Introduction

Gambling is decision making under risk. On the surface, it seems a simple exercise, of gaining a profit or losing one's wager. While that is, of course, the essence of the exercise, a closer inspection reveals great complexity and scope. Indeed, this simplicity veiling a complex and general process is what has attracted academics from a variety of disciplines to consider gambling as a forum for investigating matters of wider importance than simply wagering. For instance, psychologists have used gambling to illustrate and inform about fundamental behavior in the face of uncertainty, and numerous behavioral biases have seen support. Economists bring a more rational, utility-based perspective to decision making under uncertainty, and wagering has provided an opportunity to test numerous theories and has generated new theories as well. Financial experts view gambling markets as possessing many of the intricacies of financial markets; however, the repeated and short-lived reduced form of gambling markets allows cleaner analysis than is generally possible in most financial markets. Thus, starker views of aggregate investor behavior are sometimes possible, including investigations of market efficiency. Statisticians have used gambling to motivate improved estimates of complicated probabilistic events. Mathematicians and management scientists have developed useful gambling strategies that have drawn on and extended efforts in other domains (e.g., Kelly betting and information transmission). While gambling is of academic interest in its own right, it has clearly been demonstrated to enhance our understanding of more far-ranging environments.

This volume is concerned with racetrack betting. As evidenced by this collection of articles, our understanding of racetrack betting has clearly drawn from and correspondingly returned something to all the aforementioned fields of psychology, economics, finance, statistics, mathematics and management science.

It will be helpful to briefly describe the workings of the racetrack betting market. We adopt the terminology of U.S. racetracks. Typically, six to twelve horses are in a race and several wager types are available. A *win* bet requires one to name the winner of the race. A *place* (*show*) bet requires one's horse to finish at least second (third). Bets to win, place, and show involve naming just one horse and are termed *straight* bets. There are other wager types, called *exotic* bets, that require predicting the outcome of two of more horses. For example, and *quinella* bet requires one to name the horses finishing first and second, while an *exacta* bet requires naming these two horses in the right order. A *daily double* bet pays off when one correctly names the winners of two consecutive races. Separate pools are operated for each bet type.

The simplest wager is the win bet. Define W_i as the public's wager to win on horse i, $W \equiv \Sigma_i W_i$ as the public's win pool, and Q as the track payback. Then, in the event that horse i wins, QW/W_i is the payoff per dollar wagered on horse i. Define $O_i \equiv QW/W_i - 1$. The odds on horse i are expressed as O_i to 1, or O_i-1, and the return on a \$1 wager at odds of O_i-1 is the original \$1 plus another \$$O_i$.

The proportion of money bet on a horse, W_i/W, is called the win bet fraction. It is often interpreted as the public's estimate of the horse's win probability, which follows when bettors are expected profit maximizers. Several articles in this volume will show that the win bet fraction only roughly corresponds to the true win probability, though. The payoff functions for wagers other than win bets will be presented later.

The track take, 1-Q, is the commission collected by the management of the racetrack. A portion of that amount is paid to the state government in the form of a tax. For straight bets, the track take

varies from about 14 to 19 percent in the U.S. (There are other countries with much higher takes, such as the 26 percent in Japan.) Exotic bets typically have higher takes which are 20-25% or higher. The other transactions costs is *breakage* where the track rounds down returns to the nearest nickel or dime on the dollar. While breakage is not substantial for a single large payoff, it can be significant for a single small payoff and is always significant over a long sequence of wagers (see Hausch and Ziemba (1985)[1]). By comparison, breakage in Japan also involves rounding down payoffs while Hong Kong rounds payoffs to the nearest cent, which can mean rounding up or down. Thus, there is little if any cost due to breakage in Hong Kong.

There is typically about 20 minutes between races, during which time wagering occurs on the upcoming race. Most tracks publicly display the win, place and show bets of the public, and update these figures every minute or so during the betting period. The sheer quantity of numbers involved in exotic wagering (e.g. 90 exacta figures for a ten-horse race) means that these pools are often not publicly available or are displayed but only subsets at a time. While odds do change over the course of the betting period, actual payoffs in the parimutuel betting system are based on the final wagering of the public; thus, from the time one makes a wager, the odds can improve or worsen.

This volume is a compilation of recent research on various aspects of racetrack markets. The papers in this volume are contributed by researchers from different countries and various fields - Economics, Psychology, Finance, Mathematics, Statistics and Management Science/Operations Research. It is a blend of new articles produced for this volume and reprinted articles published from a wide variety of sources. The contributions cover empirical and theoretical studies of issues such as the favorite-longshot bias and efficiency of the markets to win, place, and show and markets for exotic betting. Other issues are ordering probability models and optimal betting strategies. Data is studied from racetracks around the world.

Part 1 presents some psychological studies of racetrack markets, including the earliest research in this area (Griffith (1949)[1] and McGlothlin (1956)[1]). Issues such as biases in probability estimation, the "gambler's fallacy," and framing are investigated at the racetrack.

Characteristics of the average racetrack bettor's utility function are studied in Part 2. Weitzman (1965)[1] and Ali (1977)[1] use racetrack data to empirically obtain a convex utility function, indicating that bettors are risk-lovers. With such a convex utility function, expected utility maximization generates the well-known favorite-longshot bias, i.e. bettors underbet favorites and overbet longshots.

Part 3 contains articles providing economic and mathematical insights about racetrack betting. There are three general issues in these papers: 1) optimal wagering strategies; 2) fundamental handicapping schemes; and 3) estimating ordering probabilities (the probabilities of various orders of finish of the horses). On the first of these, the earliest contributor was Isaacs (1953)[1] who developed an algorithm for optimal betting assuming our risk neutral bettor knows the true win probabilities of the horses. Kallberg and Ziemba (1994)[1] and Levin (1994)[1] extend this work to nonlinear utility and place and show betting. Thorp (1971)[1] describes the advantages of using the "Kelly criterion" originated by Kelly (1956)[2], which involves maximizing the expected logarithmic utility of capital. MacLean, Ziemba and Blazenko (1992)[1] discuss the tradeoff between return and security in gambling situations.

Like the stock market, investment strategies at the track are usually *technical* or *fundamental* in nature. Handicapping, or fundamental schemes, are described in several papers. Multinomial logit models are used with horse, jockey and race characteristics (Bolton and Chapman (1986)[1] and Chapman (1994)[1]). Benter's (1994)[1] scheme is both fundamental and technical in nature; to the aforementioned

fundamental characteristics, Benter's multinomial logit analysis adds the public's win odds.

Ordering probabilities are necessary for pricing place, show and exotic bets. Three probability models of running times (Harville, Henery and Stern) are the basis for most of the research on estimating ordering probabilities. An application of these models is reported in Lo (1994a)[1].

The win market is the subject of Part 4. Using Fama's (1970)[2] definition of weak form efficiency, various papers such as Ali (1977)[1], Snyder (1978)[1] and Asch, Malkiel and Quandt (1982)[1] conclude that systematic differences exist between the public's subjective win probabilities (as measured through the odds) and the objective win probabilities (as measured by actual outcomes). The bias is such that favorites tend to be underbet and longshots tend to be overbet, hence the term "favorite-longshot bias." Quandt (1986)[1] proves that the bias implies risk-seeking bettors. The empirical studies show that the bias for favorites is insufficient to overcome the track take, except for extreme favorites, which are relatively rare (Ziemba and Hausch (1986)[2]). Thus, for practical purposes, the win market is weak form efficient despite the bias. Exceptions to the favorite-longshot bias have been documented by Busche and Hall (1988)[1] and Busche (1994)[1] for tracks in Hong Kong and Japan. Other papers apply various forms of logit model to analyze win market efficiency (see Figlewski (1979)[1], Asch, Malkiel and Quandt (1984)[1], and Lo (1994b)[1]).

Part 5 discusses the place and show markets, and the existence of the most well-documented of the racetrack's inefficiencies. Ritter (1994)[1], Asch, Malkiel and Quandt (1984,86)[1] and Hausch, Ziemba and Rubinstein (1981)[1] test various approaches for identifying inefficiencies. Hausch, Ziemba and Rubinstein (1981)[1] study an optimal growth betting strategy to exploit these inefficiencies. This strategy is extended in Hausch and Ziemba (1985[1],1990a[1],1990b[1]),

Part 6 presents studies on efficiency of exotic markets. Ali (1979)[1] compares returns from a double bet with a equivalent parlay, a double bet that is self-constructed by betting to win on two consecutive races, and concludes that the two bets are "equally priced." Asch and Quandt (1987)[1] and Lo and Busche (1994)[1], however, find that the double is more profitable than the parlay. Hausch, Lo and Ziemba (1994)[1] present methods to fairly price and thus detect mispricings of various exotic wagers. Post position bias is investigated by Canfield, Fauman and Ziemba (1987)[1] and Betton (1994)[1]. Their results indicate that such a bias does exist but that, in most circumstances, it is fully accounted for in the public's odds. Other studies on exacta and quinella bets appear in Asch and Quandt (1987)[1] and Kanto and Rosenqvist (1994)[1].

Racetrack betting in a number of the countries of the Commonwealth and in Asia are discussed in Part 7. Many of these countries use a fixed-odds system rather than a parimutuel betting system. A fixed-odds system is run by bookmakers who offer odds that, upon purchase, become fixed (unlike the parimutuel betting system). While the bookmakers can change their odds at any time, the new odds apply only to new wagers. By adjusting their odds, bookmakers attempt to balance their books so as to ensure themselves of a profit regardless of which horse wins. The papers in this section show that the favorite-longshot bias is present even in fixed-odds systems, i.e., the bookmakers' odds, in order to balance their books, reflect the public's tastes for the favorite/longshot bias.

[1] included in this volume
[2] cited in the Annotated Bibliography