

**Do Financial Incentives affect the Quality of Expert Performance?
Evidence from the Racetrack***

Glenn Boyle
NZ Institute for the Study of Competition and Regulation
Victoria University of Wellington
PO Box 600, Wellington
New Zealand
glenn.boyle@vuw.ac.nz

25 May 2007

Keywords: incentives, performance, quality, experts, racehorses

JEL Categories: I28, J44, M52

* For helpful comments on early drafts of this paper, I am grateful to Graeme Guthrie, Philip McCann, Jacques Poot, and seminar participants at Waikato University. I am also indebted to Victor Rolton for assisting with the data, and to Mike Webb and Eli Grace-Webb for valuable research assistance. All remaining errors and ambiguities are my responsibility.

Do Financial Incentives affect the Quality of Expert Performance? Evidence from the Racetrack

Abstract

Does the quality of performance by experts respond to financial incentives? Or, as some psychologists argue, are experts primarily motivated by more intrinsic considerations such as professional pride? I provide some evidence on this question by examining the relationship between horse race outcomes and the level of race prize money. If financial incentives are important, then races with low prize money are more likely to see some trainers exert less than full effort, thereby upsetting the calculations of race bettors. By contrast, races with high prize money are less likely to be affected by unobservable variation in trainer effort, so bettor odds should then be a more reliable predictor of race outcomes. In a sample of 30426 horse races, I find evidence consistent with this story: average bettor payoffs in a variety of betting pools are strongly negatively related to race prize money, and the probability of a bettor-favourite horse succeeding is strictly increasing in the amount of prize money at stake. These results continue to hold when I exclude low-information races from the sample, thereby suggesting that prize money is not acting as a proxy for the quantity of information publicly available to bettors. As a group, horse trainers apparently tailor the quality of their performance to the potential size of their payout from clients.

Do Financial Incentives affect the Quality of Expert Performance? Evidence from the Racetrack

1. Introduction

A number of recent studies document a strong output response to the provision of financial incentives for workers. For example, Lazear (2000), Paarsch and Shearer (2000), and Shearer (2004) all find that productivity is significantly greater when workers are paid piece rates rather than fixed wages. This evidence suggests that the potential for greater financial returns motivates greater effort and hence enhances performance quantity, just as economic theory predicts.

By contrast, there is considerably more doubt about the sensitivity of performance *quality* to incentives - the extent to which greater payoffs motivate workers to produce *better*, as well as more, goods and services. Psychologists such as Kohn (1993a, 1993b) reject the idea that money can motivate quality-oriented tasks, a view generally supported by laboratory evidence: in a meta-analysis of 39 experiments undertaken between 1975 and 1996, Jenkins et al (1998) are unable to detect any significant relationship between financial incentives and performance quality.¹

The relationship between performance quality and financial incentives is particularly germane to the class of professional workers who use specialized knowledge, information and skills to provide quality-oriented expert services to clients. Because the actions undertaken by these experts are often difficult to observe and monitor, they can potentially tailor their effort to the magnitude of the potential payoff. For example, financial advisor firms often receive fees based on the value of client portfolios, and so have a strong incentive to devote more care and attention to high-wealth clients than to those of lesser means. Similarly, a law firm stands to gain more from a class-action lawsuit filed against a defendant with deep pockets than from a case involving a single client who wishes to dispute a parking ticket - strict financial considerations would suggest that the quality of representation is therefore likely to be higher in the former case. However, experts are typically members of a profession and thus are likely to have high intrinsic motivation related to this membership; such "professional pride" considerations should encourage experts to treat all clients equally and thus dampen the impact of financial incentives on performance quality.

Because data on expert performance and payoffs are frequently unavailable and, in the case of performance, often difficult to measure as well, empirical evidence on the economic behaviour of experts is limited to just a few studies. Levitt and Syverson (2005) and Rutherford et al (2005) find that real estate agents perform better when selling their own houses than when selling on behalf of

¹ Bonner et al (2000) and Camerer and Hogarth (1999) report broadly similar findings.

clients, but Gompers and Lerner (1999) report no difference between the secondary market performance of initial public offerings underwritten by affiliated banks and those underwritten by independent banks. Gruber and Owings (1996), Crosson et al (2001) and Gaynor et al (2004) also document examples of health professionals responding to financial incentives, although these relate more to the quantity and type of services provided rather than the overall quality of performance. Finally, Hubbard (1998) finds that auto mechanics in private firms adopt a more lenient emissions test policy than do government-employed inspectors, a finding he attributes to a desire for repeat business by the former group. Although these studies suggest that experts may tailor their performance to the potential size of the payoff they receive, and that they sometimes favour their own financial interests over those of clients, it remains uncertain whether they also serve their own interests by favouring some clients over others.

A setting that is potentially capable of providing evidence on this issue is the training of race horses. Trainers are experts who prepare horses for racing on behalf of clients, and in return receive daily fees plus a percentage of horse winnings (10% in the setting examined here). Because 10% of high winnings is worth more to trainers than 10% of low winnings, they have a financial incentive to provide better quality training services to clients whose horses compete in races that offer the greatest stakes. And, unlike most other experts, the results (race outcomes) and potential payoffs (race stakes) of trainer endeavour are both extensive and publicly available.

If trainers are motivated by financial considerations, then they can be expected to work harder and smarter leading up to races that offer high stakes. But if they are motivated by professional pride considerations, then they should exert the same "effort" for all races. However, the tournament nature of race outcomes creates a significant difficulty in trying to distinguish between these two hypotheses: even if all trainers try harder and provide higher-quality service in preparing their horses for a high-stakes race, only one horse can win that race.

One way around this problem is to use an absolute measure of performance such as the time taken to run a particular distance, i.e., if trainers provide better quality effort in preparing for high-stakes races, then a given horse should run faster in such a race than in another race where the prize money is lower. However, the sensitivity of race times to weather, track surface, and particularly, race tactics, means that such an approach is unlikely to be useful.

An alternative solution is to use the information contained in race betting odds in order to indirectly infer the relationship between race stakes and trainer performance quality. In handicapping any race, the betting public uses a vast array of information: expert opinion and advice, perceived horse quality based on prior performance, trainer quality, driver quality, barrier draw, and so on.

Most studies of racetrack betting conclude that markets are broadly efficient with respect to such publicly-available information (for overviews, see Hausch et al, 1994, and Sauer, 1998), but various anomalies suggest that private information may not be fully incorporated in odds (e.g., Crafts, 1985; Gabriel and Marsden, 1990; Shin, 1993). In this context, a particularly important piece of private information is the extent to which a horse's trainer exerts maximum effort in preparing it for that race. A trainer might, for example, under-work a horse in the days leading up to the race, or entrust the horse's care to a junior stablehand, or obtain the services of a lower-quality raceday driver or jockey. In general, such "inside information" is unobservable to bettors, and so little of it is impounded in betting odds.²

Why might a trainer exert less-than-full effort? In addition to the usual argument that effort is costly in terms of foregone leisure, three other possibilities stand out. First, the trainer may have other horses under his care and so must allocate his available time among these competing interests. Second, the trainer may be focussing on more valuable future races for the same horse and is thus using the current race largely as a practice run. Third, there may be simple corruption motivations. For example, the trainer intends to bet on another trainer's horse in the same race and so wishes to maximise the probability of that bet succeeding by making his own horse less of a factor. Or he hopes to reap a large betting payoff on the horse in a future race by first improving its odds with a lacklustre performance in the current race.³ Regardless of which of these motivations applies in any particular case, all should diminish in intensity as race stakes become higher: the opportunity cost of leisure increases, the opportunity cost of devoting time to other horses and races increases, and the risk of incurring corruption penalties becomes more difficult to justify.

Suppose that trainers are indeed more likely to choose lower effort levels in low-stakes races, and that this choice is unobservable to bettors. Then bettor odds - which primarily reflect publicly available information - will be a less reliable predictor of the outcomes of such races, since bettors are missing relevant information. By contrast, private information is less significant in high-stakes races where all trainers exert full effort, hence ensuring a closer convergence between bettor odds and race outcomes. In short, race outcomes become more "predictable" to bettors (given the information they have available) as race stakes increase, and favoured horses - which offer smaller bettor payoffs - succeed more often.

² Of course, insiders (employees and other connections of the trainer's stable) will often have this information and can be expected to utilise it. However, such investments are typically only a small proportion of the total amount bet - Coleman (2007) estimates approximately 2% - and, as previous studies indicate, do not seem to be sufficient to fully incorporate the information in betting odds.

³ Shin (1992) cites an acknowledgement by United Kingdom racing officials of the prevalence of this practice.

This situation is depicted in Figure 1. The two circles represent the total amounts of information relevant to predicting the outcome of low- and high-stakes races respectively. In the low stakes race, some trainers exert less-than-full effort because of preferred opportunities elsewhere, and so a relatively large proportion of the relevant information is unavailable to bettors. As a result, some horses will be over-bet and others under-bet, and unexpected outcomes (i.e., high payoffs to successful bettors) will be relatively common. But in the high-stakes race, all trainers do their best, so the quantity of private information is relatively small. Bettor odds thus reflect a greater proportion of the relevant information and unexpected race outcomes are less common.

[Insert Figure 1 about here]

To explore this idea, I test the following two hypotheses: if horse trainers respond to financial incentives by exerting greater effort in races with high stakes, then (i) the size of the payoffs (per unit of investment) received by successful bettors should decrease in race stakes, and (ii) the success rate of betting "favourites" (i.e., the horses with the most bettor money invested on them) should increase in race stakes. Put more succinctly, horses that are relatively heavily backed by bettors, and thus offer lower payoffs, should succeed more often in high-stakes races.

In a sample of 30426 horse races, I detect a strong relationship between race stakes and payoffs to successful bettors: after controlling for other factors that might affect race outcomes, a doubling of the race stake is associated with a decrease in bettor payoffs of between 8% and 14%, depending on the betting pool. Similarly, the success rates of bettor-favourite horses increase significantly with the size of the race stake. Importantly, I obtain essentially the same results in sub-samples that exclude races about which bettors seem likely to have relatively little information, thereby suggesting that race stake is not simply proxying for the quantity and quality of bettor information about a race. As a group, horse trainers apparently tailor the quality of their performance to the potential size of their financial payoff.

In the next section, I describe the data used in this paper and undertake some preliminary analysis. Section 3 contains the results from a series of multiple regression models that are the paper's primary focus. Section 4 offers some concluding remarks.

2. Data and Preliminary Analysis

My data sample consists of 30426 standardbred (harness) horse races contested in New Zealand (NZ) between 1993 and 2006 for which full data are available.⁴ These races were staged at

⁴ A further 330 races had missing or non-existent data, and so were deleted from the sample.

56 distinct tracks spread across the entire length of NZ, but with more than a third of them taking place at the two major tracks located in Auckland and Christchurch. Since 1996, NZ has had both fixed-odds (bookmaker) and totalisator (pari-mutuel) options available, although most betting takes place via the latter. Total betting on the races in my sample was approximately \$3.5 billion.

For each race, I recorded both the total prizemoney at stake and the payoffs to successful investors in three totalisator betting pools: win, quinella and trifecta. The win pool requires bettors to select the winning horse; the quinella pool requires selection of the first and second horses in either order; the trifecta pool requires selection of the first three horses in the correct order. Moving from win to quinella to trifecta payoffs yields an increasingly broad measure of the extent to which a race corresponds to bettors' expectations. For example, a strongly favoured horse could win the race and yield a small win pool payoff - indicating that the result was highly predictable - but be followed home by two outsiders resulting in a large trifecta pool payoff - indicating that the result was unexpected. By using all three payoff measures, I am able to better assess whether the hypothesized relationship between betting payoffs and race stakes is robust. I also recorded the success of the betting favourite (i.e., the horse that attracted the most investment) in two pools: win and place. The win favourite was deemed to be successful if and only if it won; the place favourite was successful if and only if it finished in the first three.⁵

Table 1 provides some summary statistics for the variables used in this study. Betting payoffs are expressed relative to a \$1 unit of investment, i.e., a \$10 payoff equals \$1 return of investment plus \$9 winnings.⁶ The three payoff variables display considerable dispersion, e.g., the mean trifecta payoff is \$1055, but the sample contains payoffs ranging from \$2.10 to almost \$400,000. Race stakes are similarly variable: the mean stake is \$8019, but individual stakes range from \$1000 to \$750,000. All four variables exhibit significant skewness, with means that are considerably greater than medians.⁷ Win favourites win approximately one third of all races, while place favourites successfully finish in the first three 61% of the time.

[Insert Table 1 about here]

To investigate the relation between race stakes and betting payoffs (or favourite success rates), I control for variables that potentially affect the payoff measures I use. An obvious candidate here is the number of horses in the race, since a bigger field makes it less likely that any given horse - or

⁵ Note that betting terminology varies across countries. In the United States, for example, a place bet requires a horse to finish in the first two, while an investment that requires finishing in the first three (i.e., a NZ place bet) is known as a show bet.

⁶ All monetary values are in New Zealand dollars.

⁷ The Totalisator Agency Board (the sole agency for betting on horse races in NZ) guarantees a minimum payoff of \$1.

combination of horses - will succeed. All else equal, this should have a significant impact on average payoffs, particularly in the quinella and trifecta pools.⁸

I also control for the inherent "unpredictability" of a race. Regardless of stake, some races are simply more difficult to handicap than others, resulting in a greater number of unexpected outcomes that are associated with lower favourite success rates and higher betting payoffs. For standardbred horse races, three factors stand out. First, although most races in my sample are run on all-weather surfaces, a significant minority (approximately 12.5%) are run on grass. Grass track races are often more difficult to handicap for several reasons: horses are less experienced at racing on grass, bettors have less information about horses' relative prowess on grass, and grass tracks offer a less homogeneous racing surface that may randomly affect some horses in any race. Second, approximately 10% of the races in the sample are run on rain-affected tracks. This affects different horses in different ways, but tends to neutralise the advantages of high-speed horses, thus making the outcome more difficult for bettors to predict. Third, although the majority of races in my data begin from a mobile dispatch (where the horses are already moving at close to race speed at the point where the race starts), approximately 44% are standing start races (where the horses must stand quietly and begin from a stationary position). The latter feature introduces an additional random element because horses are more likely to deviate from the correct gait (i.e., racing action) when forced to begin quickly from a stationary position, potentially resulting in considerable loss of ground or disqualification.

To quantify these features in a form amenable to empirical analysis, I construct an index that adds 1 for each of the above features that a race possesses.⁹ For example, a race that is run on a wet grass surface from a standing start has an index value of 3, while a race that is run on a dry grass surface from a mobile start is given an index value of 1. Races with none of these features have an index value of 0.

Table 2 provides some preliminary analysis of the relation between race stakes and betting payoffs. I first split the sample into races above and below the median stake and calculate average betting payoffs and favourite success rates for each group. Average payoffs are lower, and favourite success rates are higher, in higher-stakes races, and these differences are strongly statistically significant (at the 0.1% level or better). Quantitatively, the differences are greatest in the more exotic

⁸ For example, increasing the number of horses in the race from 10 to 12 raises the number of possible quinella combinations from 45 to 66, and the number of possible trifecta combinations from 720 to 1320. In my sample, 70% of the races contain between 9 and 14 horses, but as Table 1 indicates, some take on more extreme values.

⁹ I also estimated models that include these three features as separate variables. However, this has no noticeable effect on the race stake coefficient, so I report only the results from the more parsimonious models that include the index as an explanatory variable.

bet pools - the average trifecta dividend in races with stakes above the sample median is only 82% of the average dividend in races with stakes below the sample median.

[Insert Table 2 about here]

I also calculate the correlation between, on the one hand, the race stake, and on the other hand, the payoff and success rate variables. Partly to eliminate the skewness, and partly because any impact of stakes on betting payoff seems likely to be non-linear, I use the natural log of the race stake and bettor payoff variables for this purpose. As the last column of Table 2 reveals, race stakes are strongly negatively correlated with bettor payoffs and positively correlated with favourite success rates (again at the 0.1% level or better).

At the same time, however, both field size and race unpredictability are (i) greater on average in the low-stakes races, and (ii) negatively correlated with race stakes. Thus, although the above results yield the conclusion that betting payoffs are negatively related to race stakes, consistent with the view that bigger financial incentives induce greater effort and better performance from trainers, this relationship could be due to the negative relation between stakes and the two control variables. To disentangle these effects, I employ multivariate regressions that allow me to estimate the relation between race outcomes and stakes while controlling for these other factors.

3. Regression Analysis

I estimate regression models of the general form

$$y = \alpha_0 + \alpha_1(\ln \text{ race stake}) + \alpha_2(\text{field size}) + \alpha_3(\text{unpredictability index}) + \varepsilon \quad (1)$$

In some specifications, the dependent variable y equals the natural log of one of the three betting payoff variables. In other specifications, y is an indicator variable: either (i) set equal to 1 if the win pool favourite wins and 0 otherwise, or (ii) set equal to 1 if the place pool favourite finishes in the first three and 0 otherwise. The first set of models are estimated by ordinary least squares with White (1980) heteroskedasticity-robust standard errors; the second set of models are estimated using maximum-likelihood logit regression.

I also estimate an extension of (1) that incorporates fixed effects for each race track. Some tracks may be configured in such a way (e.g., smaller, with tighter and less contoured bends) that racing luck is more important at those venues, thereby resulting in a greater number of unexpected outcomes and higher average betting payoffs. To the extent that such a phenomenon is more likely to

arise on tracks used by small racing clubs that primarily run low-stakes races, the estimated coefficient on race stake will be upwardly biased. To alleviate this problem, I include a dummy variable for each of the 26 tracks on which at least 150 races were held over the sample period.

3.1 Full-sample results

I initially estimate (1) using the full sample of 30426 races and present the results in Tables 3 and 4. Table 3 summarises the results of regressions that use one of the three betting payoff measures as the dependent variable, while Table 4 focuses on the favourite-success indicator variables.

Table 3 shows that there is a strong negative relation between race stakes and bettor payoffs. On average, a doubling of the race stake is associated with an 8.2% decrease in the win payoff, a 9.8% decrease in the quinella payoff, and a 13.8% decrease in the trifecta payoff. Allowing for fixed-track effects reduces the magnitude of these relationships, but the change is minor (the fixed-effect elasticities are approximately 90% of those described above). And regardless of model specification, the estimated relationship is statistically significant at the 0.1% level.

One way of assessing the economic significance of these results is to estimate the predicted dollar impact on bettor payoffs of going from a low-stakes to a high-stakes race, holding all else equal. Specifically, I set the control variables equal to their respective sample means and then allow the race stake to go from its 20th percentile value to its 80th (approximately a 90% rise in stake). As shown in the last row of Table 3, this lowers the predicted win payoff by \$0.35, the quinella payoff by \$1.45, and the trifecta payoff by \$29.05 (approximately 6%, 7% and 10% of their respective sample medians). Allowing for fixed-track effects results in slightly smaller impacts.

[Insert Table 3 about here]

As expected, the control variables have a strongly positive impact on bettor payoffs, with all coefficients being significant at the 0.1% level. Reflecting the idiosyncratic nature of bets that rely on the success of a single horse, the explanatory power of the model is greater for the more exotic bets: the R^2 for the trifecta payoff is 0.22, but this falls to 0.07 for the win payoff.

Table 4 shows that there is also a significant positive relation between the probability of favoured horse success and race stake: the win pool favourite has an approximately 5.5% greater chance of succeeding in a race in the 80th stake percentile than in a race in the 20th stake percentile; for the place pool favourite, the corresponding margin is approximately 2.5%. Again, these estimated relationships are statistically significant at the 0.1% level, with and without track fixed effects.

[Insert Table 4 about here]

3.2 Do the full-sample results indicate that high-stakes races induce more effort from trainers, or just that high-stakes races provide more information for bettors?

The results in Tables 3 and 4 are consistent with the view that financial incentives matter for experts such as horse trainers: a high race stake induces more trainers to exert full effort in preparing their horses, thereby making the information possessed by bettors more useful in handicapping the race, and hence, all else equal, making it more likely that bettor-favoured horses will succeed. However, there is an obvious alternative explanation for Tables 3 and 4: a race with a high stake may be one in which bettors have more information about the abilities of the horses involved, and hence they are better able to handicap such a race. According to this view, races with high stakes are also races in which bettors are better informed, so any relation between race stakes and the success of bettor-favoured horses may simply reflect this information effect, and have nothing to do with trainer responses to financial incentives.

Bettors may have more and better information in a high-stakes race for at least two reasons. First, such races are typically restricted to more experienced horses (since qualification requires some minimum level of prior performance) who have a longer racing history and hence offer more information about their abilities to bettors. Second, racing media coverage is often greater for higher-stakes races, particularly when these are held in major population centres.

I attempt to control for bettor information effects by re-estimating equation (1) using sub-samples that increasingly eliminate low-information races. First, I exclude all races that were restricted to, or were primarily designed for, horses that had previously won no more than one race. This reduces the sample to 8133 races. Second, I also exclude all races that were restricted to two year old horses and those that were held outside the two principal race tracks in Auckland and Christchurch. This results in a total of 3368 races. In this way, I restrict the sample to just those races involving experienced horses at meetings receiving the most extensive media coverage. Variation in the amount of information available to bettors should be minor in these races.

Tables 5 and 6 contain the results from estimating (1) using these sub-samples. If anything, the relationship between race stake and realised bettor payoffs (Table 5) is even stronger than in the full sample: for example, a doubling of the race stake is now associated with a 19% decrease in the trifecta payoff, compared with 14% in the full sample. The relation between race stakes and bettor-favourite success (Table 6) is similarly affected: the estimated coefficients are all 50-100% greater

than in the full-sample results of Table 4. And although the standard errors are also bigger, all coefficient estimates for the stake variable remain significant at the 0.1% level. Thus, it seems unlikely that my earlier results are due to unobserved differences in the information available to bettors.

[Insert Tables 5 and 6 about here]

4. Concluding Remarks

On the one hand, experts can be expected to advance their own financial interests by, for example, favouring their most profitable clients. On the other hand, experts are more likely than other workers to be subject to professional pride considerations that motivate them to treat all clients equally. However, very little is known about the way in which experts trade off these competing incentives. In this paper, I show that racehorse trainers appear to favour clients whose horses compete in high-stakes races: such races are characterised by lower average returns to successful bettors than low-stakes races, consistent with trainers uniformly exerting full effort in the former but not the latter.

Such behaviour is not necessarily contrary to the wishes of clients: some owners of horses that race for low stakes may tolerate a lower level of service today in the expectation that they will be favoured when and if they come to own a horse that races for high stakes. Nevertheless, my findings support earlier results from real estate agents (Levitt and Syverson, 2005; Rutherford et al, 2005) and auto mechanics (Hubbard, 1998) which suggest that the financial incentives facing experts dominate the intrinsic incentives, contrary to the views of some psychologists (e.g., Kohn, 1993a, 1993b). It would of course be interesting to discover whether such dominance extends to more "upscale" expert professions such as lawyers, physicians/surgeons, dentists and accountants.

Finally, the results of this paper may also have implications for the literature on the efficiency of horse race betting markets. In a study of United Kingdom thoroughbred racing, Vaughan Williams and Paton (1997) find little evidence of betting inefficiency in high-quality races - which presumably offer relatively high stakes - but stronger evidence in other races. Those authors suggest that this difference is due to there being greater media coverage of high-quality races, but my results imply another possibility - that high-quality races induce full effort from all trainers and hence eliminate unobservable variation in trainer effort as a source of bettor error.

References

- Bonner, S., R. Hastie, G. Sprinkle and S. Young, 2000. A review of the effects of financial incentives on performance in laboratory tasks: implications for management accounting. *Journal of Management Accounting Research* 12, 19-64.
- Camerer, C. and R. Hogarth, 1999. The effects of financial incentives in experiments: a review and capital-labor-production framework. *Journal of Risk and Uncertainty* 19, 7-42.
- Crafts, N., 1985. Some evidence of insider knowledge in horse race betting in Britain. *Economica* 52, 295-304.
- Coleman, L., 2007. Just how serious is insider trading? An evaluation using thoroughbred wagering markets. *Journal of Gambling Business and Economics* 1, 31-55.
- Crosson, B., C. Propper and A. Perkins, 2001. Do doctors respond to financial incentives? UK family doctors and the GP fundholder scheme. *Journal of Public Economics* 79, 375-398.
- Ehrenberg, R. and M. Boganno, 1990. Do tournaments have incentive effects? *Journal of Political Economy* 98, 1307-1324.
- Gabriel, P. and J. Marsden, 1990. An examination of market efficiency in British racetrack betting. *Journal of Political Economy* 98, 874-885.
- Gaynor, M., J. Rebitzer and L. Taylor, 2004. Physician incentives in health maintenance organizations. *Journal of Political Economy* 112, 915-931.
- Gompers, P. and J. Lerner, 1999. Conflict of interest in the issuance of public securities: evidence from venture capital. *Journal of Law and Economics* 42, 1-28.
- Gruber, J. and M. Owings, 1996. Physician financial incentives and cesarean section delivery. *Rand Journal of Economics* 27, 99-123.
- Hausch, D., V. Lo and W. Ziemba, 1994. *Efficiency of Racetrack Betting Markets*. San Diego: Academic Press.
- Hubbard, T., 1998. An empirical examination of moral hazard in the vehicle inspection market. *RAND Journal of Economics* 29, 406-426.
- Jenkins, G., N. Gupta, A. Mitra and J. Shaw, 1998. Are financial incentives related to performance? A meta-analytical review of empirical research. *Journal of Applied Psychology* 83, 777-787.
- Kohn, A., 1993a. *Punished By Rewards*. New York: Houghton Mifflin.
- _____, 1993b. Why incentive plans cannot work. *Harvard Business Review* 71, 54-63.
- Lazear, E., 2000. Performance pay and productivity. *American Economic Review* 90, 1346-1361.
- Levitt, S. and C. Syverson, 2005. Market distortions when agents are better informed: the value of information in real estate. NBER Working Paper W11053; <http://www.nber.org/papers/w11053>

- Paarsch, H. and B. Shearer, 2000. Piece rates, fixed wages and incentive effects: statistical evidence from payroll records. *International Economic Review* 41, 59-92.
- Rutherford, R., T. Springer and A. Yavas, 2005. Conflicts between principals and agents: evidence from residential brokerage. *Journal of Financial Economics* 76, 627--665.
- Sauer, R., 1998. The economics of wagering markets. *Journal of Economic Literature* 36, 202102064.
- Shearer, B., 2004. Piece rates, fixed wages and incentives: evidence from a field experiment. *Review of Economic Studies* 71, 513-534.
- Shin, H., 1992. Prices of state contingent claims with inside traders, and the favourite-longshot bias. *Economic Journal* 102, 426-435.
- _____, 1993. Measuring the incidence of insider trading in a market for state-contingent claims. *Economic Journal* 103, 1141-1153.
- Vaughan Williams, L. and D. Paton, 1997. Why is there a favorite-longshot bias in British racetrack betting markets? *Economic Journal* 107, 150-158.
- White, H., 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica* 48, 431-460.

Figure 1

Proportion of Information Available to Bettors if Trainer Effort Depends on Race Stake

In the low stakes race, some trainers exert less-than-full effort because of better opportunities elsewhere, and so a relatively large proportion of the relevant information is unavailable to bettors. As a result, some horses will be over-bet and others under-bet, and unexpected race outcomes will be relatively common. But in the high-stakes race, all trainers do their best, so the quantity of private information is relatively small. Bettor odds thus incorporate a greater proportion of the relevant information and unexpected race outcomes are less common.

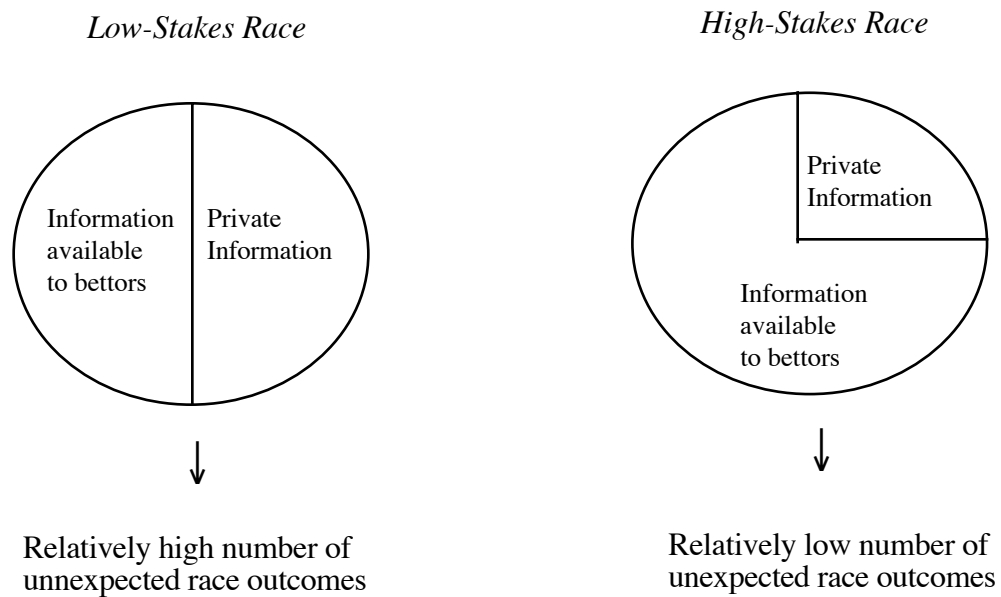


Table 1**Summary Statistics**

All data are from 30426 harness horse races run in New Zealand between 1992 and 2006. Race Stake is the total race prize money paid to the owners of successful horses. Win Payoff is the return to bettors who select the winning horse. Quinella Payoff is the return to bettors who select the first two horses home (in either order). Trifecta Payoff is the return to bettors who select the first three horses home (in finishing order). All payoffs are based on a \$1 unit of investment. Field Size is the number of horses in the race. Unpredictability Index ranges from 0 to 3 and is formed by adding 1 if the race is (i) run on grass, (ii) begun from a standing start, (iii) adversely affected by the weather. Win Favourite Success is equal to one if the favoured horse in the win betting pool wins and zero otherwise. Place Favourite Success is equal to one if the favoured horse in the place betting pool finishes in the first three and zero otherwise.

	Mean (standard deviation)	Min	Median	Max
Race Stake (\$)	8019 (15454)	1000	5000	750000
Win Payoff (\$)	8.60 (9.71)	1.00	5.45	173.70
Quinella Payoff (\$)	44.65 (72.89)	1.05	21.30	2103.35
Trifecta Payoff (\$)	1055.00 (4023.36)	2.10	307.50	398606
Field Size	11.73 (2.32)	4	12	18
Unpredictability Index	0.672 (0.71)	0	1	3
Win Favourite Success	0.335 (0.002)	-	-	-
Place Favourite Success	0.610 (0.003)	-	-	-

Table 2**Preliminary Analysis**

All data are from 30426 harness horse races run in New Zealand between 1992 and 2006. All variables are defined in Table 1. The first column of numbers calculates the mean of each variable for races offering stakes above the sample median; the second column calculates the mean of each variable for races offering stakes below the sample median; the third presents the difference between the first two columns. The fourth column gives the correlation between each variable (or its natural log in the case of the three payoff variables) and the natural log of race stake. All estimates in the last two columns are significantly different from zero at the 0.1% level.

	Races with above- median stake	Races with below- median stake	Difference	Correlation with Race Stake
Average Win Payoff (\$)	8.28	8.93	-0.65	-0.25
Average Quinella Payoff (\$)	42.13	47.15	-5.02	-0.27
Average Trifecta Payoff (\$)	951.51	1157.92	-206.41	-0.31
Average Field Size	11.45	12.00	-0.55	-0.29
Unpredictability Index Average	0.619	0.725	-0.11	-0.26
Win Favourite Success Rate	0.347	0.322	0.025	0.15
Place Favourite Success Rate	0.625	0.596	0.029	0.13

Table 3**Regressions of Betting Payoffs on Race Stakes: Full Sample**

The dependent variables are the natural log of the payoffs (based on a \$1 unit of investment) to successful bettors in three betting pools run on 30426 harness horse races run in New Zealand between 1992 and 2006. Field Size is the number of horses in the race. Unpredictability Index ranges from 0 to 3 and is formed by adding 1 if the race is (i) run on grass, (ii) begun from a standing start, (iii) adversely affected by the weather. Ln Race Stake is the natural log of total race prize money. The regressions with fixed-track effects include 26 dummy variables, one for each of the race tracks at which at least 150 races were held during the sample period. Heteroskedasticity-adjusted standard errors are in parentheses; all coefficient estimates are significantly different from zero at the 0.1% level. The final row calculates the predicted impact on betting payoffs of going from the 20th stake percentile to the 80th.

	<i>Dependent Variable</i>					
	Win Payoff	Win Payoff	Quinella Payoff	Quinella Payoff	Trifecta Payoff	Trifecta Payoff
Intercept	1.446 (0.09)	1.453 (0.09)	1.869 (0.10)	1.806 (0.11)	3.581 (0.13)	3.513 (0.14)
Field Size	0.084 (0.00)	0.083 (0.00)	0.174 (0.00)	0.176 (0.00)	0.283 (0.00)	0.286 (0.00)
Unpredictability Index	0.064 (0.01)	0.057 (0.01)	0.091 (0.01)	0.098 (0.01)	0.122 (0.01)	0.127 (0.01)
Ln Race Stake	-0.082 (0.01)	-0.075 (0.01)	-0.098 (0.01)	-0.089 (0.01)	-0.138 (0.01)	-0.124 (0.02)
R ²	0.07	0.07	0.15	0.15	0.22	0.22
Fixed-Track Effects	No	Yes	No	Yes	No	Yes
Predicted impact of going from the 20th stake percentile to the 80th	-\$0.35	-\$0.25	-\$1.45	-\$1.35	-\$29.05	-\$26.40

Table 4**Logit Regressions of Betting-Favourite Success Rates on Race Stakes: Full Sample**

All data are from 30426 harness horse races run in New Zealand between 1992 and 2006. The dependent variable is either (i) equal to one if the favoured horse in the win betting pool wins and zero otherwise (Win Favourite Success) or (ii) equal to one if the favoured horse in the place betting pool finishes in the first three and zero otherwise (Place Favourite Success). All explanatory variables are described in Table 3. All models are logit regressions. The regressions with fixed-track effects include 26 dummy variables, one for each of the race tracks at which at least 150 races were held during the sample period. Standard errors are in parentheses; all coefficient estimates are significantly different from zero at the 0.1% level. The final row calculates the predicted effect of going from the 20th stake percentile to the 80th percentile on the probability of the favourite succeeding.

	<i>Dependent Variable</i>			
	Win Favourite Success	Win Favourite Success	Place Favourite Success	Place Favourite Success
Intercept	-0.809 (0.20)	-0.907 (0.22)	1.036 (0.20)	1.098 (0.22)
Field Size	-0.080 (0.01)	-0.077 (0.01)	-0.116 (0.01)	-0.117 (0.01)
Unpredictability Index	-0.133 (0.02)	-0.100 (0.02)	-0.137 (0.02)	-0.140 (0.02)
Ln Race Stake	0.131 (0.02)	0.120 (0.02)	0.101 (0.02)	0.091 (0.02)
Pseudo-R ²	0.01	0.01	0.02	0.02
Fixed-Track Effects	No	Yes	No	Yes
Predicted impact of going from the 20th stake percentile to the 80th	5.80%	5.31%	2.55%	2.30%

Table 5**Regressions of Betting Payoffs on Race Stakes: Sub-samples**

This table estimates the relationship between betting payoffs and race stakes for two sub-samples that are increasingly restricted to races involving experienced horses for whom considerable information is available to the betting public. Sub-sample A excludes all races restricted to horses that have previously won no more than one race; sub-sample B also excludes all races for two year-old horses and all races conducted outside the two major metropolitan tracks located in Auckland and Christchurch. All variables are defined in Table 3. Heteroskedasticity-adjusted standard errors are in parentheses; all coefficient estimates except those marked with a '^' are significantly different from zero at the 0.1% level. The final row calculates the predicted impact on betting payoffs of going from the 20th stake percentile to the 80th.

	<u>Sub-sample A (n = 8132)</u>			<u>Sub-sample B (n = 3368)</u>		
	<i>Dependent Variable</i>					
	Win Payoff	Quinella Payoff	Trifecta Payoff	Win Payoff	Quinella Payoff	Trifecta Payoff
Intercept	1.575 (0.14)	2.144 (0.18)	3.818 (0.23)	1.722 (0.21)	2.022 (0.27)	3.938 (0.32)
Field Size	0.095 (0.00)	0.188 (0.00)	0.311 (0.01)	0.091 (0.01)	0.180 (0.01)	0.300 (0.01)
Unpredictability Index	0.079 (0.01)	0.113 (0.02)	0.139 (0.02)	0.068 [^] (0.02)	0.151 (0.03)	0.162 (0.04)
Ln Race Stake	-0.107 (0.02)	-0.141 (0.02)	-0.191 (0.02)	-0.118 (0.02)	-0.119 (0.03)	-0.190 (0.03)
R ²	0.09	0.18	0.28	0.07	0.15	0.24
Predicted impact of going from the 20th stake percentile to the 80th	-\$0.45	-\$2.25	-\$40.85	-\$0.55	-\$2.15	-\$43.90

Table 6**Logit Regressions of Betting-Favourite Success Rates on Race Stakes: Sub-samples**

This table estimates the relationship between favourite-horse success rates and race stakes for two sub-samples that are increasingly restricted to races involving experienced horses for whom considerable information is available to the betting public. Sub-sample A excludes all races restricted to horses that have previously won no more than one race; sub-sample B also excludes all races for two year-old horses and all races conducted outside the two major metropolitan tracks located in Auckland and Christchurch. All variables are defined in Table 4. All models are logit regressions. Standard errors are in parentheses; all coefficient estimates except those marked with a '^' are significantly different from zero at the 0.1% level. The final row calculates the predicted effect of going from the 20th stake percentile to the 80th percentile on the probability of the favourite succeeding.

	<u>Sub-sample A (n = 8132)</u>		<u>Sub-sample B (n = 3368)</u>	
	<i>Dependent Variable</i>			
	Win Favourite Success	Win Favourite Success	Place Favourite Success	Place Favourite Success
Intercept	-1.089 (0.34)	0.391 [^] (0.36)	-1.232 [^] (0.50)	-0.147 [^] (0.52)
Field Size	-0.100 (0.01)	-0.129 (0.01)	-0.078 (0.02)	-0.106 (0.02)
Unpredictability Index	-0.149 (0.04)	-0.175 (0.04)	-0.088 [^] (0.06)	-0.237 (0.06)
Ln Race Stake	0.183 (0.04)	0.185 (0.04)	0.169 (0.05)	0.217 (0.06)
Pseudo-R ²	0.02	0.02	0.01	0.02
Predicted impact of going from the 20th stake percentile to the 80th	9.82%	5.55%	10.56%	7.46%