Everyone's A Winner: The Market Impact of Technologically Advantaged Agents

Richard J. McGee[a], Johnnie E. V. Johnson[a,*]

[a]*Southampton Business School, University of Southampton, Highfield Campus, Southampton, SO17 1BJ, UK*

Abstract

There is an ongoing debate as to whether the presence of technologically advantaged agents (TAAs) in financial markets warrants regulatory intervention, because their presence may lead to an inequitable wealth transfer. We use tote betting market data (2.8 million prices from over 175,000 harness races over a five year period) to provide evidence that a market with heterogeneous utility agents can include a net transfer of wealth, from non-TAAs to TAAs, with the transaction proving beneficial to both in terms of their realized utility. Our findings are consistent with TAAs receiving a premium for providing liquidity and supplying payoffs in popular states. We conclude that a transfer of wealth to TAAs may not be a sufficient criteria for regulatory intervention and suggest that a more nuanced approach that considers representative agent utilities should be applied.

Keywords: Market Making, Market Regulation, Heterogeneous Agent Utility

1. Introduction

We show that the consistent transfer of trading returns from one agent to another does not necessarily imply an inequitable or unfair market, even when the beneficiary agent is exploiting a technological advantage to make profit.

There is an ongoing debate around high frequency trading practices in financial markets, as some of these practices are seen as predatory. Agents with faster access to bid/offer prices, through faster hardware connectivity and/or co-location on an exchange, may be deemed to have an unfair advantage. Regulatory bodies in different countries have responded in a variety of ways to the challenges posed by technologically advantaged agents (TAAs), with some considering transaction taxes and restrictions on order cancellations or algorithm usage (see Linton et al. (2013), O'Hara (2015)). On the other hand, Menkveld (2013) provides direct evidence of how the entrance of TAAs can be beneficial (through reduced spreads).

We analyze a market where technology can afford TAAs a timing advantage in placing trades. We examine the actions of these TAAs in order to investigate whether their technological advantage results in an inequitable market. We achieve this by looking at the price impact of their actions on the non-TAA's utility. What we find is that rather than exploiting the non-TAAs, the TAAs are meeting a demand for payoffs in the most popular payoff states, resulting in an improved realized utility score for the non-TAAs.

The tote betting market provides the opportunity to separate investors with different utility functions. The conventional wisdom is that informed bettors bet late (see e.g. Asch et al. (1982)) and we present strong evidence to support this. Late bettors gain two main advantages: firstly, they obtain the most accurate estimate of the final odds, secondly, they hide their own probability estimates which might otherwise be revealed through their bets. The tote publishes prices based on updated pool totals at a frequency of 1 or 2 updates per minute. The penultimate tote cycle is the last set of odds displayed after which valid bets can be placed. A level of sophistication in bet timing and transmission is required to guarantee bet placement after this cycle but before the pools close.

The discussion on market regulation issues in betting markets to date has focused on the use of insider information rather than TAAs (see Peirson (2011)). Using time-stamped betting amounts on each horse, we separate the tote pool into those amounts bet before and after the penultimate tote cycle (the 'early' and 'late' betting periods). This allows us to separate the amount bet by the TAAs and the non-TAAs and to examine the utility functions implied by the distributions of these agents' bets associated with different odds.

2. Materials & Methods

2.1. Data

The data set includes tote market win pool betting for each runner over 174,000 races and includes over 2.8 million price quotes for 1.4 million runners across 39 tracks in the USA and Canada from May 2011 to August 2016. It includes the amounts bet on each runner to produce the final odds and the last predicted set of odds published by the tote (the end of the penultimate tote cycle). The difference between the two amounts being largely due to the amount placed by TAAs.

*Corresponding author

Email addresses: R.J.McGee@soton.ac.uk (Richard J. McGee), jej@soton.ac.uk (Johnnie E. V. Johnson)

Track-specific takeout rates (i.e. the percentage of the betting pool deducted by the tote to to cover operating costs and profit) are used to reconstruct the odds implied by the amounts bet. The tote also applies 'breakage', whereby odds implied by the amounts bet are rounded down to one decimal place before display, and bets are settled on this final figure.

The dividends, including breakage, are calculated as:

$$div_i = \frac{\left\lfloor 10 * \frac{\sum_i wamt_i}{wamt_i} * (1 - takeout) \right\rfloor}{10}, \qquad (1)$$

where $wamt_i$ is the amount bet on horse i.

2.2. The Favorite Longshot Bias

The favorite-longshot bias (FLB) is a phenomenon observed in betting markets whereby longer/shorter odds prices are over-/under-bet (relative to the probability of a payoff).[1] We use the FLB to test for evidence of risk-loving preferences among two distinct agents: TAAs and non-TAAs. To do this we report the coefficient from a conditional logistic regression (CL) (as per Bolton & Chapman (1986)) of the race outcome on the log of the odds.

The probability of runner i winning the race is given as:

$$p_i = \frac{\exp(\beta \log(div_i))}{\sum_i \exp(\sum_i \beta \log(div_i))} \qquad (2)$$

The β value is obtained by maximizing a log-likelihood (LL) score, $\mathcal{L}(\beta)$:

$$exp(\mathcal{L}(\beta)) = \prod_m \frac{\exp(\beta log(div_m)^w)}{\sum_i \exp(\beta log(div_m)^i)}, \qquad (3)$$

where $log(div_m)^i$ is the log of the dividend for runner i in race m and $log(div_m)^w$ is the log of the dividend for the winning runner in race m, $(-1 \leq m \leq 174,000)$.

An unbiased coefficient is given by $\beta = -1.0$, values more negative/positive than this indicate the presence of the FLB/reverse FLB.

In our data set, of the $1.18 billion bet in the win market, $396 million was bet in the final tote cycle (i.e. TAAs bet 33.6% of the total market volume). The results of estimating (2) for odds determined by the early and late bettors separately and combined are given in Table 1. The final (combined) odds demonstrate a clear FLB ($\beta = -1.11$), and this is more pronounced for the odds set by early bettors ($\beta = -1.152$). The move toward an unbiased measure is coming from the TAAs whose bets are focused heavily on the favorites. The average bet return for all agents corresponds to the average track take plus breakage (-18.4%). The non-TAAs do worse than the track take by 1.4%, while the TAAs are recovering some of the take, losing only 15.7%.

[1] See Ottaviani & Sørensen (2008) for a review of the main explanations for causes of the bias. Williams & Paton (1998) defined two separate bettor types to explain variation in the FLB and we also split the betting pool into two separate representative agents.

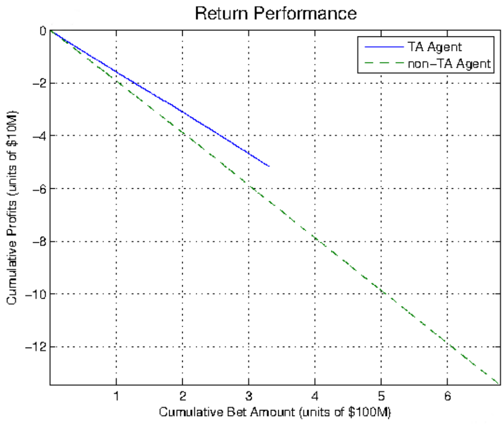

Figure 1: A comparison of the cumulative returns of the TAAs and non-TAAs.

The higher LL score associated with the estimations based on the final (cf. early) odds suggest that the bets of TAAs improve the accuracy of the early odds (LL-ratio test score: 13,441, $p < 0.0001$).

2.3. Implied Utility Functions

We use Weitzman's (1965) method to imply two different representative agent utility functions for the TAAs and non-TAAs. We first separate the data into return bins based on the final payoffs. We only include bins with at least 100 prices present in the sample and this leads to 961 return values in the range 0.1 to 102.6. An empirical probability of winning - the ratio of the number of victories divided by the number of entries - was then calculated for each return bin. This yields 961 points of the form (x, p), where x is the value of return to the dollar and p is the empirical probability of a payoff associated with that return. The bins are estimated over 1.4 million entrants. We use the same corrected hyperbolic form as Weitzman and this is fit to the data using non-linear least squares regression. The resulting curve represents the empirical probability as a function of the return implied by the tote payoffs. The function parameters are similar to the ones obtained in Weitzman (1965)[2].

$$p_i = \frac{0.950}{x_i} - 0.097 \frac{\log(1 + x_i)}{x_i} \qquad (4)$$

The resulting function is displayed in Figure 2. To evaluate the utility of wealth, we must transform the analysis from the return domain to actual monetary values. To do this we assume an average bet size of $5, in line with Weitzman (1965). This value is not important as we are looking at the relative change in utility (i.e. not absolute values). The money won by an agent in dollar terms is then given by $m = 5x_i$. The utility, K, of every point (p, m) on the agent's utility curve is

[2] $p_i = \frac{1.011}{x_i} - 0.087 \frac{\log(1+x_i)}{x_i}$

Table 1: The conditional logistic $log(div_i)$ parameter, β, and likelihood scores for the three sets of odds: early bettors (non-TAAs), late bettors (TAAs) and the final combined odds. The pool percentage breaks down the percentage of the total amount wagered into amounts up to and after the odds at the penultimate cycle are calculated. Test statistics for the β coefficients are given in square brackets, we can reject the null hypothesis of an unbiased coefficient, $H0 : \beta = -1.0$, with p values < 0.00001 in each case.**The late bettor LL value is NA because the late bettors bet zero on some horses that win races, resulting in a likelihood score of zero. In this case they are completely avoiding betting on some horses as they represent bad value, rendering the difference in amounts from the penultimate to the final cycle an incomplete market.

	Early	Late	Final (Combined)
Ave. Pool Percentage	66.4 %	33.6 %	100 %
CL Coefficient (β)	-1.152	-0.807	-1.111
	[-363.3]	[-342.9]	[-356.4]
Average Bet Return	-19.82%	-15.7%	-18.4%
LL score	-252,808	*NA***	-246,207

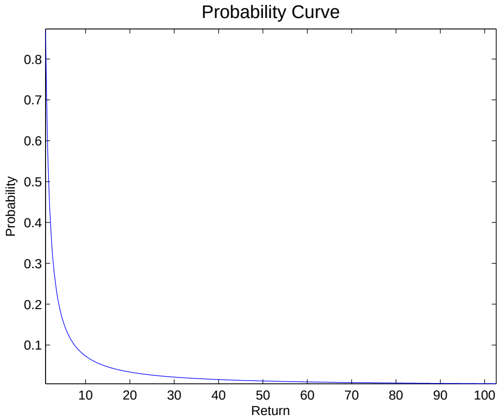

Figure 2: The empirical average of the probability of horses finishing first where the return on the x–axis is the return implied by the odds on the horse on the penultimate tote cycle.

given by $K = U(p, m) = pU(m)$. This gives the utility function:$U(m) = \frac{K}{p(m)}$. The K value is set to $5p(\$5)$ such that the utility of \$5 is 5 utils. The resulting utility functions evaluated using the penultimate and final pool amounts, and the TAA amounts, are illustrated in figure 3. To calculate the realized utility of the non-TAAs we evaluate the utility function in figure 3 using the payoffs received by the non-TAAs based on their actual bet amounts and the dividend that they receive: (i) when the TAA amounts are removed and (ii) when they are included. From Table 2 it can be seen that the utility of the non-TAAs actually increases when the TAA's bets are included (despite the greater loss the non-TAAs incur: see Table 1).

This apparent paradox is explained by the fact that the non-TAAs are risk-loving and place greater value on high dividend payoffs. The TAAs, as shown by the β of -0.8 in Table 1, are betting more on shorter odds runners, lengthening the longer odds. Consequently, the TAAs are supplying more of the payoffs that the non-TAAs want and receive a premium to do so.

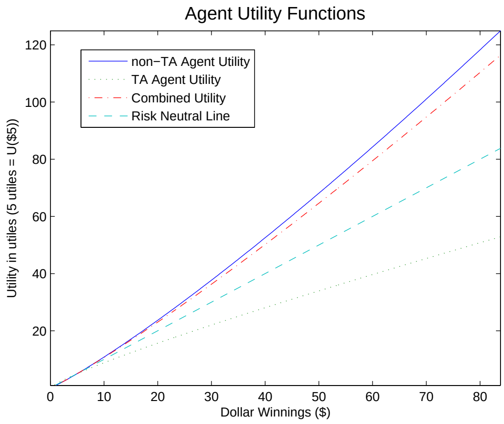

Figure 3: The agent utility functions implied by the win dividends on the penultimate and final tote cycles and the amounts bet in between by TAAs.

Table 2: The average utility using the non-TAA, TAA and combined utility functions over the 174,000 race outcomes using 1. the dividends at the penultimate cycle excluding the TAA price impact and 2. the final dividends including the TAA's price impact.

	1. Pen. Prices	2. Final Prices
non-TAA Utility	6.16	6.31
TAA Utility	5.28	5.28
Combined Utility	6.08	6.22

3. Conclusion

We find that the price impact of TAAs improves the expected utility of non-TAAs in US and Canadian tote betting markets. The TAAs appear to perform a role analogous to market makers in options markets.[3] They receive a premium from the other market participants and in return provide better returns for the most popular state payoffs. We demonstrate that a net transfer of wealth to a TAA does not necessarily imply an inequitable market. Our findings point to the need for a more nuanced approach, including agent utility analysis, when considering regulatory intervention in markets containing TAAs and non-TAAs.

Acknowledgements

The authors would like to thank both Bruce Soulsby and the Las Vegas Dissemination Company (LVDC) for providing the data.

Asch, P., Malkiel, B. G., & Quandt, R. E. (1982). "Racetrack betting and informed behavior". *Journal of Financial Economics*, 10(2), 187-194.

Bolton, R., & R. Chapman (1986), "Searching for positive returns at the track: A multinomial logit model for handicapping horse races." *Management Science* 32(8), 1040-1060.

Johnstone, D., (2013) "A Simple Automated Market Maker for Prediction Markets" in *The Oxford Handbook of the Economics of Gambling*, ed. L Vaughan Williams and D S Siegel , Oxford University Press, New York, United States, pp. 543-59.

Linton, O., O'Hara, M., & Zigrand, J. P. (2013). "The regulatory challenge of high-frequency markets". *High-Frequency Trading*. Risk Books.

Menkveld, A. J. (2013). "High frequency trading and the new market makers". *Journal of Financial Markets*, 16(4), 712-740.

O'Hara, M. (2015). "High frequency market microstructure". *Journal of Financial Economics*, 116(2), 257-270.

Ottaviani, M., and Sørensen, P. N. (2008). "The favorite-longshot bias: An overview of the main explanations". *Handbook of Sports and Lottery markets*, 83-101.

Peirson, J., (2011). "The economic analysis of sports betting by expert gamblers and insiders" in *Prediction Markets: Theory and Applications*, ed. L Vaughan Williams, Routledge International Studies in Money and Banking, 66, p189.

Weitzman, M. (1965), "Utility analysis and group behavior: An empirical study", *The Journal of Political Economy*, 18–26.

Williams, L. V., & Paton, D. (1998). "Why are some favourite-longshot biases positive and others negative?". *Applied Economics*, 30(11), 1505-1510.

[3]Market maker algorithms have also been developed for prediction markets (see Johnstone (2013)).