ECON30580 Economics of Betting Markets
9. Monopoly and Oligopoly

Karl Whelan

School of Economics, UCD

Spring 2026

Part I

A Monopoly Bookmaker Model

Another Micro 101 Recap: Monopoly

- In a perfectly competitive market, firms set prices equal to marginal cost.
- In a monopolistic market, however, prices are higher.
- Specifically, the monopoly sets prices as a mark-up on marginal cost.
- The size of this markup depends on the **elasticity of demand** for the monopolist's product.
  - If the elasticity of demand is relatively low, then the monopolist maximises profits by setting a relatively high price, taking advantage of its customers not being very price sensitive.
  - If the elasticity of demand is relatively high, then the monopolist maximises profits by setting a relatively low price because setting a higher price would reduce volumes too much.
- Bookmakers might seem like a very different business to textbook firms—the prices they set (the odds) determine both their costs (via payouts to customers) and their revenues (since higher odds attract more bets).
- But it turns out that monopoly bookmakers follow the same principles as normal Micro 101 monopolists.

Risk-Neutral Monopoly Bookmakers

- Consider a risk-neutral bookmaker who knows p, the probability that a bet on a specific outcome will win and who has costs per each unit bet of μ.
- Demand for bets is assumed to be $B(D)$ where D is the decimal odds the bookmaker offers and $B'(D) = \frac{dB}{dD} > 0$.
- Revenue is $B(D)$ and expected costs from offering the bet are $(\mu + pD) B(D)$.
- So the bookmaker's expected profit from offering odds of D on this bet is

$$E(\Pi) = (1 - \mu - pD) B(D)$$

- What is the expected-profit-maximising odds to set? Taking derivatives with respect to D (and using the product rule of differentiation) the condition for optimal odds is

$$\frac{\partial E(\Pi)}{\partial D} = (1 - \mu) B'(D) - pB(D) - pDB'(D) = 0$$

- This solves to give

$$D = \frac{1 - \mu}{p} - \frac{B(D)}{B'(D)}$$

Odds and Elasticity

- Optimal odds are

$$D = \frac{1-\mu}{p} - \frac{B(D)}{B'(D)}$$

- Defining the elasticity of demand for the bet as

$$\epsilon = \frac{DB'(D)}{B(D)}$$

- Optimal odds are

$$D = \frac{\epsilon}{\epsilon + 1} \frac{1-\mu}{p}$$

- $\frac{1-\mu}{p}$ are the odds that would be offered in a zero-profit competitive bookmaking market. The fraction $\frac{\epsilon}{\epsilon+1}$ is less than one but approaches one as ϵ gets bigger.
- So the odds are a "markdown" on zero-profit competitive odds and bets with higher elasticities of demand will have lower markdowns.
- If you think of the inverse odds as the "price" of the bet, you get the standard price markup formula.

Demand for Bets With Disagreement

- We will assume that beliefs about the probability p are uniformly distributed over $[L, H]$.
- Bettors are risk-neutral and bet when $\tilde{p}D > 1$.
- When they bet, they place a bet of size 1.
- We showed before that if everyone places the same size bet, the demand for bets at odds of D (where $\frac{1}{H} < D < \frac{1}{L}$) is just the fraction of people with beliefs above $\frac{1}{D}$. This is

$$B(D) = \frac{H - \frac{1}{D}}{H - L}$$

- The monopolist's expected profit on the bet is

$$E(\Pi) = (1 - \mu - pD)\, B(D)$$

- We can use calculus to figure out the expected profit maximizing odds.

Maximizing Profits

- The monopolist's expected profit on the bet is

$$E(\Pi) = (1 - \mu - pD)\left[\frac{H - \frac{1}{D}}{H - L}\right]$$

- Let's multiply this out so we can identify the separate terms involving D and then differentiate them.

$$E(\Pi) = \frac{(1 - \mu)\left(H - \frac{1}{D}\right)}{H - L} - \left[\frac{pDH - p}{H - L}\right]$$

- Differentiating this with respect to D gives

$$\frac{dE(\Pi)}{dD} = \frac{1 - \mu}{H - L}\frac{1}{D^2} - \frac{pH}{H - L}$$

- Setting this equal to zero and cancelling the $H - L$ terms, we get

$$\frac{1 - \mu}{D^2} = pH$$

Profit-Maximising Odds

- From

$$\frac{1-\mu}{D^2} = pH$$

we get the optimal odds

$$D = \sqrt{\frac{1-\mu}{pH}}$$

- It is useful to re-write this as

$$D = \sqrt{\frac{1-\mu}{p}\frac{1}{H}}$$

- This means the profit-maximising odds are a **geometric average** (square root of the product) of
  1. **The breakeven odds**: At $D = \frac{1-\mu}{p}$, the expected profit is zero.
  2. **The lowest acceptable odds**: $\frac{1}{H}$, the odds that will be accepted by the most optimistic person.

Features of Profit-Maximising Odds: 1

1. **Bookmakers earn profits and these increase with optimism of bettors**
   - The geometric mean is always lower than the arithmetic mean but, for realistic values of p, μ and H, the two are very similar here.
   - So odds are approximately half way in between the break-even odds and the odds the most optimistic person would accept.
   - The higher H is, the lower the odds will be.
   - The higher H is, the bigger the bookmaker's profits are.

Features of Profit-Maximising Odds: 2

2. **Pessimists don't matter**
   - The lower value for beliefs L does not matter.
   - It does not appear in the odds formula.
   - The bookmaker will never set odds higher than $\frac{1-\mu}{p}$ and people with $\tilde{p} < \frac{p}{1-\mu}$ won't accept any odds offered by the bookmaker.
   - All that matters for pricing is the upper bound on beliefs H.
   - This specific result depends on the assumption of a uniform distribution but remember that elasticity of demand is very similar when beliefs are Normally distributed so the optimal odds are almost the same in that case.

Features of Profit-Maximising Odds: 3

3. **No need for wisdom of crowds**
   - The crowd might be right on average. For example, beliefs could be uniformly distributed over $[p - \sigma, p + \sigma]$.
   - But that is irrelevant for pricing by monopolists.
   - The odds would be the same if beliefs were uniform on $[p, p + \sigma]$
   - Pricing is designed for unwise bettors.

Features of Profit-Maximising Odds: 4

4. **Higher costs reduce odds and profits**
   - A higher value for unit costs, μ reduces the odds.
   - This was also true in for perfectly competitive pricing.
   - But here the odds depend on the square root $\sqrt{1-\mu}$ rather than with $1-\mu$ as they do when there is competition.
   - Higher costs get split between lower odds (which hurts bettors) and lower profits for bookmakers.

Favourite-Longshot Bias

- Let $H = p + \sigma$ so σ indicates the extent of over-optimism. The optimal odds are

$$D = \sqrt{\frac{1-\mu}{p\,(p+\sigma)}}$$

- These odds also imply favourite-longshot bias. The expected payout on a bet of 1 unit is

$$pD = \sqrt{\frac{(1-\mu)\,p}{(p+\sigma)}}$$

  which depends positively on p. Higher bookmaker costs and more disagreement means higher losses for bettors.

- Suppose $\mu = 0.02$ and $\sigma = 0.06$. The expected loss rate $(pD - 1)$ for a bettor that chooses a bet with probability of success $p = 0.8$ is 6.7%. The corresponding expected loss for $p = 0.2$ is 13.2%.
- See the graph for expected loss rates on the next page.
- Note the "nonlinear" pattern of losses: **Loss rates accelerate as p falls.**

Monopoly Loss Rates with $\mu = 0.02$ and $\sigma = 0.06$

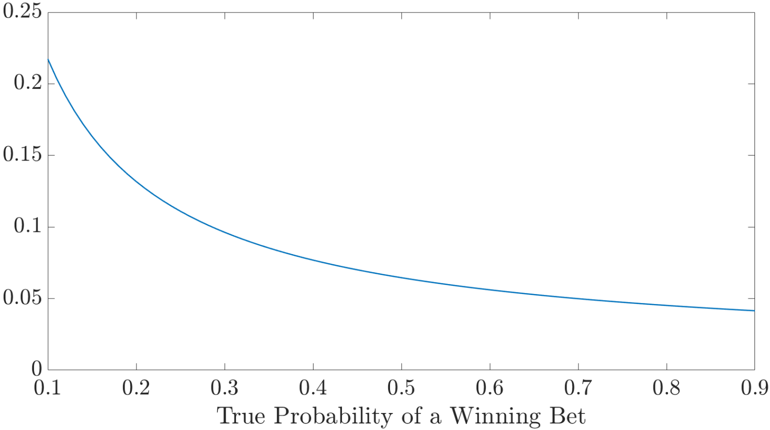

What Drives This? Variations in Elasticity of Demand

- Why do longshot bettors get worse odds than those who bet with favourites?
- The answer comes from something we worked out in the last lecture.
- We showed that the demand curve for bets on longshots was much less steep than for bets on favourites.
- In other words, demand for favourites is **more elastic** than demand for longshots.
- Monopoly bookmakers use this to set profit-maximising odds that get worse as p falls.

Example: Two Possible Outcomes

- Now consider a two-team event, with a favourite with probability p of winning.
- The odds are

$$D_F = \sqrt{\frac{1-\mu}{p\left(p+\sigma\right)}} \qquad D_L = \sqrt{\frac{1-\mu}{\left(1-p\right)\left(1-p+\sigma\right)}}$$

- We can use the odds to calculate the overround

$$\frac{1}{k} = \frac{1}{D_F} + \frac{1}{D_L} = \sqrt{\frac{p\left(p+\sigma\right)}{1-\mu}} + \sqrt{\frac{\left(1-p\right)\left(1-p+\sigma\right)}{1-\mu}}$$

- We can also calculate the bookmaker's expected profit rates on both bets ($1 - pD_F$ and $1 - \left(1-p\right)D_L$) and on the match itself.

$$\frac{1 - pD_F B_F\left(D_F\right) - \left(1-p\right)D_L B_L\left(D_L\right)}{B_F\left(D_F\right) + B_L\left(D_L\right)}$$

Bookmaker's expected profit rate on a match with two teams for different values of p with $\mu = 0.02$ and $\sigma = 0.06$

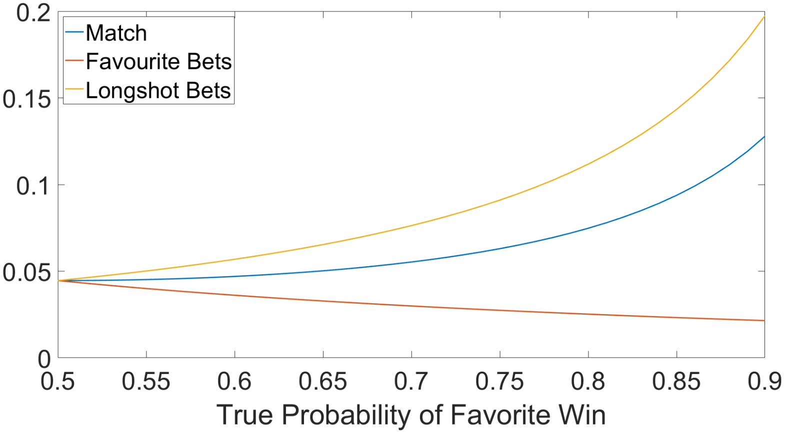

Overrounds in two-team contests for different values of p with $\mu = 0.02$ and $\sigma = 0.06$

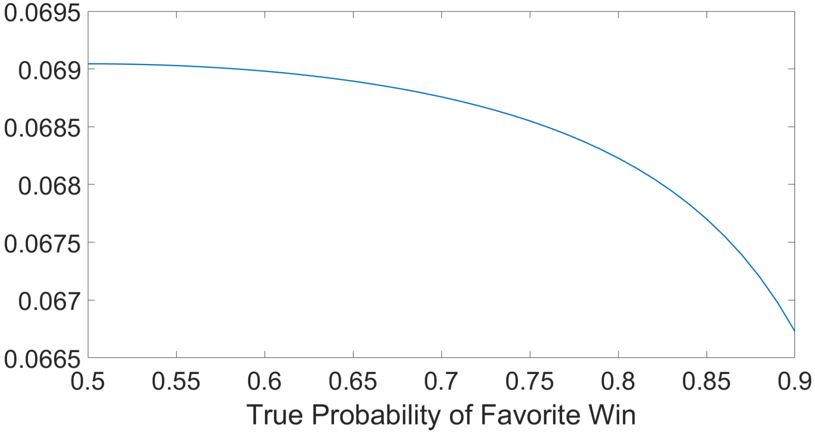

Profits and Overrounds

- Bookmakers make higher profits on games that are mis-matches (high value of p) than on games that are closer to toss-ups (p close to 0.5).
- The nonlinear pattern of increased loss rates for longshots (and thus increased profits for bookmakers) means the additional gains from bets on more extreme longshots offsets the slightly lower profits obtained from bets on a favorite with a high win probability.
- Normally, bettors assess whether a bookmaker is making a lot of profit on a game by calculating the overround.
- But look at the chart for the overround. It actually shows overrounds being lower for games with high values of p, so bettors would get the wrong signal.
- When there is a favourite-longshot bias, the overround calculation is not actually telling you how much profit the bookmaker is making.
- The overround calculation of the margin assumes strong betting market efficiency—equal returns on all bets—and that does not hold here.
- And the normalised probabilities implied by the model do not equal the actual probabilities: Too high for low values of p and too low for high values of p.

Actual win probabilities and normalised probabilities with monopoly odds with $\mu = 0.02$ and $\sigma = 0.06$

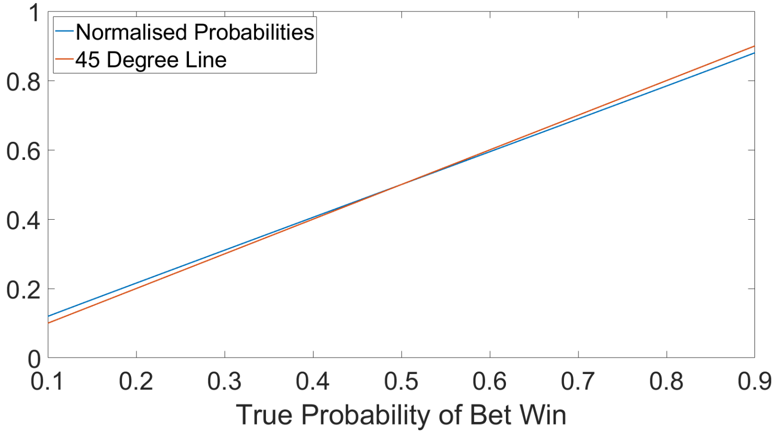

Part II

Oligopoly: Cournot Competition Among Bookmakers

A Cournot Oligopoly Model: Setup

- What happens when, instead of a monopoly, there is a small number of bookmakers?
- We could assume they compete on price and drive odds down to their competitive level
- But bookmakers report substantial profits, so this would not be a good model.
- Instead use a different model – the **Cournot oligopoly** model – where bookmakers set quantities but do not compete on price.
  - There are N identical bookmaking firms
  - All firms post the same decimal odds, D
  - Firms compete by choosing quantities of bets, B_i
  - Bets are split across firms, so total betting is

$$B = \sum_{i=1}^{N} B_i$$

  - We solve for a symmetric Cournot equilibrium with all firms choosing the same B_i

Demand for Betting at Odds D

- Bettors differ in their subjective beliefs about the probability of winning
- Beliefs are uniformly distributed on $[L, H]$
- Total betting demand at odds D is

$$B = \frac{H - \frac{1}{D}}{H - L}$$

- Higher odds attract more bettors by improving the expected payoff.
- Inverting the demand function gives odds as a function of total betting

$$D = \frac{1}{H - (H - L)B}$$

- Higher odds are associated with higher total betting.
- This captures that the industry can only increase total betting by offering higher odds.

Expected Profit for Firm i

- Firm i accepts betting volume B_i at odds D and operating costs are a fraction μ of stakes.
- Expected profit is

$$
\begin{aligned}
E(\Pi_i) &= (1-\mu-pD)\,B_i \\
&= \left(1-\mu-\frac{p}{H-(H-L)B}\right) B_i \\
&= \left(1-\mu-\frac{p}{H-(H-L)\sum_{j=1}^{N} B_j}\right) B_i
\end{aligned}
$$

- Each firm chooses B_i to maximise expected profit, internalising how its betting volume affects odds via total B.

First-Order Condition and Symmetric Equilibrium

- Differentiating expected profit with respect to B_i yields

$$\frac{dE(\Pi_i)}{dB_i} = 1 - \mu - \frac{pH - p(H-L)\sum_{k\neq i} B_k}{\left(H - (H-L)\left(B_i + \sum_{k\neq i} B_k\right)\right)^2} = 0$$

- This condition balances scale against the effect of betting volume on odds.
- In a symmetric equilibrium all firms choose the same betting volume

$$B_i = B_j \quad \text{for all } i, j$$

- Total betting is $B = NB_i$
- The first-order condition becomes

$$1 - \mu - \frac{pH - p(H-L)(N-1)B_i}{(H - (H-L)NB_i)^2} = 0$$

Solving for Equilibrium Betting Volume

- Rearranging gives

$$pH - p(H-L)(N-1)B_i = (1-\mu)\left(H - (H-L)NB_i\right)^2$$

- This can be re-written as a quadratic equation in B_i

$$aB_i^2 + bB_i + c = 0$$

where

$$a = (1-\mu)(H-L)^2N^2$$

$$b = -2(1-\mu)H(H-L)N + p(H-L)(N-1)$$

$$c = (1-\mu)H^2 - pH$$

- It yields two algebraic solutions for B_i.
- One solution implies positive odds and is economically meaningful
- The other implies negative odds and is discarded
- There is therefore a unique relevant Cournot equilibrium

Competition and Loss Rates

- From B_i and N, we can calculate total betting $B = NB_i$ and thus the equilibrium odds D and expected loss rates $1 - pD$.
- We can see how the extent of competition affects outcomes for bettors by holding L, H, and μ fixed and varying N.
- Results:
  - Loss rates are lower than monopoly loss rates
  - Loss rates fall as N increases
  - As $N \to \infty$, loss rates converge to μ
  - As N goes closer to one, loss rates converge towards their monopoly levels.
- For realistic N, favourite-longshot bias remains but is weaker than under monopoly.
- If lots of bookmakers compete in a market, it could have no favourite-longshot bias.
- More generally, we can say that more competition between bookmakers implies lower loss rates and reduced favourite-longshot bias.

Loss rates for bettors for $\mu = 0.02$ and $\sigma = 0.06$ and various values of p and N

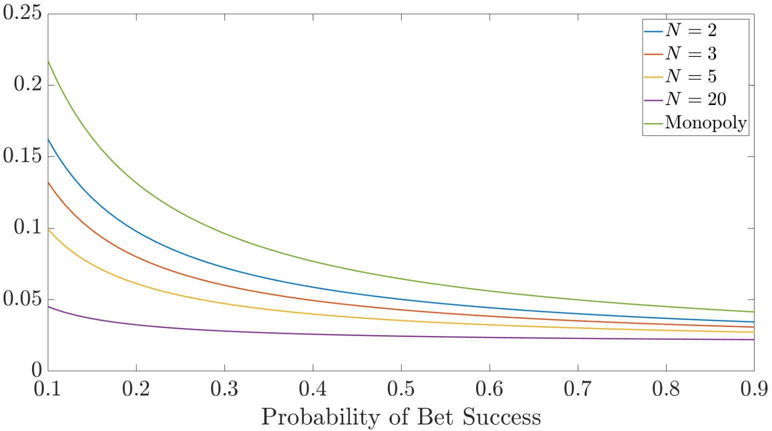

Supplementary Material in the Draft Book

- Chapter 13: How Far Will You Stretch? Bookmakers and Elasticity