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

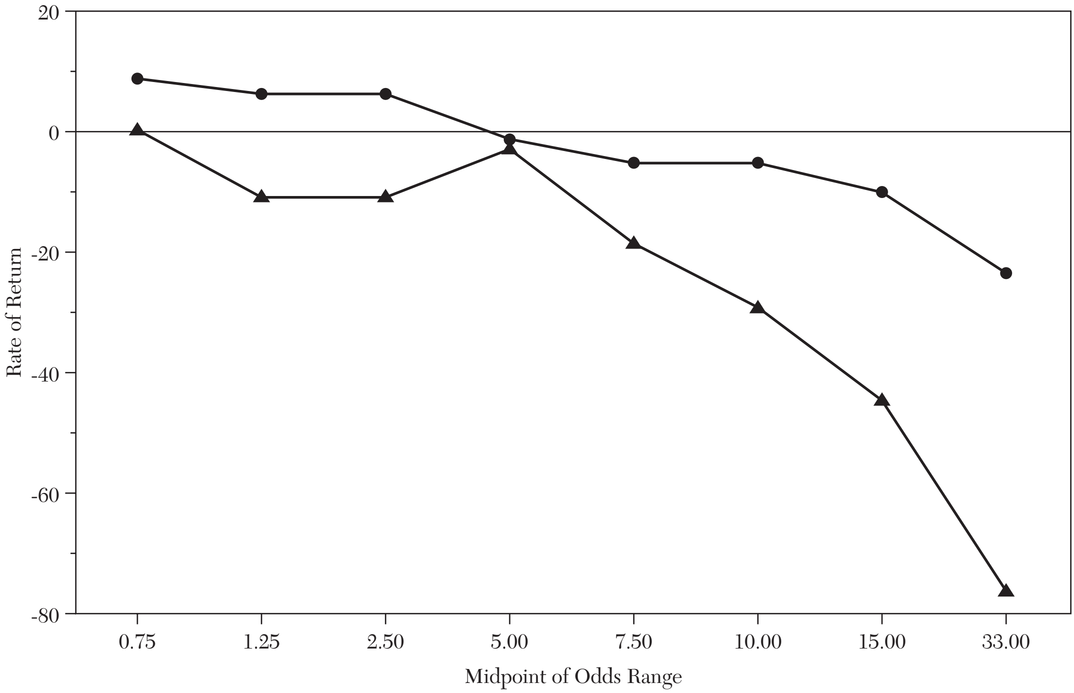

Figure 1. The Favorite-Longshot Bias in the U.S. Pari-mutuel (Snyder) and U.K. Bookmaking (Dowie) Markets
Note: Snyder's are pari-mutuel returns with the takeout added back; hence the norm is zero. Dowie's are pre-tax returns at SP odds. Hence the norm for Dowie is less than zero to account for the costs 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 before-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

	Ali (1977)		Asch, Malkiel, & Quandt (1982)		Busche & Hall (1988)	
Rank	$\hat{p}_k$	$w_k - \hat{p}_k$	$\hat{p}_k$	$w_k - \hat{p}_k$	$\hat{p}_k$	$w_k - \hat{p}_k$
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 4^{th}–8^{th} 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.

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_{FINAI}/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_{FINAI}/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 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

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.

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

$$E(DP - PS \mid \Omega) = 0 \tag{7}$$

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

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

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 $[\psi - (DP - \text{PSHAT}) \mid \psi\ ] = 0$, 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