

Examining Agency Conflict in Horse Racing

Alasdair Brown*

We study UK horse racing for signs of conflict between horse owners (principals) and trainers (agents). Trainers often prepare their own horses for races in addition to having outsiders' horses in their care. Utilizing betting market data to infer the expected performance of a horse, we find that owner–trainer horses outperform outsider–trainer horses, indicating that this principal–agent relationship is characterized by agent shirking. If the owner holds a large proportion of the horses in the trainer's stable, the shirking effect may be mitigated but not eradicated. In a separate result, we find that outsider–trainer horses are more inconsistent than their owner–trainer peers. As inconsistency is a sign of betting market manipulation, this suggests that the agent in this setting extracts a second, informational rent from the principal.

JEL Classification: D82, J32, M52, G14

1. Introduction

We employ agents to undertake a number of frequent tasks on our behalf. Agents are entrusted to find us a job, manage our investments, and school our children. As the effort and performance of agents is imperfectly observed, this creates ample opportunities for conflict and often welfare losses for the principal. For example, Rutherford, Springer, and Yavas (2005) and Levitt and Syverson (2008) find that estate agents sell their client's houses for less than their own houses after controlling for observable house characteristics. Similarly, Ang, Cole, and Wuh Lin (2000) find that managers who have little stake in a small firm spend more on expenses and run the firm less efficiently than firms with a single owner-manager.

In this article, we examine agency conflict in horse racing. Racehorse trainers are responsible for the welfare of horses in their stable, supervising how the animal is fed, how hard the horse is run prior to a race, and often deciding in which races the horse will run. Many trainers (agents) divide their time between the preparation of their own horses and those of outsiders (principals). While trainers receive all of the win purse if their horse is successful, they receive only a fraction (typically 10%)¹ for training an outsider's horse to victory. In addition—and crucially for any study of agency conflict—much of the trainer's effort is unobservable; moreover, the asymmetric information problem is exacerbated by there being an animal at the center of the conflict. To illustrate the opacity of the trainer's actions, Fox (2005) likens the

* School of Economics, University of East Anglia, Norwich NR4 7TJ, United Kingdom; E-mail alasdair.brown@uea.ac.uk.

I would like to thank the Editor, Laura Razzolini, and two anonymous referees for their comments on the article. This work also benefited from suggestions from Fabrizio Adriani, Pasquale Scaramozzino, Sasha Talavera, and audiences at SOAS and UEA. All remaining errors are my own.

Received July 2011; accepted November 2011.

¹ Boyle, Guthrie, and Gorton (2010) indicate that 10% is the typical commission the trainer can earn, but Scott (1968) suggests that the commission can be as much as 50% if the owner wishes to align the trainer's incentives with his or her own.

racehorse trainer to a tribal shaman or witch doctor, whose success is ascribed to impressive work and whose failure is due to factors outside of their control.

There are a number of elements to horse racing that make it an interesting arena for the study of agency conflict. First, in contrast to the infrequent sale of a house, horses run on a regular basis over the course of their careers, yielding a large sample of performance data. We use data on every horse race in the United Kingdom between 2005 and 2010. Second, the interaction of horse trainer and owner is dynamic, as the owner can withdraw a horse from a trainer's stable if performance is deemed unsatisfactory; the estate agent is unlikely to fret to the same degree over repeat business from a house seller. Third, there is no need to infer the intrinsic ability of a horse from its observable characteristics; there is a betting market from which to gauge the prerace expected performance of a horse. Finally, the horse owner is relatively empowered compared to the owner of a small firm; whereas removing a firm manager can be problematic and time consuming, the horse can often be removed from a trainer's stable on short notice and without compensation.

Utilizing betting market data to infer the expected performance of a horse, we find that owner–trainer horses outperform when compared with outsider–trainer horses. This suggests that this particular agency conflict is characterized by the agent shirking. Alternatively, it is plausible that trainers—who undoubtedly possess private information on the racing game—are better able to identify the good horses unappreciated by the betting public. To control for this, we examine the performance of older, higher-class horses for which the private information set should be smaller. As in the broader study, owner–trainer horses outperform outsider–trainer horses in this subset of animals.

One way in which an owner can mitigate the shirking effect is to entrust the preparation of a number of horses to the same trainer. This way, the trainer should be careful not to disappoint the owner for fear of losing a substantial proportion of their stable and livelihood. We calculate the dominance of a trainer's stable by each owner and find that owning a large proportion of the horses under a trainer's care can sometimes mitigate but not eradicate the shirking effect.

In horse racing, there is potentially a second source of income besides the prize money. If trainers or owners have private information on the upcoming performance of a horse, they can bet for profit; indeed, in the United Kingdom, it is legal for them to do so. Trainers have the potential to create such private information, as they control the preparation of horses in their care. One sign of manipulation of horse performance for betting market gain is inconsistency—initiating a poor performance by the horse one week and thereby causing the betting public to downgrade their estimate of the horse's ability only for the horse to outperform relative to the betting market's downgraded expectations the next week when in fact the horse was simply performing at its true level in the second race.

We examine whether such performance manipulation is a second, informational source of agency conflict. After controlling for other factors that affect the consistency of performance, we find that horses trained by an outsider are more inconsistent than horses trained by their owner. This suggests that trainers not only exert less effort when training an outsider's horse but also, on average, exploit these horses for betting market gain. These results give empirical weight to anecdotal evidence that horse racing is at least partially manipulated and that the trainer, with privileged access to information, is front and center in this manipulation.²

² Scott (1968), in an anthropological study of a U.S. racetrack, details how widespread performance manipulation is, and the central role that trainers play in this due to their informational advantage over other actors.

The evidence of performance manipulation in horse racing also corresponds with evidence of betting market-induced corruption in other sports. For example, Wolfers (2006) finds evidence of point shaving—where a heavily favored team wins but fails to beat the bookmaker’s spread—in college and professional basketball. Borghesi (2008) and Bernhardt and Heston (2010) suggest that such indirect studies of corruption, which rely on subtle inferences, should test an alternative hypothesis to ensure the robustness of the result.

In our case, an alternative explanation for the greater consistency of owner–trainer horses is that the trainer—with superior knowledge of racing—is able to select the most consistent horses. If the trainer does have private information, it is much more likely to be related to younger, lower-profile horses. This explanation can be discounted, however, because the inconsistency of outsider–trainer horses is apparent even after controlling for the class in which a horse runs and its age.

An article by Boyle, Guthrie, and Gorton (2010) examined horse racing results data from New Zealand to search for conflict between horse owners and trainers. On average, they find that the performance of horses is unaffected by the owner–trainer relationship, though horses trained in small stables do show some signs of agency-induced underperformance. Their study differs from ours in that betting market data are not used to infer a horse’s expected performance, nor do the authors consider the possibility that the trainer’s manipulation of the performance of outsiders’ horses—and the resultant opportunity for profitable betting—may be a second source of agency conflict.

The rest of this article is structured as follows. Section 2 describes the data set and how the expected performance of a horse is inferred. Section 3 contains the empirical analysis of horse performance and consistency, and section 4 concludes.

2. Data

We collected data on every horse race run in the United Kingdom between January 1, 2005, and December 31, 2010, inclusive, from Betwise, a betting information company. After discarding data on races without a designated class—which is an important control variable for our later analysis—we were left with 559,383 horse-level performance observations. The race data include the date of the meeting, the type of race (e.g., flat, national hunt), the distance over which the race is run, the handicap system, the class of the race (which ranges from 1 [top] to 7 [bottom]), and the winning prize money (which ranges from GBP 0 to GBP 1,000,000). Summary statistics on horse and race characteristics can be found in Table 1.

For each horse, we have their age and the identity of their owner and trainer at the time of the race. There are 55,763 unique horses in our sample, with 27,861 unique owners and 2944 unique trainers. To determine whether the owner was also the trainer of the horse, we took the surname of the trainer, then ran a regular expression to test whether any sequential part of the owner name was an exact match on the trainer surname. If it was, our *owner=trainer* dummy variable equals 1 and 0 otherwise. This algorithm captures, at least in part, the frequent occurrence of a trainer registering a horse one owns in one’s wife’s name.³

³ To ensure that our results are not driven by common names (e.g. Smith, Brown), we also altered the algorithm so that only owners and trainers with an exact match (i.e., A Brown) were classified as one individual. Our results were unaffected.

Table 1. Summary Statistics

	Mean	Median	Maximum	Minimum	Standard Deviation	Skewness	Kurtosis
Age	4.95	4	22	2	2.33	0.84	3.28
Starting price (odds)	22.6	13	501	1.02	29.2	3.64	24.32
Class	4.41	5	7	1	1.41	-0.71	3.09
Distance (thousands of yards)	2.72	2.2	7.92	1.1	1.4	0.66	2.27
No. of horses in race	11.53	11	40	2	3.74	0.89	6.16
Winning prize money (thousands of GBP)	12.65	5.24	1000	0	38	13.83	275.4

Summary statistics for horse and race characteristics for all UK races between January 1, 2005, and December 31, 2010.

We are also interested in the extent to which an owner dominates a trainer's stable. This is calculated by counting the number of horses in a trainer's stable in any year and defining our *owner dominance* variable as the proportion (between 0 and 1) of these horses that the owner in question holds. For example, suppose a trainer had 10 horses, of which two were his own and eight belonged to one other individual. For the trainer's two horses, we would have the *owner=trainer* dummy equal to 1 and the *owner dominance* variable equal to 0.2. For the remaining eight horses, the *owner=trainer* dummy would be 0, with the *owner dominance* variable equal to 0.8. The expectation is that the outsider-owner can offset part of the agency costs by holding a large proportion of the remaining horses under the same trainer's supervision.

In order to infer the expected performance of each horse, we collected the Starting Price, which is a summary measure of the odds quoted by British bookmakers at the time the race began. From here, we can infer a predicted finishing position in the race for each horse. For example, a horse with odds of two to one (GBP 2 returned for every GBP 1 staked if the horse wins) can be expected to finish above a horse with odds of four to one.⁴

We then compared the predicted finishing position of each horse with its actual finishing position. All horses that failed to finish a race were classified as finishing last. We then used the actual and predicted finishing positions and number of horses in the race to construct a measure of horse performance:

$$\text{Relative Performance} = \frac{\text{Predicted Finish} - \text{Actual Finish}}{\text{Number of Horses}}.$$

Relative performance has a mean of -0.0266, a median of 0, a minimum of -0.97, and a maximum of 0.96. The standard deviation is 0.3084, with skewness of -0.145 and kurtosis of 0.1980. The measure is negatively skewed because nonfinishing horses are classified as having finished last. If there is more than one horse that fails to finish, the average *relative performance* is less than 0. If no more than one horse fails to finish, the average for that race is 0.

In order to infer the consistency of the horse, we are also interested in the absolute value of *relative performance*:

⁴ There is the possibility of errors in this estimation of predicted finishing position. For example, consider a temperamental horse that will either win or fail to finish. This horse may be expected to finish behind a slower, more consistent horse but the odds of winning will not capture this. Nevertheless, in the absence of place (top 2/3 finish) odds, the odds of winning are a good proxy for the predicted finishing position.

$$\text{Absolute Relative Performance} = \left| \frac{\text{Predicted Finish} - \text{Actual Finish}}{\text{Number of Horses}} \right|.$$

This measure treats a horse that outperforms exactly the same as a horse that underperforms. Suppose a horse runs in an eight-horse race. The horse is tipped to finish first but instead finishes fifth. This leads to an *absolute relative performance* observation of 0.5. If another horse is predicted to finish fifth but actually finishes first, the *absolute relative performance* observation would also be equal to 0.5.

3. Empirical Analysis

In the first part of this section, we are interested in the effect that the separation of ownership and management (training) in horse racing has on performance. Further, we examine the extent to which an outsider-owner can mitigate the agency problem by placing a number of horses with the same trainer, thereby increasing the trainer's dependence on the owner.

In Table 2, we display the results of a regression of our *relative performance* measure on the *owner=trainer* dummy and the *owner dominance* variable for the full six-year sample. The intercept term is negative because of the negative skew in *relative performance*. More important, we observe that owner-trainer horses outperform outsider-trainer horses at the 0.1% significance level. As the average horse runs in a race of 11.53 horses (see Table 1), the coefficient in regression 1 in Table 1 implies that owner-trainer horses, on average, outperform outsider-trainer horses by 0.35 places. This agency-induced problem can be mitigated by an owner's dominance of a trainer's stable, as this coefficient is positive and significant at the 0.1% level. However, examining the coefficients of the model, we can see that an owner who completely dominated an outsider-trainer's stable could not expect to completely avoid agency costs in the form of underperformance. Indeed, the coefficient associated with *owner dominance* is positive and significant in only 7 out of 13 specifications used in this article, so its effect is not altogether consistent.

We ran supplementary regressions using fixed effects, also displayed in Table 2, estimated using the White method to allow for clustering of errors at the handicap-type, horse, trainer, owner, and race-type level. The fixed effects for handicap type, horse, and race type have little effect on the remaining coefficients. On the other hand, the trainer fixed effect, captured in regression 4, serves to reduce but not eliminate the coefficient associated with the *owner=trainer* dummy. This suggests that most of the agent shirking is concentrated with a number of "bad-egg" trainers. Similarly, the coefficient is reduced when owner fixed effects are incorporated, suggesting that a number of "naive" owners are being disproportionately hit by agency costs. In race terms, owner-trainer horses outperform outsider-trainer horses by only 0.07 places once trainer fixed effects are incorporated and only 0.15 places when owner fixed effects are incorporated.

It is also instructive to consider the interaction between the *owner=trainer* dummy and the *owner dominance* variable. Trainers may be expected to exert minimal effort for their own horses if they make up only a small proportion of their livelihood (stable). Their own horses could, in effect, be their hobby and run alongside their professional training of outsider horses. We use regression 7 in Table 2 to capture this effect. As expected, the strongest outperformance of *owner=trainer* horses comes when the trainer owns a greater proportion of the horses they train.

Table 2. Performance Analysis

Dependent Variable: Relative Performance	1	2	3	4	5	6	7
Intercept	-0.032284 (0.000478)****	-0.033817 (0.002876)****	-0.032298 (0.027893)	-0.02598 (0.002201)	-0.029335 (0.014377)**	-0.033429 (0.000546)****	-0.030279 (0.000507)****
Owner = trainer dummy	0.030231 (0.001432)****	0.034705 (0.001682)****	0.035163 (0.011492)****	0.006956 (0.001642)****	0.012966 (0.006780)*	0.03283 (0.001434)****	0.016565 (0.001822)****
Owner Dominance	0.016303 (0.002120)****	0.024165 (0.002137)****	0.011872 (0.017179)	-0.011473 (0.004134)***	0.009124 (0.010253)	0.022859 (0.002125)****	-0.005841 (0.002835)**
Owner = trainer dummy * Owner Dominance							0.050752 (0.004261)****
Includes handicap fixed effects	No	Yes	No	No	No	No	No
Includes horse fixed effects	No	No	Yes	No	No	No	No
Includes trainer fixed effects	No	No	No	Yes	No	No	No
Includes owner fixed effects	No	No	No	No	Yes	No	No
Includes race-type fixed effects	No	No	No	No	No	Yes	No
R ²	0.001474	0.00486	0.115379	0.033707	0.074864	0.003844	0.001727
No. of observations	559,383	559,383	559,383	559,383	559,383	559,383	559,383

Coefficient estimates when relative performance was regressed on an owner-trainer dummy variable (equaling one when the owner is also the trainer) and owner dominance (between 0 and 1, equaling 1 when the horse owner holds all horses in a trainer's stable). Regression 1 is pooled, while regressions 2-6 incorporate handicap-type, horse, trainer, owner, and race-type fixed effects, respectively. Regression 7 incorporates the interaction between the owner-trainer dummy and owner dominance. White heteroskedasticity-consistent standard errors, including allowances for clustering of errors in fixed effects, are in parentheses. ****, ***, **, and * indicate significance at the 0.1%, 1%, 5%, and 10% levels, respectively, in this table and Tables 3 and 4.

One plausible explanation for the significance of the *owner=trainer* dummy is that trainers have private information on horses and therefore buy horses that will, on average, outperform relative to betting market expectations. To control for this, we rerun the first regression, this time sampling only horses running in races of class 1–4 and of age five or over. Older, higher-grade horses are more likely to be well known by the betting public, and therefore the amount of private information on which the trainer can capitalize is likely to be less. This choice of class and age removes around half of the sample in each case. As shown in Table 3, the *owner=trainer* dummy is still a positive and significant determinant of horse performance for these two subsamples at the 0.1% level.

Another concern may be that the focus on horse-level performances is dominated by horses that run often and stables that contain a large number of horses. To a certain extent, this is controlled for by the horse and trainer fixed effects, but in regression 4 in Table 3, we consider only horses that are part of a stable of 149 or fewer horses, and in regression 5 we consider only horses that ran in 18 or fewer races during our sample period. Again, both subsamples make up approximately half of the full sample size. The *owner–trainer* effect is still present in these two subsamples, highlighting the robustness of this particular result.

Our results rely on the bookmakers—and implicitly the betting public—not fully incorporating information on horse ownership (and its effect on performance) into betting prices. It is possible, however, that a degree of learning about the implications of horse ownership takes place over the course of our sample. To investigate this, we restrict our sample to those horses that ran more than 18 times over the period, the idea being that, for these horses, there is ample opportunity (at least by the end of the sample) for the betting market to digest information on the ownership of this horse. The results are displayed in regression 6 in Table 3. There is some evidence of learning, as the coefficient associated with the *owner=trainer* dummy is reduced but the significance of the effect remains.

Our second search for agency conflict focuses on the consistency of horses. A horse can be inconsistent because of age, the class of the race, the distance over which they run, the prize money on offer, and the number of horses in a race. In particular, we expect that horses running in higher-class races, with greater prize money, should be more consistent. Boyle (2008) and Johnson, Bruce, and Yu (2010) indeed find that high-class races finish closer to betting market expectations than low-grade races.⁵

In an opaque environment where the horse's true ability is difficult to observe, inconsistency can also be a sign of betting market manipulation. A poor performance one week can lead the betting market to downgrade their estimate of the horse's ability, driving up the odds and the potential returns to betting the next race the horse runs. We assess whether, after controlling for observable determinants of inconsistency, outsider–trainer horses are more inconsistent than owner–trainer horses. If so, this suggests the existence of a second, informational form of agency conflict.

We regress the *absolute relative performance* measure on the *owner=trainer* dummy and the *owner dominance* variable, with results displayed in Table 4. We find that outsider–trainer horses are more inconsistent, as *absolute relative performance* is lower when the owner is also

⁵ This could be due to the greater inherent consistency of better horses or due to the greater monetary incentives for maximum effort at the higher-grade races. Alternatively, the explanation could be betting market related in that higher-grade races attract more betting volume and information, and therefore the finishing order is more accurately predicted.

Table 3. Performance Analysis (Robustness)

Dependent Variable: <i>Relative Performance</i>	1		2		3		4		5		6	
	All		Class ≤ 4		Age ≥ 5		Stable Size ≤ 149		No. of Runs ≤ 18		No. of Runs > 18	
Intercept	-0.032284 (0.000478)****	-0.034883 (0.000719)****	-0.040987 (0.000714)****	-0.003882 (0.000691)****	-0.045664 (0.000651)****	-0.018219 (0.000703)****						
<i>Owner = trainer dummy</i>	0.030231 (0.001432)****	0.045714 (0.002399)****	0.027539 (0.001944)****	0.009799 (0.001763)****	0.038783 (0.001951)****	0.016937 (0.002109)****						
<i>Owner Dominance</i>	0.016303 (0.002120)****	0.001286 (0.003142)	0.035887 (0.002829)****	0.018473 (0.002564)****	0.000528 (0.002735)	0.042066 (0.003323)****						
Includes handicap fixed effects	No	No	No	No	No	No						
Includes horse fixed effects	No	No	No	No	No	No						
Includes trainer fixed effects	No	No	No	No	No	No						
Includes owner fixed effects	No	No	No	No	No	No						
Includes race-type fixed effects	No	No	No	No	No	No						
<i>R</i> ²	0.001474	0.001833	0.002749	0.000633	0.001662	0.001622						
No. of observations	559,383	261,192	277,838	284,561	290,550	268,833						

Coefficient estimates when relative performance was regressed on an owner-trainer dummy variable (equaling one when the owner is also the trainer) and owner dominance (between 0 and 1, equaling 1 when the horse owner holds all horses in a trainer's stable). Regression 1 is pooled; regression 2 considers only horses when run in races of class 4, 3, 2 or 1; regression 3 considers only horses age five and over; regression 4 considers only horses in a trainer's stable of less than or equal to 149 horses; regression 5 considers only horses that ran 18 or fewer times over the sample period; and regression 6 considers only horses that ran more than 18 times over the sample period. White heteroskedasticity-consistent standard errors are in parentheses.

Table 4. Inconsistency Analysis

Dependent Variable: Absolute Relative Performance	1	2	3	4	5	6	7
Intercept	0.235279 (0.000313)****	0.201061 (0.001547)****	0.231978 (0.005208)****	0.192526 (0.043023)****	0.183706 (0.003499)****	0.184213 (0.015061)****	0.195775 (0.001898)****
Owner = trainer dummy	-0.003716 (0.000932)****	-0.005640 (0.000942)****	-0.003405 (0.001014)****	-0.006724 (0.002882)**	-0.002702 (0.001074)**	-0.007342 (0.003758)*	-0.00402 (0.000939)****
Owner Dominance	0.001981 (0.001398)	-0.004603 (0.001402)***	-0.003833 (0.001444)**	-0.003120 (0.002806)	0.004728 (0.002211)**	0.00323 (0.002949)	-0.004954 (0.001401)****
No. of horses	0.002389 (7.56E-05)****	0.002389 (7.56E-05)****	0.002224 (9.14E-05)****	0.002361 (0.000140)****	0.00244 (7.73E-05)****	0.002484 (8.66E-05)****	0.003052 (7.81E-05)****
Distance (thousands of yards)	0.002902 (0.000254)****	0.002902 (0.000254)****	1.22E-03 (0.000480)**	0.008347 (0.002356)****	0.004543 (0.000390)****	0.005312 (0.000585)****	0.004376 (0.000449)****
Class	-0.004147 (0.000222)****	-0.004147 (0.000222)****	-0.00339 (0.000402)****	-0.004734 (0.000477)****	-0.00278 (0.000269)****	-0.003148 (0.000416)****	-0.00376 (0.000234)****
Winning prize money (thousands of GBP)	-0.000025 (8.39E-06)***	-0.000025 (8.39E-06)***	-1.66E-05 (1.47E-05)	0.000033 (9.30E-06)****	-2.19E-05 (8.93E-06)***	-1.62E-05 (8.82E-06)*	-5.56E-05 (8.42E-06)****
Age	0.003729 (0.000144)****	0.003729 (0.000144)****	-0.00198 (0.000742)***	0.002886 (0.007210)	0.004677 (0.000376)****	0.004514 (0.001940)**	0.002145 (0.000166)****
Includes handicap fixed effects	No	No	Yes	No	No	No	No
Includes horse fixed effects	No	No	No	Yes	No	No	No
Includes trainer fixed effects	No	No	No	No	Yes	No	No
Includes owner fixed effects	No	No	No	No	No	Yes	No
Includes race-type fixed effects	No	No	No	No	No	No	Yes
R ²	0.000028	0.006617	0.024986	0.116652	0.015558	0.064459	0.009586
No. of observations	559,383	559,383	559,383	559,383	559,383	559,383	559,383

Coefficient estimates when absolute relative performance (a measure of horse inconsistency) was regressed on an owner-trainer dummy variable (equaling one when the owner is also the trainer) and owner dominance (between 0 and 1, equaling 1 when the horse owner holds all horses in a trainer's stable). Regression 1 is pooled, regression 2 is pooled and includes control variables, and regressions 3-7 incorporate handicap-type, horse, trainer, owner, and race-type fixed effects, respectively. White heteroskedasticity-consistent standard errors, including allowances for clustering of errors in fixed effects, are in parentheses.

the trainer. This result is robust to the introduction of control variables for consistency and handicap-type, horse, trainer, and owner fixed effects. The introduction of trainer fixed effects reduces the coefficient associated with the owner–trainer dummy, again suggesting that any manipulation of horse performance is concentrated with a number of bad-egg trainers. Further, dominance of a trainer’s stable does not appear to offset this agency cost.

A natural next step would be to develop a betting strategy that capitalizes on the results outlined in this article. This has not proved feasible for two reasons. First, the over-round (the extent to which the sum of implied win probabilities quoted by a bookmaker exceeds 1) renders most strategies unprofitable, even if they increase returns above the blanket strategy of betting on all horses. A second reason—more closely associated with the betting market manipulation outlined in this article—is that any trainer manipulating the performance of horses would, by necessity, attempt to disguise their actions. A horse may be deliberately underperformed for a number of races until the odds on a win have lengthened sufficiently for the trainer to push it for a win. Alternatively, the advent of betting exchanges—where bettors can wager on a horse to lose—could mean that the payoff is actually reaped when the horse underperforms. Therefore, while (unnecessary) inconsistency indicates betting market manipulation, it does not lead us to a profitable strategy.

There are two other issues that warrant discussion. The first is the role of the jockey, specifically, whether the jockey would be in a position to manipulate horse races without the consent of the trainer. The jockey would be responsible for implementing the parts of performance manipulation that the trainer does not control (i.e., in the race itself), but it is unlikely that the jockey can manipulate horse performance without the trainer’s knowledge. A jockey does not have a significant informational advantage over the trainer, who, after all, is in charge of all elements of the preparation of the horse for races.

A second issue is that, under certain circumstances, it may also be in the interests of trainers to manipulate the performance of their own horses. For this to occur, however, we would require the prize money to be very small compared to the potential betting income. As trainers receive only a percentage of any prize money, it must be optimal for the trainer to manipulate horse performance if it is optimal for the owner to do so, but it is not necessarily optimal for the owner to engage in manipulation if it optimal for the trainer to do so. Therefore, a trainer with both one’s own horses and those of outsiders in one’s care will first find it optimal to manipulate the performance of the outsiders’ horses.

4. Conclusion

We examine horse racing for signs of conflict between the principal—the horse owner—and the agent—the trainer entrusted with preparing the horse for racing. We find that owner–trainer horses outperform outsider–trainer horses, though this effect can, in some cases, be mitigated by the owner placing a number of horses with the same trainer. We also find a second agency cost where outsider–trainer horses are more inconsistent than owner–trainer horses, suggesting that trainers, on average, use outsiders’ horses for betting market gain.

In race terms, the separation of ownership and management leads to, on average, underperformance of 0.35 places for a horse trained by an outsider. It is perhaps surprising that the betting market does not incorporate information on the ownership of horses and its

implications for performance. It is because of this oversight, however, that we are able to identify the costs that owners bear when delegating management of their asset to an outsider.

References

- Ang, J. S., R. A. Cole, and J. Wuh Lin. 2000. Agency costs and ownership structure. *Journal of Finance* 55:81–106.
- Bernhardt, D., and S. Heston. 2010. Point shaving in college basketball: A cautionary tale for forensic economics. *Economic Inquiry* 48:14–25.
- Borghesi, R. 2008. Widespread corruption in sports gambling: Fact or fiction? *Southern Economic Journal* 74:1063–9.
- Boyle, G. 2008. Do financial incentives affect the quality of expert performance? Evidence from the racetrack. *Journal of Gambling Business and Economics* 2:43–60.
- Boyle, G., G. Guthrie, and L. Gorton. 2010. Holding onto your horses: Conflicts of interest in asset management. *Journal of Law and Economics* 53:689–713.
- Fox, K. 2005. *Racing tribe: Portrait of a British subculture*. New Brunswick, NJ: Transaction.
- Johnson, J., A. Bruce, and J. Yu. 2010. The ordinal efficiency of betting markets: An exploded logit approach. *Applied Economics* 42:3703–9.
- Levitt, S. D., and C. Syverson. 2008. Market distortions when agents are better informed: The value of information in real estate transactions. *Review of Economics and Statistics* 90:599–611.
- Rutherford, R. C., T. M. Springer, and A. Yavas. 2005. Conflicts between principals and agents: Evidence from residential brokerage. *Journal of Financial Economics* 76:627–65.
- Scott, M. B. 1968. *The racing game*. Chicago: Aldine.
- Wolfers, J. 2006. Point shaving: Corruption in NCCA basketball. *American Economic Review* 96:279–83.