than measures produced by an individual or group of experts. This suggests that the betting market aggregates disparate sources of information into a superior probability estimate of the race's outcome. Second, the odds adjustment stops short of achieving constant returns; to equalize returns (at SP odds) between horses whose odds have fallen and those that have risen would require additional reductions and additional increases for each group.[41] Third, and this is Crafts' main point, these returns clearly point to the existence of an informed class of bettors. The published descriptions of the betting are helpful in this regard, for they establish that bets were made at odds substantially greater than SP odds for many of these winners. Consider these descriptions, first for a winner that had never

[41] Incomplete adjustment is related to the favorite-long shot bias, since horses whose odds shorten are more likely to be favorites, and those who lengthen long shots. This feature is closely related to Hurley and McDonough's (1995) model of the bias.

previously finished better than fourth—"dropped dramatically from 12–1 to 7–2," and second on a winner with a similar record—"following some 8–1 and 6–1 was reduced from 4–1 to 11–4 favorite" (Crafts 1985, p. 301). It seems quite unlikely that the average punter was responsible for the tumbling odds in these cases.

Adi Schnytzer and Yuval Shilony (1995) provide indirect evidence of inside information. Schnytzer and Shilony study the betting market in Australia, in which on-track and off-track pari-mutuel betting coincides with an on-track bookmaking market. Their data consists of the win pool shares from the on- and off-track pari-mutuel markets for all horses competing in 168 races run in 1984. The betting from both sources is aggregated into a single pool for determining payoffs, with about 3/4 of the money originating from off-track bets. Let w_i^{on} represent the share of the on-track win pool for horse i, and w_i^{off} represent that of the off-track market. Betting at the two sources is quite different: the correlation coefficient of w_i^{on} and w_i^{off} is only .48; the rank order correlation is .40. On-track bettors appear to be better informed, obtaining a (pre-takeout) rate of return of 9.7 percent compared with a loss of 3 percent in the off-track market.

Schnytzer and Shilony consider the hypothesis that on-track pari-mutuel bettors can capitalize on inside information created in the bookmaking market. Their story, like that of Asch et al., is based on common racetrack wisdom. An individual possessing inside information will generally be better off placing large bets with bookmakers rather than in the pari-mutuel pool, for two reasons. First, bookmakers will accept large wagers at the odds posted the moment the bet is placed, whereas a large wager reduces the odds in the pari-mutuel pool. Sec-

ond, large bets signal information to others when the odds on a horse fall sharply, and can create bandwagon effects in the betting. Bandwagon effects will further reduce the pari-mutuel odds, while having no effect on bets already contracted for in the bookmaking market. Schnytzer and Shilony hypothesize that on-track pari-mutuel bettors observe tumbling odds generated by insiders in the bookmaking market and can capitalize on this activity, whereas the off-track bettors are unable to do so.

Schnytzer and Shilony use the difference between w_i^{on} and w_i^{off} as a proxy for "plunges" in the on-track market odds based on inside information. Let $X_{1i} = 1$ when $w_i^{on} - w_i^{off} > 0.1$; $X_{1i} = 0$ otherwise; and $X_{2i} = 1$ when $0 < w_i^{on} - w_i^{off} \leq 0.1$; $X_{2i} = 0$ otherwise. Hence, when $X_{1i} = 1$, the on-track market's subjective probability that horse i wins exceeds that in the off-track pool by .1 or more; less so when $X_{2i} = 1$.

The authors estimate equations similar to Figlewski (1979) in order to test for the presence of inside information. Their maintained hypothesis is that all public information is fully reflected in the (much larger) off-course pools. Consider then a set of regressors $Z = [Z_1 \ Z_2 \ X]$ with associated coefficients $\beta = [\beta_1 \ \beta_2 \ \beta_3]'$ for the Figlewski model in which $Z_{1i} = w_i^{off}$, $Z_{2i} = w_i^{on}$, and X_i contains the proxies X_{1i} and X_{2i} defined above. The basic model contains just public information (Z_1). If private information is either not present or not useful, then the exclusion restriction $\beta_2 = 0$ should be accepted, but it is decisively rejected. If X is added to the basic model, one obtains a similar result, rejecting the hypothesis that $\beta_3 = 0$. If one accepts Schnytzer and Shilony's hypothesis that the off-track pools fully reflect all public information, this evidence suggests that inside information

is present at the track and important to predicting the outcome of the race.

The startling finding in Schnytzer and Shilony is the following. Suppose one obtained measures of $w_i^{on} - w_i^{off}$ just prior to the race. The authors show that bets placed on horses in proportion to this difference would have realized an after-tax profit of between 26.1 and 32.5 percent—very large numbers, indeed.[42] Although the sample of races is small and the exercise should be replicated, this is substantial evidence of inside information in the on-track betting market.

4.4 *Are There Profitable Wagering Rules?*

The answer to this question is a qualified "Yes." The qualifications stem from the fact that (1) some opportunities arise infrequently; and, (2) in some cases the betting pools are small, limiting the ability of net earnings to cover the opportunity cost of time. With these caveats in mind, there are some interesting cases of wagering rules which have been found to be profitable in these markets. We discuss them in turn.

Trends in the Odds. Based on his earlier paper, Crafts (1994) proposed the following trading rule in the U.K. market: bet when a horse has not yet started this year, $s_{SP}/s_{FP} > 1.5$, and the SP odds are less than 7–1. The heavier than anticipated betting indicates insider knowledge of the horse's fitness and the stable's intentions to put forth a winning effort. Bets at SP earned rate of profit of .558, and a phenomenal 2.619 had one made a bet at FP (the first insider to move?). Although the odds have tumbled, the market has not completely adjusted to the presence of

inside information by the close of betting. Crafts (1994, p. 548) appropriately cautions those seeking to capitalize: "in a period of five years only 88 bets would have resulted from a very substantial investment of time and effort." Based on their similar findings on late betting, Asch, Malkiel, and Quandt (1984) seek to determine if following the late money in the North American pari-mutuel market could be profitable. The rule they explore is much less specialized than that of Crafts (1994), and does not fare so well. Bets are placed on horses whose odds drop late in the betting. Although they initially report evidence of positive profits, this was the result of a programming error.[43] The results reported in Asch, Malkiel and Quandt (1986) correct the original programming error and also test the system on a new set of data. Although following the late money yields a loss which beats the track's takeout rate, the scheme is not profitable.

Pseudo-Arbitrage Using the Win Pool's Probability Estimates. Ali (1977) examined daily double bets, which require the bettor to select the winner of consecutive races. Assuming the win pool provides efficient estimates of probabilities and that the two races are independent, efficient odds for the double bet are simply the inverse of the multiple of the two probabilities less 1. Ali's data conform to this prediction.

David Harville (1973) presented a set of formulas for determining the probability that horse j finishes second or third in an n-horse race. These formulas are based on p_k, the probability that horse k wins the race, and a critical assumption concerning ρ_j, the probability that horse j beats all horses other than horse k. The joint probability that horse

---

[42] The authors note that the information required to exploit this was not available to on-track bettors in the real time setting.

[43] Place and show odds were converted to payoffs with faulty computer code.

k wins with horse j second is assumed to be $p_k \cdot \rho_j$, where $\rho_j = p_j/(1 - p_k)$.[44] Extending this logic, combinations of joint probabilities can be used to estimate the probability of any rank order of finish. This can be used to estimate the probabilities of exactas, bets where the order of the first two finishers must be selected, and place and show bets, which pay off if the selected horse finishes in the top two or top three, respectively. The resulting equations have been widely applied and are now known as the "Harville formulas."

In a clever application of efficient markets theory, Hausch, Ziemba, and Mark Rubinstein (1981) constructed a wagering rule for place and show betting. Hausch et al. obtain efficient estimates of the probabilities of winning from the win pool and use the Harville formulas to calculate the probabilities of running second or third. The question, then, is whether the place and show pools are efficient. Assuming the Harville formulas measure the true place and show probabilities, Hausch et al. show that the win pool shares should equal the place and show pool shares; i.e. $w_i = \text{place}_i = \text{show}_i$, where place_i and show_i are the shares of the respective pools bet on horse i. If these ratios differ, either the Harville assumption is wrong or the market is inefficiently pricing one of the wagers. Hausch et al.'s prior is that place and show pools are likely to be less efficient than the win pool for two reasons. First

is the complexity of determining the efficient payoff. Second, assuming that the typical bettor is risk-loving, few will pay attention to these pools, in which the payoffs are relatively low.

The proposed betting rule is to wait until the last two minutes of the betting period, calculate the ratios for likely (low odds) horses and bet on horse i in the place or show pool when w_i/place_i or w_i/show_i is 1.15 or greater. Hausch et al. report that a positive, significant rate of return was realized in several simulations at various racetracks.

This wagering rule is the basis of the "Dr. Z System," which has received considerable attention in the racing world. The formulas have been encoded on computer chips and detailed in books for laymen (Ziemba and Hausch 1987). The widespread dissemination of this information should have a depressing effect on realized returns. This issue should be explored systematically, although there is already some evidence that this is the case. Financial columnist Daniel Seligman (1985) tested the Dr. Z System at Belmont Park in New York. In a two-week period, the system generated 33 bets from 108 races, resulting in a rate of return of –0.067. Jay Ritter (1994) argues that the racetrack market is one in which Dr. Z-type bettors efficiently arbitrage the implied price differences between different pools. Ritter obtains ex post profitability for a Dr. Z-type system using the known pool shares from the close of the market, but he shows in a simulation that a modest degree of uncertainty in the pool shares is sufficient to drive the rate of return to zero.

Cross-Track Arbitrage. Wagering takes place at multiple locations, notably for races of national interest such as the Kentucky Derby and the Triple Crown series. In some cases the pools at each locale are independently managed, with the consequence that the payoffs at one

---

[44] Harville thus assumes that there is a "race within the race," where the probability the horse j beats the remaining n-2 horses is independent of where k finishes. Indeed, it is independent of whether or not horse k even runs in the race. Although this is a useful assumption, there are cases in which it will mislead, since outcome of a race can be influenced by the strategies of the runners and the number of runners with each strategy. See Donald Hausch and William T. Ziemba (1994) for further discussion of the systematic biases in probabilities derived from this assumption, and Hal Stern (1990) for an alternative to the Harville model.

place, say Hollywood Park in California, might differ from another, say Churchill Downs in Kentucky.[45] Differences in payoffs appear to reflect local information or perhaps a rooting interest in a locally based horse. For example, the 1986 Derby winner Ferdinand (a horse based in California) paid off at odds of 7.4.–1 at Hollywood and 17.7–1 at Churchill. Price differences this large are obvious candidates for arbitrage.

Hausch and Ziemba (1990) provide a systematic analysis of price differences for nine Triple Crown races in the early 1980s. The prices come from separate wagering pools at between four and eleven race tracks scattered throughout North America.[46] Hausch and Ziemba record the final odds at each location, and construct a portfolio of wagers designed to yield a certain return of $1. By wagering the appropriate amount on each horse at the track where its odds are the highest, they can calculate the minimum cost of winning $1. This cost is less than $1 in seven of the nine cases, with an average rate of return of about 5 percent. These opportunities do not arise frequently however, and some of the wagering pools are small enough (Hausch and Ziemba 1990, p. 72) to limit the bet size and hence the magnitude of the returns. These considerations make it a leap to consider cross-track arbitrage as an unexploited profit opportunity. Still, the evidence does represent another crease in what is predominantly a smooth pattern of efficiency in the racetrack betting markets.

---

[45] Advances in telecommunications and pari-mutuel systems now enable such wagering pools to be merged. Merged pools have since become the norm.

[46] The number of tracks offering such wagering varied during the period, as did Hausch and Ziemba's success in obtaining information on the wagering pools. Note that I ignore for purposes of this discussion one race where their sample contains data from just two race tracks.

4.5 *Summary of the Literature on Racetrack Betting*

Before turning to point spread betting, it is worthwhile to summarize the key findings of the literature on racetrack betting.

1. Win pool shares w_i are good approximations to p_i, the probability that a horse i wins the race, the favorite–long shot bias notwithstanding.

2. There is ample evidence that the betting market creates improved estimates of p_i through trading. Disparate sources of information are aggregated in the market, and the opinions of "experts" appear to be fully discounted in market prices.

3. Odds adjustments improve prediction, but adjustment stops short of being complete. Due to transaction costs, there is little incentive for additional trading to complete the process, since final odds in both pari-mutuel and fixed odds markets generally yield negative expected returns.

4. Late money is smart money. In addition, there is evidence that suggests that informed traders earn positive profits in the bookmaking markets of the U.K. and Australia.

5. Representative-agent models emphasizing risk-loving behavior can provide an equilibrium explanation of the favorite-long shot bias. However, there is abundant evidence (points 2–4 above) that bettors are different, and that these differences are integral to establishing efficient prices. Models with heterogenous bettors, in which risk preference is ignored, are consistent with some features of the favorite-long shot bias, but more work needs to be done. Work documenting the source of variation in the favorite-long shot bias would be particularly useful.

6. Trading rules documenting positive profits exist, but in three of the

four cases studied here—Crafts (1985), Hausch, Ziemba, and Rubinstein (1981), and Hausch and Ziemba (1990)—the returns from following these rules may not cover the fixed costs required to execute them. The fourth rule (Schnytzer and Shilony 1995) used data not available in a real-time setting, but is suggestive of a very large discrepancy between returns based on public and private information.

In sum, models of efficient pricing can account for much of the variation in prices and rates of return at the racetrack. Yet numerous outstanding puzzles remain, and fundamental modeling problems remain to be solved. The evidence suggests that an informed class of bettors is responsible for altering prices in these markets, yet observed returns are uniformly non-positive. We do not yet have an equilibrium model that motivates the acquisition of information and at the same time remains consistent with these facts.

5. *Point Spread Betting on Team Sports*

5.1 *Simple Tests of Point Spread Efficiency*

Although fixed odds betting exists on some team sports, most papers have focused on the point spread market, in which bettors attempt to "beat the spread" through superior prediction of the score difference of games.[47] Most studies focus on the restrictions imposed on point spreads by models of efficient pricing. Two basic questions have been extensively explored. First, is it possible to identify systematic profit opportunities in point spread betting? Second, provided that certain distributional assumptions are satisfied, effi-

ciency implies that the point spread (*PS*) is an unbiased, minimum variance estimator of the difference in points (*DP*) scored in a game—does this property hold? Beyond this, some interesting and creative applications of the efficiency concept have been employed, but we begin with the basics.

Efficient pricing of wagers is generally asserted by the absence of profit opportunities condition, as defined in Section 2.2. Let $p(1-p)$ represent the probability of winning a point spread bet by wagering on the favorite (underdog). Using equation (2), efficiency restricts p (also $1-p$) to the interval

$$p \in [.476, .524] \tag{5}$$

since $\tau = .1$ in the Las Vegas market. Again, $\tau = 0$ would imply that $p = .5$, and thus that *PS* is the median of the distribution of *DP*.

An alternative to the favorite/underdog condition redefines the score difference and point spread to a home team minus visiting team basis. This procedure has been adopted in much of the recent literature. This is useful, since point spreads under this ordering measure the market's estimate of the home court advantage, a factor (in contrast to a team's favoritism) which is fundamental in its own right to the outcome of the game.

Provided that the distribution of score differences is symmetric, an additional implication of efficiency is that *PS* is an unbiased forecast of *DP*:[48]

$$PS = E(DP) \tag{6}$$

Continuing within this framework, a stronger definition of efficiency implies that the point spread fully incorporates

---

[47] A notable exception is Woodland and Woodland (1991), which shows that no favorite long-shot bias exists in the odds market for baseball games.

[48] Again, this follows if transactions costs (τ) are zero; with symmetry the median and mean of point differences coincide. This is strong but allows the use of standard hypothesis testing procedures. The exchange between William Even and Nicholas Noble (1992), and Gandar et al. (1993) covers distributional issues.

TABLE 7
SCORE DIFFERENCES AND POINT SPREADS FOR NBA GAMES

A. Sample Frequencies

Differencing Method/ Sample Partition	Games	Bets	Wins	Ties	Wins/Bets
A1. Home-Away					
All Games	5636	5510	2789	126	.506
Home Favorites	4341	4243	2148	98	.506
Home Underdogs	1209	1181	600	28	.508
Pick 'em Games	86	86	41	0	.477
A2. Favorite-Underdog	5550	5424	2729	126	.503

B. Sample Means and Standard Deviations

Differencing Method/ Sample Partition	DP	PS	DP-PS	t-stat
B1. Home-Away				
All Games	4.62(12.42)	4.38(5.59)	0.24(11.15)	1.62
Home Favorites	6.87(11.82)	6.81(3.62)	0.06(11.07)	0.37
Home Underdogs	−3.09(11.74)	−4.05(2.30)	0.96(11.45)	2.91
Pick 'em Games	−0.91(10.58)	0.00(0.00)	−0.91(10.58)	−0.79
B2. Favorite-Underdog	6.05(11.83)	6.21(3.56)	−0.16(11.16)	1.06

Notes: (i) Sample Characteristics: The sample encompasses all regular season NBA games played in the six seasons from 1982–83 through 1987–88. Score differences were obtained from the annual edition of the *Sporting News NBA Guide.* Point spreads were obtained from *The Basketball Scoreboard Book.* These point spreads are those prevailing in the Las Vegas market about 2.5 hours prior to the start of play (5 PM Eastern time on a typical night). No point spread is reported for 22 games during this period, which reduces the sample from 5658 (all games played) to 5636 (all games with point spreads).

(ii) Panel A: This panel lists the number of games, bets (the number of games in which DP ≠ PS, which are ties), and the number of bets won by wagering on the team in the first position of the score difference. Wins/Bets is the sample estimate of p, the proportion of such bets won. Since this proportion always lies inside the bounds given by (2), no test statistic is required to evaluate this implication of efficient pricing.

(iii) Panel B: Standard deviations are given in parentheses. The t-statistic tests the null hypothesis that the mean forecast error (DP – PS) is zero. Although the null is rejected in the case of home underdogs, the failure to reject efficient pricing in panel A for this partition indicates that the rejection in B is caused by a departure from the symmetry assumption.

all relevant information. Denoting the set of all relevant information by Ω, this requires that

$$E(DP - PS \mid \Omega) = 0 \qquad (7)$$

which says that the forecast error is unrelated to relevant information (because this is already present in equal amounts in both *DP* and *PS*).

Table 7 presents data from the betting market on NBA games. This table can be used to assess the implications listed above, using both the home team-away team and favorite-underdog ordering of score differences. Also included are all partitions of these orderings, including games in which there was no favorite (*PS* = 0). Panel A examines equation (5). In no partition is the ratio of winning to total bets outside the bounds implied by (5).

Sample means and standard deviations of *DP*, *PS*, and *DP-PS* are presented in panel B of Table 7. Note the

relatively high variability in *DP*: for pick 'em games in which *PS* = 0, the average score difference is –0.91; but the range including just one standard deviation is [–11.49, 9.67], or about 21 points. NBA score differences thus appear to be very noisy. The hypothesis that *PS* is the expected value of *DP* can be examined in panel B. The right-hand column contains *t*-statistics for testing the null hypothesis that $E(DP - PS) = 0$. Note that $t = 2.91$ in the case of home underdogs, which is sufficient to reject (6). The information in panel A suggests that this rejection is due to a violation of the symmetry condition, and not a violation of efficient pricing, since the proportion of wins from bets is .508 in this partition. It makes sense to examine the data with simple procedures such as these so that subsequent errors are avoided, an issue we revisit in the following section.

Vergin and Scriabin (1978) examine a number of simple betting rules for potential profitability, and conclude that "discernable biased patterns" existed in NFL point spreads during the 1969–74 seasons. An inaccurate but common interpretation of their results is that "Vergin and Scriabin have identified exploitable biases . . . [and] show that arbitrage possibilities exist" (Peter Pope and David Peel 1989, p. 325). The arbitrage claim can be traced to a single anecdote in Vergin and Scriabin (p. 816) on the existence of different point spreads in different cities for the same contest. What Vergin and Scriabin actually show is that significant profits could have been realized *if* they obtained a point spread favoring a given betting strategy by one or two points over the published point spread. A proper interpretation of this evidence is that differences in the spread as small as one or two points can have a tangible effect on wagering profitability.

The success of the Vergin and Scriabin strategies against actual point spreads is summarized in Table 8 for two samples. The table lists the sample sizes, proportion of winning bets, and *p*-values for two hypothesis tests for each strategy.[49] Column 3 lists the p-value from testing the hypothesis that the proportion equals 0.5. This does not test the profitability proposition however, as pointed out originally by Tryfos et al. (1984). Column 4 lists the *p*-value from testing the hypothesis that the proportion < 0.524. While these strategies were profitable in Vergin and Scriabin's sample, the *p*-values in column 4 are too high to reject the null hypothesis that $p < .524$ for two strategies, and a third is marginal (the *p*-value is .112).

Tryfos et al. (1984) point out that one difficulty with this type of analysis is the large number of strategies "that can be defended on reasonably plausible a priori grounds." Hence, it is not clear what a small number of biases really tells us about the point spread market. The acid test is if the strategies found profitable in an original study remain so out of sample. Tryfos et al. conducted out-of-sample tests on the Vergin and Scriabin strategies. The results for the 1975–81 period are presented in panel B of Table 8. The strategy of betting big underdogs remained profitable, but the result is not statistically significant (the *p*-value is .277); the remaining strategies incurred losses. The Vergin and Scriabin paper and its extension by Tryfos et al. established a pattern that would be repeated: sightings of profitable wagering rules are occasionally reported, but they often disappear on subsequent investigation.

---

[49] Of the rules considered in Vergin and Scriabin, only the betting rules most favorable to the profitability proposition are presented in Table 8.

TABLE 8
THE PROFITABILITY OF SIMPLE BETTING RULES

A. Vergin and Scriabin's Sample 1969–74

	N	$\hat{p}$	p-value (p = 0.5)	p-value (p < 0.5238)
Bet on big underdogs	674	.546	.017	.249
Bet against big winner	78	.538	.499	.795
Bet on turnaround team	59	.627	.051	.112
Bet on strongest team	57	.667	.012	.012

B. Tryfos et al.'s Sample 1975–1981

	N	$\hat{p}$	p-value (p = 0.5)	p-value (p < 0.5238)
Bet on big underdogs	735	.535	.060	.277
Bet against big winner	N/A	N/A	N/A	N/A
Bet on turnaround team	76	.447	.359	—
Bet on strongest team	71	.459	.486	—

Notes: $\hat{p}$ is the proportion of winning bets from N tries. The p-values in the final two columns test the hypothesis that the true probability of winning is .5, and < .5238, respectively. Big underdogs are predicted to lose by more than 5 points by the point spread. The big winner is the team with the largest margin of victory in the prior week. The turnaround team is the team that beat the spread by the largest amount over the prior 4 weeks. The strongest team is the team with the largest victory margin over the prior 4 weeks. Tryfos et al. did not reexamine the big winner strategy.

5.2 *The Simple Linear Prediction Model*

Equation (6) has been repeatedly examined in the context of a linear prediction model. The basic form of this model is

$$DP = \alpha \cdot H + \beta \cdot PS + \varepsilon \qquad (8)$$

where H is a vector of ones, α and β are regression coefficients, and ε is an error term. Market efficiency is examined by testing the joint null hypothesis that $\alpha = 0$ and $\beta = 1$; i.e., that PS is an unbiased linear predictor of DP. With scores and spreads ordered on a home team minus away team basis, it is clear that the intercept term α reflects advantages of the home team that are not priced in the betting market.

The first versions of equation 5.4 were estimated by Ben Amoako-Adu, Harry Marmer, and Joseph Yagil (1985) and Richard Zuber, John Gandar, and Benny Bowers (1985), and were presented as evidence that point spreads were very poor and probably inefficient predictors of score differences of professional football games. Zuber et al. estimated separate regressions for each of the 16 weeks of the 1983 NFL regular season, and failed to reject the non-informative null hypothesis that DP is *unrelated* to PS (i.e. that $\alpha = \beta = 0$) in 15 of the 16 weeks.[50] They conclude (p. 802) that the noninformative null hypothesis "is as consistent with the sample data as is the efficiency hypothesis."

Amoako-Adu et al. reversed the dependent and independent variables in the regression. Their estimated equation is $PS = -4.47 + 0.04 \cdot DP$, with an R^2 of .04 from a sample of 233 games.[51]

---

[50] Note that even in the absence of point spread effects (let $\beta = 0$), Zuber et. al's non-informative null hypothesis could have been rejected if α were sufficiently non-zero. Hence, one interpretation of their results is that it is difficult to establish the well-known positive effect of playing at home using a single week's sample of games.

[51] In contrast to most studies, the ordering of point differences in Amoako-Adu et al. is as

This apparently weak relation lead them to conclude (p. 376) that "there is very little connection between spreads and actual game outcomes."

These conclusions are simply wrong. The lack of a tight statistical relation between *DP* and *PS* in these studies is due to the relatively large variation in *DP*. Sauer et al. (1988) argue that the weekly samples in Zuber et al. are too small (14 games each) to reject a non-informative null hypothesis. Amoako-Adu et al.'s method dilutes the signal to noise ratio. In the standard setup, the estimate of β in equation (8) is cov(*PS*,*DP*)/var(*PS*). By reversing the dependent and independent variables, the denominator becomes var(*DP*). This adds the noise source (ε) to the denominator, which reduces the slope coefficient in proportion to the variance of ε, and makes for a more complicated test of efficiency.[52]

Subsequent papers estimating equation 5.4 in large samples (Gandar et al. 1988, and Sauer et al. 1988) find that *DP* and *PS* are indeed related and were unable to reject the efficiency null. Indeed, pooling data from six NFL seasons, Gandar et al. obtained point estimates for α and β of 0.01 and 1.02, respectively. Although the equation ex-

plained little of the variance in *DP* (R^2 was .14), the *F*-statistic testing the joint hypothesis that α = 0, β = 1 was also tiny (0.03). What seemed to be clear was that (a) score differences were noisy and hard to predict, and (b) the point spread was a weak but unbiased predictor of the score difference.

Or was it? Joseph Golec and Maurry Tamarkin (1991, p. 312) argued that these "tests have low power to reject the null hypothesis of market efficiency because they are mis-specified." The lack of power stems from the naivete of regressions based on (8), since it can test for mispricing only in the home team-visitor dimension. Golec and Tamarkin argue that these regressions ignore potentially relevant information; for example, most of the time one team is favored and the other is an underdog, and this information is not fully accounted for in (8).

Potential biases can potentially cancel each other out, in the following way. Suppose that 2/3 of the games consist of favorites playing at home, and that favorites at home are overbet on average by one point (i.e. that E(*DP* − *PS*) = −1 in this sub-sample). In the remaining 1/3 of the games the home team is an underdog, and suppose these teams are underbet by two points. The reciprocal nature of point differences implies that favorites playing away are overbet by two points. Thus, in the full sample, all favorites are overbet, on average, by 1.67 points. But when point differences are constructed on a home-away basis, the sample means of *DP* and *PS* are the same—the two biases posited above cancel exactly. It follows that a regression based on (8) using home-away point differences will fail to detect what is, by construction, an obvious case of mispricing. This is the concern of Golec and Tamarkin.

But their cure for this problem is

---

printed in the daily papers; i.e. if the Cowboys are favored by 7 points over the Steelers, *PS* = −7. Hence there are no positive values for *PS* in the sample.

[52] The ratio of the slope coefficient to its standard error (*t*-statistic) in Amoako-Adu et al. of 2.93 implies there is a significant relation between *DP* and *PS*, but the low value of the regression coefficient makes the relation appear weak. Using the definitions of the least squares coefficients, their estimates, and the reported values of the *t*-statistic and R^2, one can calculate the coefficients and *t*-statistic for the regression in its standard form. This equation is *DP* = −.527 + 0.89 · *PS*; by definition the value of R^2 and the *t*-statistic (for β = 0) are the same. The *t*-statistic testing the hypothesis that β = 1 is −0.36; hence this evidence from Amoako-Adu et al. is actually consistent with efficient pricing.

worse than the disease. A simple approach would be to partition the data as done in Table 7. Any potential bias in these dimensions will be revealed in the appropriate partition of the sample. Golec and Tamarkin stick with the regression model, and seek to estimate the following equation:

$$DP = \alpha^I \cdot I + \alpha^H \cdot H + \alpha^F \cdot F + \beta \cdot PS + \varepsilon \quad (9)$$

where I is a vector of ones, H is one for home teams and zero otherwise, and F is one for favorites and zero otherwise. Either the home-away or favorite-underdog method of differencing this specification yields a singular matrix of regressors. Golec and Tamarkin's solution is to randomize the differencing method (adjusting the values of H, F, and PS accordingly) and estimate the parameters using the randomized data. Their procedure yields nonzero estimates of α^I in two specifications, leading Golec and Tamarkin to conclude that "unspecified biases" exist, "demonstrating that our statistical tests are more powerful" (p. 313). This statement makes no sense. The estimates of α^I are roughly the mean (conditioned on values of other regressors) of DP, where DP represents randomly determined point differences and hence has no meaningful interpretation.

William Dare and Scott McDonald (1996) address the differencing problem by taking appropriate account of the reciprocal relations inherent in point differences. The following equation is a simplified version of their model, appropriate for home-away point differences:

$$DP = \alpha^H H + \alpha^F Z + \beta PS + \varepsilon \quad (10)$$

In this equation Z is a composite favorite/underdog dummy variable which is 1 when the home team is favored and −1 when the home team is an underdog.[53] The simplest way to see how

such an equation can be estimated is to posit a primitive 0/1 dummy variable model for the points scored by a team i:

$$P_i = \gamma^H H_i + \gamma^A A_i + \gamma^F F_i + \gamma^U U_i + \varepsilon_i$$

where $A_i = 1$ when team i plays away from home and $\gamma^A A_i$ is the contribution of this factor to points scored by team i, with a similar interpretation of $\gamma^U U_i$, where $U_i = 1$ when team i is an underdog. Other variables are as previously defined.

The key is Dare and McDonald's imposition of symmetry on the home/away and favorite/underdog effects, so that $\gamma^U = -\gamma^F$ and $\gamma^A = -\gamma^H$. Symmetry clearly makes sense since the model focuses on the possibility of bias in point differences: if favorites are overbet by 2 points then underdogs must be underbet by two points. Let $DP = P_i - P_j$ where i is the home team and j is the visitor. Then $DP = \gamma^H(H_i - H_j) + \gamma^F(F_i - F_j) + \varepsilon_i - \varepsilon_j$. $(H_i - H_j) = 2$ under home-away differencing; $(F_i - F_j)$ is 2 when the home team is favored and −2 when the home team is an underdog. Straightforward transformations to this equation yield equation (10), with efficiency implying $\beta = 1$, $\alpha^H = \alpha^F = 0$.

Dare and McDonald estimate a version of (10) which allows for slope effects in addition to intercept shifts, along with the ability to handle games with $PS = 0$, and games played at neutral sites. In contrast to Golec and Tamarkin, Dare and McDonald's specification results in a failure to reject the efficiency null on these dimensions of the data.[54]

---

[53] For simplicity assume there are no "pick em" or neutral site games in the data. The full specifi-

cation in Dare and McDonald handles these games appropriately.

[54] Their sample of 6685 college games over the same period does yield an estimate of 0.77 points for the condition of being favored (α^F). The null is that this coefficient is zero, which leads Dare and McDonald to conclude that "lines of favored teams were downwardly biased by 0.77 points" ($t = 2.46$). This result probably stems from a violation

The misspecification charges levied at equation 5.4 by both Golec and Tamarkin and Dare and McDonald are overstated—the "differencing problem" is a red herring. The Dare and McDonald specification incorporates the favorite/underdog information in the model in an appropriate way and thereby adds power against the null, but only in this dimension of the data. Numerous issues, notably the relation of fundamental team characteristics to outcomes and point spreads, are omitted from these models. In pursuing these issues, the Dare and McDonald result can be relied on to simplify matters; i.e., it is appropriate to ignore the favorite/underdog status of the team, and home-away differencing can be used in these regressions.[55]

5.3 *Pricing of Fundamental Factors*

Zuber et al. (1985) investigated the market's incorporation of fundamental factors into point spreads of NFL football games. Their procedure contains three components. First, an OLS model is estimated in which the score difference of a game is explained by differences in the characteristics and contemporaneous performance of the opposing teams. These variables include differences in yards gained rushing, yards passing, fumbles lost, interceptions, penalties incurred, and two indexes of team strength: the number of rookies, and the prior number of wins. Denote the difference in characteristics between teams playing the nth game in week m by the vector $X_{mn} = X_{mn}^h - X_{mn}^a$, where X^h and X^a represent characteristics and performance attributes for the home team and away team, respectively. The contemporaneous least squares model of the score difference is then

$$DP_{mn} = \gamma + \delta X_{mn} + \omega_{mn}. \qquad (11)$$

where γ is a constant and ω_{mn} is a game-specific disturbance. Home-away differencing is used, hence the estimate of γ measures the home field advantage, and δ is the marginal effect of the performance measures on the score difference. The regression achieved a good fit (R^2 of .733) and the coefficient estimates are sensible.

The second component is an attempt to predict future score differences. Team-specific moving averages of passing yards, rushing yards, etc. from prior weeks are used to forecast X_{mn}^h and X_{mn}^a for the following week's games. These forecasts are combined with coefficient estimates from (11) to yield predicted score differences for subsequent games. This amounts to a simple fundamentals-based pricing model. The third component tests market efficiency with a simulation in which appropriate bets are placed when the market point spread differs from the model's prediction. Zuber et al. found that their method earned a profit during the 1983 NFL season. On this basis they reject efficient pricing: point spreads seemingly ignore relevant information contained in their list of fundamentals.

This result proved to be fragile. Sauer et al. (1988) showed that while

---

of the symmetry condition; blowouts are relatively common in college football, which can lead to this result. This conjecture is supported by Golec and Tamarkin's data, where they report that the percentage of winning bets on favorites was 49.8% in their sample of 6514 college games. Dare and McDonald do not report such a statistic. Dare and McDonald also refer to a possible violation of the efficiency condition for Super Bowl games, but this attention seems unwarranted given there are only 14 such games in their sample. As discussed earlier, the large variability of point differences renders samples of this size uninformative.

[55] As a general econometric matter, the imposition of true restrictions improves the efficiency of parameter estimation. Thus, unless there is evidence to expect a nonzero coefficient estimate, the "factor" *should* be excluded from the regression.

the contemporaneous regression model held up well out of sample, the Zuber et al. performance forecasts yielded no improvement over the point spread in predicting score differences. Denote the forecasts of X_{mn} by X^f_{mn}, and add this variable to the equation (8):

$$DP_{mn} = \alpha + \beta \cdot PS + \theta \cdot X^f_{mn} + \varepsilon_{mn} \quad (8')$$

Efficiency implies that the regressor X^f will not improve the fit of this equation. Hence the complete efficiency null is that $\alpha = 0$, $\beta = 1$, and $\theta = 0$. Sauer et al. fail to reject this restriction for both the 1983 and 1984 NFL seasons.[56] The inference is that the betting market makes efficient use of the information contained in teams' past performances. In addition, Sauer et al. show that the Zuber et al. betting method produced heavy losses out of sample. The profitability result may thus be spurious.

5.4 *Point Spread Variation and the Volatility Debate*

An important debate in financial economics concerns the ability of fundamentals-based pricing models to explain variation in the prices of financial instruments. It is well established that aggregate stock market returns exhibit volatility which is inconsistent with textbook models. John Cochrane (1991) shows that simple modifications can account for much of the excess volatility result. In contrast, Stephen LeRoy (1989) and Robert Shiller (1989) argue that the models based on fundamentals are incapable of explaining observed volatility.

The debate over the sources of excess volatility will be difficult to resolve, in part due to what LeRoy (1989) calls "the curse of dimensionality"—a change

in expectations tied to one of many factors at many future dates can cause the price of a financial asset to change. This multiplicity of possibilities makes it difficult to tie price changes to changes in expectations. Richard Roll's (1984, 1986) celebrated but failed attempts to obtain ex post explanatory power in empirical models of asset prices exemplify this problem.

This problem is greatly simplified in the wagering market, since here the set of relevant fundamentals is reduced to those affecting a single, easily observed outcome. Pricing issues in which expectations depend on the expectations of others, for example, are largely irrelevant in point spread markets. This provides a unique opportunity to study the relation between fundamentals, price changes, and outcomes. Sauer (1991), William Brown and Sauer (1993a), and Jim Dana and Michael Knetter (1995) exploit this feature using explicit models of score differences, and examine the restrictions these models impose on efficient point spreads.[57]

Sauer (1991) isolates a fundamental factor with observable changes and important consequences: injuries to star basketball players in the NBA. To investigate the effects of injuries on game scores and point spreads, Sauer estimated a fixed effects model of the point spread. The score difference for a game between teams i and j is assumed to be a function of four factors: luck, the home court advantage, the difference in the strengths of each team relative to the overall league, and an idiosyncratic factor relating to this particular matchup between teams i and j. Denote the home court advantage by α_H, team strengths by S_i and S_j, the idiosyncratic factor by ε, luck by ω, and let the

---

[56] Note that regressing the forecast error ($DP - PS$) on X^f (i.e. imposing the efficiency restrictions on α and β alone) yields negative values of RBAR^2 for both seasons.

[57] These models share a common structure first developed in the statistics literature by Harville (1975).

function $g(\cdot)$ map these factors into the score difference. Assuming i is the home team and $g(\cdot)$ is linear, the score difference is:

$$DP = \alpha^H + S_i - S_j + \varepsilon + \omega \qquad (12)$$

Using (6) and noting that luck cannot be anticipated,

$$PS = \alpha^H + S_i - S_j + \varepsilon \qquad (13)$$

A fixed effects/dummy variable model is used to obtain estimates of α^H and the team strength variables S. Efficiency implies that the coefficients in equations (12) and (13) are the same. These restrictions are satisfied by NBA data (Sauer 1991, Brown and Sauer 1993a).

Coefficient estimates from the PS equation are used to predict point spreads out of sample PSHAT. Sauer (1991) shows that PSHAT is an unbiased predictor of PS for games unaffected by injuries, and that the forecasts account for 80 to 90 percent of the variation in PS. It is important that the forecasts be accurate since heavy emphasis is placed on $PS -$ PSHAT. For the sample of games affected by player injuries, Sauer interprets this difference as the market's valuation of the injured player.

Now define the order of differencing such that team i has the injured player and team j is at normal strength (games with multiple injuries are excluded). For injury spells in which players miss a lengthy series of games, the mean of PS $-$ PSHAT is -2. That is, the estimated ability of the team with the injured player is two points lower, which seems a small amount for players of all-star caliber. However, PS is unbiased for these games, whereas PSHAT is significantly biased. Hence, the 2-point adjustment in PS represents an efficient response to the change in a fundamental factor.

For short injury spells, the adjust-ment in PS is smaller. The value of the player is roughly the same, however, so that PS is significantly biased in games which the injured player misses (the team does worse than predicted by the spread). Sauer shows that this bias changes magnitude and sign throughout the injury spell in ways that are predicted by a model in which player participation is uncertain. In addition, the participation probability implied by the bias closely matches the probability estimate obtained from a hazard function explaining the length of injury spells.[58] Hence, given the uncertainty of playing, the change in PS can be viewed as an efficient pricing response to the injury.

Brown and Sauer (1993a) examine the properties of $PS -$ PSHAT for non-injury games. Brown and Sauer consider two opposing hypothesis. One hypothesis is that $PS -$ PSHAT is meaningless noise. A second is that the residual represents unobserved fundamentals which are difficult to include in a systematic regression. Table 9 presents the mean forecast errors (MFE) of PSHAT conditional on the value of $PS -$ PSHAT. Define $\psi = PS -$ PSHAT as the idiosyncratic component in the market point spread. Consider cases in which ψ is 3 points. There are 211 such games in Brown and Sauer's sample, with a MFE(PSHAT) of 2.75 points. For these games, the idiosyncratic component is essential to unbiased prediction. The story is much the same for other values of ψ, although the information content in the tails is weak given the paucity of observations.

The predominance of positive values

---

[58] Injury spells are defined in the paper by games which players miss due to injury, which creates selection bias if the probability that the player plays in the game is non-zero. Hence, in assessing the question of efficiency the extent of the selection bias must be modeled explicitly, which is the approach taken in the paper.

TABLE 9
THE INFORMATION CONTENT OF THE
IDIOSYNCRATIC COMPONENT IN NBA POINT
SPREADS

ψ	Number of Observations	MFE (PSHAT)	Standard Error
0.0	432	−0.35	10.14
0.5	744	0.92	11.19
1.0	624	0.88	10.17
1.5	535	1.26	11.37
2.0	398	1.07	11.06
2.5	307	2.61	11.77
3.0	211	2.75	11.34
3.5	162	1.44	10.53
4.0	90	2.98	12.27
4.5	57	2.97	9.60
5.0	35	5.69	12.35
5.5	21	5.88	8.84
6.0	14	2.29	12.65
6.5	7	−0.74	7.05
7.0	5	8.00	4.85
7.5	8	4.50	13.69
≥8.0	4	−3.50	6.90

Source: Brown and Sauer 1993a. The data on scores and spreads are the same as in Table 5.1. ψ = PS − PSHAT is the idiosyncratic component of point spreads, based on out of sample forecasts using parameters obtained from estimating equation 13. Note that negative values of PS − PSHAT (and the corresponding forecast errors) have been multiplied by −1. This conserves table space and statistical power. For example, suppose that MFE(PSHAT) = −3 for observations where ψ = −4. The "adjustment" in the spread of −4 points accounted for by the unobserved component is thus 1 point too large. Converting these numbers to 3 and 4 points, respectively, yields the same interpretation and allows the positive and negative observations to be pooled.

of MFE(PSHAT) in Table 9 indicates that ψ does not represent meaningless noise; if it were, these values would be centered on 0. A fundamentals-based null hypothesis is that $[\psi − (DP − PSHAT) \mid \psi] = 0$, which requires that MFE(PSHAT) increase with the value of ψ. Brown and Sauer fail to reject this hypothesis. Brown and Sauer interpret these results in light of Roll's inability to explain stock price changes with a small scale model. Since the forecast errors in Table 9 tend to increase with ψ, this indicates that ψ represents unobserved fundamentals which are excluded from the team strength model.

Brown and Sauer also identify cases in which the estimates of team strength change from season to season, and show that the changes in these estimates are essential to unbiased prediction. Brown and Sauer do not model the adjustment in team ability estimates, however.

Dana and Knetter (1995) model this process explicitly using NFL data. Dana and Knetter examine how the market incorporates information from score differences in revising its estimates of team abilities. They estimate a version of equation (12) in which the team strength estimates follow a random walk. Efficient updating of these estimates is complicated by the low signal to noise ratio in score differences. Past score differences are used to estimate the team strength parameters, with an endogenously estimated threshold level beyond which increases in the score difference are discounted rather than attributed to relative ability. Other indicators of noise—net turnovers and penalties—are used as regressors in the model in an attempt to clean the ability estimates from these factors.

Dana and Knetter calculate a discount factor of .25 and a threshold of 8.3 points, implying that score differences beyond 8 points add about 1/4 of the information on relative ability compared to score differences within 8 points. The key question is whether market participants efficiently discount the noise in large score differences. Dana and Knetter conduct betting simulations both in and out of sample to examine this question. These simulations fail

to reach the threshold of profitability. By focusing on games in which noise may be a relatively large factor, Dana and Knetter do a bit better. They calculate a variable, *DPAR*, which stands for "difference in points attributable to randomness." When a current game involves two teams with a large value of *DPAR* from the prior week, their model's success at picking winning bets against the spread increases to the verge of profitability. Dana and Knetter interpret this as an "indication that learning is inhibited" by the noise inherent in game outcomes. Their conclusion stands in contrast to interpretations of excess volatility in financial markets, however. Dana and Knetter observe a form of "excess stability," where noise causes agents to "weight recent observations less than our statistical model suggests is optimal."

Ultimately, Dana and Knetter report "scant evidence" against efficiency. Regardless of this conclusion, what distinguishes their paper from run-of-the-mill tests is the explicit modeling of an important process—learning about a noisy world—and the pointed questions that can be asked when such a model is constructed. More work along these lines would be a welcome addition to the literature.

5.5 *Point Spread Changes During the Trading Period*

In the horse racing context, we have already seen how changes in odds during the trading period improve the betting market's estimates of the probability of winning. A natural question for the point spread market is whether trading generates similar improvement.

In the NFL market, Gandar et al. (1988) compared the mean square error of point spread forecasts using spreads from the beginning of the week with those from the end for the period 1980–85, and found virtually no difference in forecast accuracy. They also conducted an analysis of simple betting rules; the most successful are three behavioral rules which propose making wagers in opposition to recent public sentiment. Rule #5 reads as follows: "Bet on the team that becomes less favored (more of an underdog) over the course of the week's betting. In effect, one is generally betting against the direction bet by the majority of the public" (p. 1004). This betting rule achieved a success rate of 54.9 percent. The two other behavioral rules (similarly motivated) achieved similar success, all being statistically profitable in the sample. This evidence stands in contrast to the bulk of the evidence on point spread pricing thus far. One can think of several factors that may cause the efficient point spread to change over the course of a week.[59] Gandar et al.'s (1988) evidence is inconsistent with the proposition that NFL point spread changes are efficient responses to changes in fundamentals.

Gandar et al. (1998) examine intraday changes in point spreads for NBA games, and here the conclusion is quite different. Gandar et al. (1998) contrast the success of wagers made at the opening and closing spread with a method analogous to Brown and Sauer (1993a). They show that, conditional on the change in the spread, bets at the opening spread are profitable and depart significantly from the bounds implied by equation (5), whereas bets at the closing spread are a coin flip. In this market, observed point spread changes are necessary to eliminate the existence of

---

[59] For example, a textbook on sports book management asserts that bad weather reduces scoring and the score differences between good and bad teams, calling for adjustments in betting lines (Roxborough and Rhoden, p. 29). Changes in the status of injured players during the week also represent a change in fundamentals.

profit opportunities—changes in NBA spreads make them more accurate predictors of score differences.[60]

It is worth emphasizing that changes in NBA point spreads provide additional evidence on the importance of differentially informed agents. Bookmakers are experts who obviously have very strong incentives to be at least as well informed as the general betting public. Hence it is reasonable to conjecture that it is not the typical bettor, but rather individuals with scarce, highly specialized information that cause informative adjustments in NBA point spreads.

5.6 *The Hot Hand*

Gambling markets are especially fertile ground for contests between the behavioral and economic approaches to behavior. The behavioral approach emphasizes the effects of persistent misperceptions on individual choice and market outcomes. The economic approach emphasizes the role of economic forces in determining outcomes. For example, the average sports fan might have biased perceptions which overstate the probability that his team wins a game, but that observation does not establish that teams with more fans are poor bets in the point spread market.

Gilovich, Robert Vallone, and Amos Tversky (1985) make the case that behavior is influenced by biased perceptions, in particular, that belief in the hot hand—non-random streakiness in performance—is a "widely shared cognitive illusion" (p. 313). They examine field goal shooting by the members of the Philadelphia 76ers, free throw shooting by the Boston Celtics, and free throw shooting by Cornell University basket-

ball players, and in each case there is no evidence of serial dependence. Yet in the Cornell case, the players were willing to bet larger amounts at lower odds after making a shot, despite the fact that these shots were no more likely to be successful than shots after a miss. Belief in the hot hand, assert Gilovich et al., is belief in a myth.

Colin Camerer (1993) uses data on basketball betting to address the mythical hot hand argument in a repeated, nonexperimental setting. Camerer finds that bets on teams with recent winning streaks (against the point spread) are slightly more likely to be losers than winners. Camerer argues on this basis that the point spread market is an example where biased perceptions affect market prices. The proportion of winning bets, however, does not differ significantly from coin-flipping, hence the evidence is not fully persuasive.

Brown and Sauer (1993b) study the question using a point spread pricing model based on equations (12) and (13). The advantage offered by the pricing model is its ability to measure any adjustments in point spreads that stem from winning and losing streaks. Brown and Sauer find strong evidence that point spreads adjust by about 1/4 of a point (on average) for teams on a two or three game winning streak, and by about 2/3 of a point for teams on a streak of four or more games. The model thus confirms the market's belief that streaks may represent something real. Again, efficiency imposes equality restrictions on the coefficients in the point spread and score difference equations. Brown and Sauer fail to reject these restrictions, a result consistent with the existence of real streak effects. Note however that in single equation estimation, the score differences themselves provide no evidence against the hypothesis that streaks are purely

---

[60] The contrast between these results and those for the NFL market is a puzzle that calls for explanation.

random.[61] Brown and Sauer thus cannot reject the hypothesis that the hot hand is mythical.

Statisticians, with no axe to grind in the debate between the behavioral and economic approaches, have also examined the hot hand phenomenon. S. Christian Albright (1993) analyzed hitting by professional baseball players. The number of streaky hitters is consistent with a model in which hitting success is serially uncorrelated, with no persistence in the identity of who is streaky across seasons. Albright concludes that the stationary, purely random model of performance provides a "reasonably consistent" description of the data.

Jim Albert (1993) and Stern and Carl Morris (1993) agree that real (as opposed to random) streak effects are difficult to detect in hitting data. But they do not accept the argument that the hot hand is a myth. Stern and Morris argue that the problem is due to low power against the random null. Suppose that the hot hand were to raise the probability of a hit by as much as .1 during sub-periods of a season. Even with such a strong effect, their simulations show that Albright's methods lack power and likely will fail to reject the random null. Stern and Morris (p. 1189) argue that emphasis on the random null is misplaced in this context: "The evidence that all of us have obtained as participants and fans, which suggests that there are streaks in sports, should be preferred to the null hypothesis, at least until reasonably powerful approaches fail to reject plausible alternatives."

---

[61] An obvious problem for modeling which Brown and Sauer do not address is that random outcomes in an efficient market will generate a two game streak (either winning or losing) in about half of all two game sequences. An objective basis for ferreting out which streaks are candidates for non-randomness and which are not seems crucial for gaining power against the random null.

6. *Conclusion*

Wagering markets offer a unique setting for economists to study models of market pricing and choice under uncertainty. They provide a pricing mechanism for financial instruments in a context where outcomes are readily revealed and the scope of the pricing problem is reduced.

These features simplify the analysis of questions that are important in other contexts, and hence the findings of this literature have a bearing on economic issues elsewhere. Early researchers such as McGlothlin (1956) and Weitzman (1965) were not interested in the racetrack betting market per se, but in the information available from this market on choice under uncertainty. Although the literature on wagering has since become more specialized, the inherent advantages for attacking questions of general interest remain.

Most of the evidence surveyed here can be given a coherent equilibrium interpretation, and standard definitions of market efficiency are generally satisfied. The odds generated by racetrack betting efficiently predict the order of finish, and provide reasonably good estimates of the probability of winning, the favorite-long shot bias notwithstanding. As a rule, point spreads are efficient estimates of the median of the distribution of score differences. There is a strong tendency in these markets for price changes to move in the direction of outcomes.

Nevertheless, there are empirical regularities that are inconsistent with generic notions of efficiency. Although the representative agent model with convex preferences can rationalize the favorite-long shot bias, it fails to explain why this bias is not present in some markets. Differences in the rate of return associated with price changes are

inconsistent with representative agent models of asset pricing. In addition, systematic differences in the rates of return to wagering on the same event at different locations suggest the importance of differentially informed agents.

These findings present a challenge to equilibrium models of the wagering market. Models emphasizing heterogeneous agents, information asymmetries, and transaction costs have enjoyed some success in addressing these issues. More work along these lines is certainly called for and may advance our understanding of the forces that determine market prices.

REFERENCES

Albert, Jim. 1993. "Comment (on A Statistical Analysis of Hitting Streaks in Baseball)," *J. Amer. Statist. Assoc.*, 88:424, pp. 1184–88.

Albright, S. Christian. 1993. "A Statistical Analysis of Hitting Streaks in Baseball," *J. Amer. Statist. Assoc.*, 88:424, pp. 1175–83.

Ali, Mukhtar M. 1977. "Probability and Utility Estimates for Racetrack Bettors," *J. Polit. Econ.*, 85:4, pp. 803–15. Reprinted in Hausch et al., 1994.

———. 1979. "Some Evidence of the Efficiency of a Speculative Market," *Econometrica*, 47, 387–92. Reprinted in Hausch et al., 1994.

Amoako–Adu, Ben; Harry Marmer, and Joseph Yagil. 1985. "The Efficiency of Certain Speculative Markets and Gambler Behavior," *J. Econ. Bus.*, 37, pp. 365–78.

Applebaum, Elie and Eliakim Katz. 1981."Market Constraints as a Rationale for the Friedman–Savage Utility Function," *J. Polit. Econ.*, 89:4, pp. 819–25.

Arrow, Kenneth. 1951. "Alternative Approaches to the Theory of Risk in Risk-Taking Situations," *Econometrica*, 19, pp. 404–35.

Asch, Peter; Burton G. Malkiel, and Richard E. Quandt. 1982. "Racetrack Betting and Informed Behavior," *J. Finan. Econ.*, 10, pp. 187–94. Reprinted in Hausch et al., 1994.

———. "Market Efficiency in Racetrack Betting," *J. Business*, 1984, 57, pp. 165–175. Reprinted in Hausch et al., 1994.

———. 1986. "Market Efficiency in Racetrack Betting: Further Evidence and a Correction," *J. Bus.*, 59, pp. 157–60. Reprinted in Hausch et al., 1994.

Asch, Peter, and Richard E. Quandt. 1987. "Efficiency and Profitability in Exotic Bets," *Economica*, 54, pp. 289–98. Reprinted in Hausch et al., 1994.

Bassett, Gilbert W. Jr. 1981. "Point Spreads versus Odds," *J. Polit. Econ.*, 89:4, pp. 752–68.

Benjamin, Daniel K. and William R. Dougan. 1997. "Individual Assessments of the Risk of Death: A Reassessment of Previous Evidence," *J. Risk Uncertainty*, pp. 115–33.

Brenner, Reuven and Gabriele A. Brenner. 1990. *Gambling and Speculation: A Theory, a History, and a Future of Some Human Decisions*, New York: Cambridge U. Press.

Brown, William O. and Raymond D. Sauer. 1993. "Does the Basketball Market Believe in the Hot Hand? Comment," *Amer. Econ. Rev.*, 83:5, pp. 1377–86.

———. 1993. "Fundamentals or Noise? Evidence from the Point Spread Betting Market," *J. Finance*, 48:4, pp. 1193–209.

Busche, Kelly, and Christopher Hall. 1988. "An Exception to the Risk Preference Anomaly," *J. Bus.*, 61, pp. 337–46. Reprinted in Hausch et al., 1994.

Busche, Kelly. 1994. "Efficient Market Results in an Asian Setting," in *Efficiency of Racetrack Betting Markets*. D. B. Hausch, V. S. Y. Lo, and W. T. Ziemba, eds., San Diego: Academic Press.

Camerer, Colin. 1989. "Does the Basketball Market Believe in the Hot Hand?" *Amer. Econ. Rev.*, 79:5, pp. 1257–61.

Clotfelter, Charles T. and Phillip J. Cook. 1989. *Selling Hope: State Lotteries in America*, Cambridge, MA: Harvard U. Press.

Cochrane, John H. 1991. "Volatility Tests and Efficient Markets: A Rev. Essay," *J. Monet. Econ.* 27, 463–85.

Conlisk, John. 1993. "The Utility of Gambling," *J. Risk Uncertainty*, 6:3, 255–75.

Crafts, N.F.R. 1985. "Some Evidence of Insider Knowledge in Horse Race Betting in Britain," *Economica*, 52, pp. 295–304.

———. 1994. "Insiders and Outsiders in British Horse Race Betting," in Hausch, Lo, and Ziemba., pp. 545–49.

Dana, Jim, and Michael Knetter. 1994. "Learning and Efficiency in a Gambling Market," *Manage. Sci.*, 40:10, pp. 1317–28.

Dare, William H., and S. Scott McDonald. 1996. "A Generalized Model for Testing the Home and Favorite Team Advantage in Point Spread Markets," *J. Finan. Econ.*, 40, pp. 295–318.

Dolbear, F. Trenery. 1993. "Is Racetrack Betting on Exactas Efficient?" *Economica*, 60, pp. 105–11.

Dowie, Jack. 1976. "On the Efficiency and Equity of Betting Markets," *Economica* 43, pp. 139–50. Reprinted in Hausch et al., 1994.

Dwyer, Gerald P., Jr. et al. 1993. "Tests Of Rational Expectations In A Stark Setting," *Econ. J.*, 103, pp. 586–601.

Even, William, and Nicholas Noble. 1993. "Testing Efficiency in Gambling Markets," *Applied Econ.*, 24, pp. 85–88.

Fama, Eugene F. 1970. "Efficient Capital Markets: A Rev. of Theory and Empirical Work," *J. Finance*, 25:2, pp. 383–417.

Figlewski, Stephen. 1979. "Subjective Information and Market Efficiency in a Betting Market." *J. Polit. Econ.*, 87:1, pp. 75–89.

Findlay, John M. 1986. *People of Chance: Gambling in American Society from Jamestown to Las Vegas*, Oxford U. Press: New York.

Friedman, Milton and Leonard Savage. 1948. "The Utility Analysis of Choices Involving Risk," *J. Polit. Econ.*, 56:4, pp. 279–304.

Gabriel, Paul E. and James R. Marsden. 1990. "An Examination of Market Efficiency in British Racetrack Betting," *J. Polit. Econ.*, 98 (4), pp. 874–85.

———. 1991. "An Examination of Market Efficiency in British Racetrack Betting: Errata and Corrections," *J. Polit. Econ.*, 99:3, pp. 657–59.

Gandar, John; W. Dare, Craig Brown and R. Zuber. 1998. "Informed Traders in the Betting Market for Professional Basketball Games," *J. Finance,* 53, pp. 385–401.

Gandar, John; Richard Zuber, Thomas O'Brien, and Ben Russo. 1988. "Testing Rationality in the Point Spread Betting Market," *J. Finance,* 43, pp. 995–1007.

Gandar, John; Richard Zuber, and Ben Russo. 1993. "Testing Efficiency in Gambling Markets: A Comment," *Appl. Econ.*, 25, pp. 937–43.

Gilovich, Thomas. 1983. "Biased Evaluation and Persistence in gambling," *J. Personality and Psych.*, 44:6, pp. 1110–26.

———, Robert Vallone, and Amos Tversky. 1985. "The Hot Hand in Basketball: On the Misperception of Random Sequences," *Cognitive Psych.*, 17, pp. 295–314.

Golec, Joseph, and Maurry Tamarkin. 1991. "The Degree of Inefficiency in the Football Betting Market," *J. Finan. Econ.*, 30, pp. 311–23.

Griffith, Richard M. 1949. "Odds Adjustment by American Horse Race Bettors," *Amer. J. Psych.*, 62, pp. 290–94. Reprinted in Hausch et al., 1994.

Harris, H. A. 1972. *Sport in Greece and Rome*. Ithaca, N.Y.: Cornell U. Press.

Harville, David A. 1973. "Assigning Probabilities to the Outcomes of Multi–entry Competitions," *J. Amer. Statist. Assoc.*, 68, pp. 312–16. Reprinted in Hausch et al., 1994.

———. 1975. "Predictions for NFL Games via Linear Model Methodology," *J. Amer. Statist. Assoc.*, 75, pp. 516–24.

Hausch, Donald B.; Victor S. Y. Lo, and William T. Ziemba, eds. 1994. *Efficiency of Racetrack Betting Markets*, San Diego: Academic Press.

Hausch, Donald B.; William T. Ziemba, and Mark Rubinstein. 1981. "Efficiency of the Market for Racetrack Betting," *Manage. Sci.*, 27:12, pp. 1435–52. Reprinted in Hausch et al., 1994.

Hausch, Donald B. and William T. Ziemba. 1990. "Arbitrage Strategies for Cross-Track Betting on Major Horse Races," *J. Bus.*, 63, pp. 61–78. Reprinted in Hausch et al., 1994.

———. 1985. "Transactions Costs, Extent of Inefficiencies, Entries, and Multiple Wagers in a Racetrack Betting Model," *Manage. Sci.*, 31, pp. 381–94. Reprinted in Hausch et al., 1994.

———. 1992. "Efficiency of Sports and Lottery Betting Markets," in *The New Palgrave Dictionary of Money and Banking,* P. Newman, M. Milgate, and J. Eatwell, eds. London: Macmillan.

———. 1994. "Efficiency of Sports and Lottery Betting Markets," in R. Jarrow, V. Maksimovic, and W. T. Ziemba, eds., *Finance Handbook*, Elsevier-North Holland.

Hirshleifer, Jack. 1966. "Investment Decision Under Uncertainty: Applications of the State-Preference Approach," *Quart. J. Econ.*, 66, pp. 252–77.

Hoerl, Arthur E. and Herbert K. Fallin. 1974. "Reliability of Subjective Evaluations in a High Incentive Situation," *J. Roy. Statist. Soc.*, 137, pp. 227–30.

Humphrey, John H. 1986. *Roman Circuses: Arenas for Chariot Racing*, U. California Press, Berkeley.

Hurley, William and Lawrence McDonough. 1995. "A Note on the Hayek Hypothesis and the Favorite Long Shot Bias in Parimutuel Betting," *Amer. Econ. Rev.*, 85:4, pp. 949–55.

———. 1996. "The Favorite Long Shot Bias in Parimutuel Betting: A Clarification of the Explanation that Bettors Like to Bet Long Shots," *Economic Letters*, 52:3, pp. 275–78.

———. 1996. "An Unbiased Explanation of the Favourite Long Shot Bias in Parimutuel Betting," manuscript, The Royal Military College of Canada.

Issacs, Rufus. 1953. "Optimal Horse Race Bets," *Amer. Math. Monthly*, 60, pp. 310–15, Reprinted in Hausch et al., 1994.

LeRoy, Stephen F. 1989. "Efficient Capital Markets and Martingales," *J. Econ. Lit.*, 27, pp. 1583–621.

Longrigg, Robert. 1972. *The History of Horse Racing*, NY: Stein and Day.

Losey, Robert L. and John C. Talbott, Jr. 1980. "Back on the Track with the Efficient Markets Hypothesis," *J. Finance*, 35:4, pp. 1038–43.

Machina, Mark J. 1987. "Choice Under Uncertainty: Problems Solved and Unsolved," *J. Econ. Perspectives*, 1:1, pp. 121–54.

MacLean, L. C., W. T. Ziemba, and G. Blazenko. 1992. "Growth Versus Security in Dynamic Investment Analysis," *Manage. Sci.*, 38, pp. 1562–85.

Markowitz, Harry. 1952. "The Utility of Wealth," *J. Polit. Econ.*, 56:4, pp. 151–58.

McGlothlin, W. H. 1956. "Stability of Choices Among Uncertain Alternatives," *Amer. J. Psych.*, 69, pp. 604–15. Reprinted in Hausch et al., 1994.

Pankoff, Lynn. 1968. "Market Efficiency in Football Betting," *J. Bus.*, 41, pp. 203–14.

Pope, Peter F. and David A. Peel. 1989. "Information, Prices, and Efficiency in a Fixed-Odds Betting Market," *Economica*, 56, pp. 322–41.

Quandt, Richard E. 1986. "Betting and Equilibrium," *Quart. J. Econ.*, 101:1, pp. 201–207.

Ritter, Jay. 1994. "Racetrack Betting—An Example of a Market with Efficient Arbitrage," in Hausch et al.

Robertson, William H. P. 1964. *The History of Thoroughbred Racing in America*, Englewood Cliffs, NJ: Prentice Hall.

Roll, Richard. 1984. "Orange Juice and Weather," *Amer. Econ. Rev.,* 74, pp. 861–80.

———. 1988. "R2," *J. Finance*, 43, pp. 541–66.

Rosett, Richard N. 1975. "Gambling and Rationality," *J. Polit. Econ.*, 73:6, pp. 595–607. Reprinted in Hausch et al., 1994.

Roxborough, Michael and Mike Rhoden. 1991. *Race and Sports Book Management.* Las Vegas: Las Vegas Sports Consultants.

Russo, Ben J.; John Gandar, and Richard Zuber. 1989. "Market Rationality Tests Based on Cross-Equation Restrictions," *J. Monet. Econ.*, 24, pp. 455–70.

Samuelson, Paul. 1952. "Probability, Utility, and the Independence Axiom," *Econometrica*, 20, pp. 670–78.

Sauer, Raymond D. 1991. "Do Price Changes Accurately Reflect Changes in the Distribution of Outcomes? A Study of the Effects of Player Injuries on Basketball Scores and Spreads," unpublished manuscript, Clemson U.

Sauer, Raymond D.; Vic Brajer, Stephen Ferris, and M. Wayne Marr. 1988. "Hold Your Bets: Another Look at the Efficiency of the Betting Market for NFL Games," *J. Polit. Econ.*, 96:1, pp. 206–13.

Schnytzer, Adi; and Yuval Shilony. 1995. "Inside Information in a Betting Market," *Econ. J.,* 105, pp. 963–71.

Shiller, Robert. 1989. *Market Volatility*. Cambridge, MA: MIT Press.

Shin, Hyun Song. 1991. "Optimal Odds Against Insider Traders," *Econ. J.*, 101, pp. 1179–85.

———. 1992. "Prices of State Contingent Claims With Insider Traders, and the Favourite Long Shot Bias," *Econ. J.*, 102, pp. 426–35. Reprinted in Hausch et al., 1994.

———. 1993. "Measuring the Incidence of Insider Trading in a Market for State–Contingent Claims," *Econ. J.*, 103, pp. 1141–53.

Slovic, Paul; Baruch Fischoff, and Sarah Lichtenstein. 1982. "Facts versus Fears: Understanding Perceived Risk," in *Judgement Under Uncertainty: Heuristics and Biases*. D. Kahneman, P. Slovic, and A. Tversky, eds., Cambridge: Cambridge U. Press.

Smith, Vernon L. 1971. "Economic Theory of Wager Markets," *Western Econ. J.*, 9, pp. 242–55.

Snyder, Wayne W. 1978. "Horse Racing: Testing the Efficient Markets Model," *J. Finance*, 33, pp. 1109–18. Reprinted in Hausch et al., 1994.

Stern, Hal S. 1990. "Models for Distributions on Permutations," *J. Amer. Statist. Assoc.*, 85, pp. 558–64.

——— and Carl N. Morris. 1993. "Comment (on A Statistical Analysis of Hitting Streaks in Baseball)," *J. Amer. Statist. Assoc.*, 88:424, pp. 1189–93.

Thaler, Richard H. and William T. Ziemba. 1988. "Anomalies—Parimutuel Betting Markets: Racetracks and Lotteries," *J. Econ. Perspect.*, 2, pp. 161–74. Reprinted in Hausch et al., 1994.

Tryfos, P.; S. Casey, S. Cook, G. Leger, and B. Pylypiak. 1984. "The Profitability of Wagering on NFL Games," *Manage. Sci.,* 30:1, pp. 123–32.

Vamplew, Wray. 1976. *The Turf. A Social and Economic History of Horse Racing*. London: Allen Lane.

Vaughan Williams, Leighton and David Paton. 1997. "Why is There a Favourite Long Shot Bias in British Racetrack Betting Markets," *Econ. J.*, 107, pp. 150–58.

Vergin, Roger C. and Michael Scriabin. 1978. "Winning Strategies for Wagering on National Football League Games," *Manage. Sci.*, 24, pp. 809–18.

Viscusi, W. Kip and Charles O'Connor. 1984. "Adaptive Responses to Chemical Labeling: Are Workers Bayesian Decision Makers?" *Amer. Econ. Rev.*, pp. 942–56.

Weitzman, Martin. 1965. "Utility Analysis and Group Behavior: An Empirical Study," *J. Polit. Econ.*, 73:1, pp. 18–26. Reprinted in Hausch et al., 1994.

Woodland, Linda M. and Bill M. Woodland. 1991. "The Effects of Risk Aversion on Wagering: Point Spread versus Odds," *J. Polit. Econ.*, 99:3, pp. 638–53.

———. 1994. "Market Efficiency and the Favorite Long Shot Bias: The Baseball Betting Market," *J. Finance*, 49:1, pp. 269–80.

Ziemba, William T. and Donald B Hausch. 1987. *Dr. Z's Beat the Racetrack*. New York: William Morrow.

Zuber, Richard; John M. Gandar, and Benny Bowers. 1985. "Beating the Spread: Testing the Efficiency of the Gambling Market for NFL Games," *J. Polit. Econ.*, 93:4, pp. 800–806.