University of Antwerp

Faculty of Business and Economics

Efficient Prediction Markets

Jonas Vandenbruaene

Doctoral dissertation submitted to obtain the degree of Doctor in Applied Economics

Supervisors: Prof. dr. Jan Annaert & Prof. dr. Marc De Ceuster

June 2022

Efficient Prediction Markets

Jonas Vandenbruaene, Antwerpen, 2022

ISBN: 9789057287503

Depotnummer: D/2022/12.293/22

Het grootste geluk dat ik ooit had,

was dat ik de juiste ouders koos.

Opgedragen aan Marc en Linda

Doctoral Jury

Supervisors

Prof. dr. Jan Annaert	(University of Antwerp)
Prof. dr. Marc De Ceuster	(University of Antwerp)

Members of the Doctoral Jury

Prof. dr. Marc Deloof	(University of Antwerp) (Chairman)
Prof. dr. Koen Vandenbempt	(University of Antwerp)
Prof. dr. Thomas Peeters	(Erasmus University Rotterdam)
Prof. dr. Niklas Wagner	(University of Passau)

Acknowledgments

Academics like to refer to their home university as their alma mater. Their nourishing mother. I am very thankful that I could spend 10 years, from 2012 to 2022 at my alma mater, the University of Antwerp. The first 5 years as a student in business engineering and the last 5 years as a PhD student and teaching assistant. I entered the first bachelor year as an 18-year-old who only had a vague idea of what economics was about and I leave 10 years later as a doctor in the domain. I have grown a lot academically and intellectually but also personally and socially. I have bumped into the people who would become (and still are) my best friends and on 11 March 2015, I serendipitously met Sarah, my lovely future wife in *Aula Rector Dhanis*. It is hard to imagine how different life would be without everyone I met here. The University also offered me the opportunity to explore the world. I got the chance to study in the United States and India, co-organized a yearly summer school to Georgetown University and Johns Hopkins and attended international conferences. Furthermore, the university environment stimulated me to try bridge the gap between academia and the wider public by writing op-eds and book reviews for major Belgian news outlets and by discussing economic topics on podcasts. Looking back on my days in and around *Het Hof Van Liere,* it is clear that I owe a huge debt of gratitude to my alma mater and it is due time to thank the people who were closely involved.

Writing a PhD is a team effort, so I would first and foremost like to thank my supervisors: Jan Annaert and Marc De Ceuster. Thank you for letting me start a PhD in the first place and coaching me diligently along the way. Thank you for giving me the time to freely roam across the vast landscape of financial economics and beyond. It gave me the opportunity to develop my own voice and interests while I could always rely on your guidance whenever I wandered off too far. Also, as a teaching assistant you gave me the opportunity, even in the very first year of my PhD, to teach to large groups of 1st year bachelor students and to master students at the Antwerp Management School. This was scary and out of my comfort zone at first, but I got a lot out of it. Thank you for the confidence. Furthermore, I want to thank all the members of my doctoral jury: Marc Deloof, Niklas Wagner, Koen Vandenbempt and Thomas Peeters. Thank you for taking the time to read my dissertation diligently. The vivid discussion during the internal defense and your comments on the empirical papers were very helpful.

A big thank you goes out to all colleagues who colored my days at the faculty. Thank you Fynn Gerken for the countless discussions and being a great *Resonanzboden*. We will write that insider trading paper we often talked about one day! Thank you Kevin Van Mencxel and Jef Van Cappellen to be my partners in crime for the PhD courses at KU Leuven and our long conversations on the train. Thank you Kim Van Overvelt and Peter De Bruyn for our trips to Washington DC. The summer school was always a highlight in the academic year for me. Thank you Yanou Ramon for almost going to Washington DC with me, twice. Luckily, we can somewhat make up for it on the third attempt now the pandemic is finally over. Thank you Yanou for the countless coffee breaks, running sessions and the legendary after works! The

timing of our PhDs was very much in sync so that we always understood each other's ups and downs really well. Thanks for all the conversations on our dissertations and later on what we were going to do afterwards. Thank you Joris Beckers, with just a bit more training I'm confident that we will win the ASL competition. Thank you Pieter De Rijck, one day I will be able to finally return your ping-pong serve. Thank you Dieter Brughmans. Our brief momentum project will always give us the bragging rights that we worked on a supercomputer during our PhDs. Thank you Gertjan Verdickt and Leentje Moortgat. You were among the most senior PhD students at our department when I started so I picked up a lot from you. Thank you Tristan De Blick for our lively debates on everything ranging from stock markets to universal basic income to Milton Freedman and Ayn Rand. Your entrepreneurial spirit is contagious! Thank you Laura Caluwe for the mental support that followed out of all the talks we had about life in academia. Thank you Tijmen van Kempen for instigating many great evenings on *het Mechelseplein* and for your wonderful gastronomical skills. Thank you Kris Hardies for coordinating the brown bags and pre-discussions. I always looked forward to these meetings and got a lot out of them. Thank you to all ACF colleagues who participated in these sessions and livened up the discussion. Thank you Hairui Zhang for the many exam supervisions we spent together and your very skillful support with Educloud issues. Thank you Kelly Martens, Stiene Praet, Sven Buyle, Bram Roosens, Sisangile Nduna, Mohammadamin Javaheri, Laura Caluwe and Floor Verbiest for being such a nice Young Research Network team. I loved organizing events and contributing to the social life at the faculty. Next, I would like to thank the members of the executive committee, the faculty council and the department council. It was a pleasure to be one of the AAP/BAP delegates in the respective bodies and to get an insider look into the mechanics of the faculty. Last but not least, thank you to all other B3 inhabitants for all the lunches, coffee breaks and after work drinks. Thank you Sahar Salehi, Jonas Vandennieuwenhuysen, Nina Marien, Michiel Van Roy, Michiel Dierckx, Carlos Tubbax, Jo Mentens, Tu Luong, Freek Van Doninck, Anil Berkin, Tom Vermeire, Jovana Cadenovic, Pedro Moraya Barros, Gilles Oorts and Tim Huygh. I hope to see all of you back soon and as they say, I buy the first drink for anyone mentioned in these acknowledgements and the first three for anyone who should be mentioned but isn't!

Dear friends, you know who you are, thanks for the support. The last years, we followed an implicit rule that as our conversations went on and the night got darker, every discussion we had eventually morphed into a debate on the stock market and market efficiency. It was very entertaining and instructive to discuss the issues I think about on a daily basis with you, especially since market efficiency is such a controversial topic in our circles.

Mama, Papa and Dorien, you were with me since the very beginning. Thank you first and foremost for your most unconditional love. The older I get, the more I start to understand how exceptional it is. Thank you for being such a warm family. Looking back, our home was the ideal environment to carefreely grow up in and to thrive. With two ferocious readers as parents, our house was always filled with books and from early on, I developed the habit of

reading the newspaper every day as it was always on our living room table anyway. That I now read and write for a living undoubtably has its origins there.

Thank you Papa for being such a great teacher. Thank you for sparking my enthusiasm in science and your super wide range of interests. Discussing anything from supernovas to natural selection or mirror neurons with you made me eager to go out to study and learn more myself. I still vividly remember when I was six or seven, standing in the dark basement with you while you were shining with a flashlight on a globe to explain day, night and time zones to me. Even now you still sometimes take out a piece of paper to sketch a bacteriophage or explain the difference between DNA and RNA vaccines. That I now underline sentences and make notes in every book I read is behavior I totally copied from you.

Thank you Mama for being such a great coach. Thank you for your dedication and perseverance. You were always very involved in anything we did. In my first year of saxophone classes, you were attending the first lessons to make sure I picked up everything I needed to get started. Or when I had to memorize long lists of Italian musical terms, you sat beside me and helped me come up with mnemonics that made it much easier to remember them. We didn’t stop until I knew all the terms by heart and I still remember some of those mnemonics today. Later in high school, you used to go through my textbooks the day before the exam to test whether I knew its contents as well as I said I did. Even in university you often knew my exam schedule better than me. You made absolutely sure I developed the character needed to study hard and achieve dreams.

Mama and Papa, thank you for your total devotion to Dorien and me. The greatest luck we ever got was choosing the right parents. This PhD is dedicated to you.

Last but certainly not least, dear Sarah, what would I do without you? You have experienced my entire PhD journey from up close. Thank you for your patience whenever I was working on a paper late at night or in the weekend. Thank you for always listening with interest to all my new crazy ideas (or at least convincingly pretending to). Thank you for your love and your enormous support. Thank you for being my best friend.

Jonas Vandenbruaene

Antwerp, 2022

Abstract

A core question in financial economics is whether markets are informationally efficient i.e., whether asset prices accurately reflect all available information. Although the empirical literature on market efficiency is vast, the subject is still heavily debated among academics and practitioners. A fundamental issue with informational efficiency is that it is virtually untestable in traditional financial markets. For example, in an ideal world, researchers would compare stock prices on stock markets with their true values to check whether they are aligned or not. However, as true values of stocks are never available, this is not possible. This untestability of market efficiency is referred to as the joint hypothesis problem.

In this dissertation, we try to make an original, unconventional contribution to the market efficiency debate by studying prediction markets. Prediction markets are platforms where people can bet on the outcome of future events, like a presidential election or a football game. Prediction markets have many characteristics that make them interesting research labs. Their main advantage is that the outcomes of the events are exogenously revealed, the market prices collide with reality. This allows researchers to systematically compare market prices with terminal values to detect mispricing, which is not possible on stock markets and circumvents the joint hypothesis problem.

This dissertation contains three empirical chapters. In the first, we review 40 years of literature on mechanical trading strategies in sports prediction markets. Many individual studies claim to have found profitable trading strategies which implies inefficient market pricing. However, when we consider the entire literature, the evidence is consistent with an efficient market where profit opportunities are chance results. Furthermore, we argue for more meta-scientific reflexes and put forward a hurdle rate of $|z|>3$ to benchmark the statistical significance of empirical results.

The second empirical chapter studies the evolution of the UK fixed odds betting market between 2000 and 2018. This period is of particular interest as it coincides with the rise of online gambling. We find that over this period, transaction costs decreased very significantly, both statistically and economically. Furthermore, we document a decrease in the favorite-longshot bias, a persistent anomaly in prediction market research.

The third empirical chapter tests whether time series momentum, a well-known irregularity in traditional financial markets, is also present in prediction market data. We find that a time series momentum effect is indeed present and by leveraging the prediction market characteristics, we show it is consistent with behavioral underreaction and not a rational premium for variance or skewness.

Dutch Abstract

EFFICIENTE VOORSPELLINGSMARKTEN

Een kernvraag in de financiële economie is of markten informatie-efficiënt zijn, d.w.z. of activaprijzen alle beschikbare informatie accuraat weerspiegelen. Hoewel de empirische literatuur over marktefficiëntie omvangrijk is, wordt het onderwerp nog steeds hevig besproken door academici en beroepsbeoefenaars. Een fundamenteel probleem met informatie-efficiëntie is dat het op traditionele financiële markten vrijwel niet kan worden getest. In een ideale wereld zouden onderzoekers bijvoorbeeld de koersen van aandelen op de aandelenmarkten vergelijken met hun werkelijke waarden om na te gaan of zij al dan niet op elkaar zijn afgestemd. Aangezien de werkelijke waarden van aandelenkoersen echter nooit beschikbaar zijn, is dit niet mogelijk. Dat informatie-efficiëntie niet te testen is wordt de *joint hypothesis problem* genoemd.

In dit proefschrift proberen we een originele, onconventionele bijdrage te leveren aan het debat over marktefficiëntie door voorspellingsmarkten te bestuderen. Voorspellingsmarkten zijn platformen waar mensen kunnen wedden op de uitkomst van toekomstige gebeurtenissen, zoals een presidentsverkiezing of een voetbalwedstrijd. Voorspellingsmarkten hebben vele kenmerken die hen tot interessante onderzoekslaboratoria maken. Hun belangrijkste voordeel is dat de uitkomsten van de gebeurtenissen exogeen worden onthuld, de marktprijzen komen in aanraking met de realiteit. Hierdoor kunnen onderzoekers systematisch marktprijzen vergelijken met de uitkomsten van de gebeurtenissen om efficiëntie te testen, wat niet mogelijk is op aandelenmarkten, en zo de *joint hypothesis problem* omzeilen.

Dit proefschrift bevat drie empirische hoofdstukken. In het eerste hoofdstuk bestuderen we 40 jaar aan literatuur over mechanische handelsstrategieën in sportvoorspellingsmarkten. Veel afzonderlijke studies beweren dat ze winstgevende handelsstrategieën hebben gevonden die inefficiënte marktprijzen impliceren. Wanneer we echter de hele literatuur in ogenschouw nemen, is het bewijs consistent met een efficiënte markt waar winstopportuniteiten toevallig voorkomen. Bovendien pleiten we voor meer metawetenschappelijke reflexen en stellen we een drempelwaarde van $|z|>3$ voor om de statistische significantie van empirische resultaten te benchmarken.

Het tweede empirische hoofdstuk bestudeert de evolutie van de Britse markt voor *fixed odds* weddenschappen tussen 2000 en 2018. Deze periode is interessant omdat ze samenvalt met de opkomst van online gokken. We stellen vast dat de transactiekosten in deze periode zeer significant gedaald zijn, zowel statistisch als economisch. Bovendien documenteren we doorheen onze steekproef een afname van de *favorite-longshot bias*, een anomalie in voorspellingsmarktonderzoek.

In het derde empirische hoofdstuk wordt getest of tijdreeksmomentum, een bekende onregelmatigheid in traditionele financiële markten, ook aanwezig is in voorspellingsmarkten. We vinden een tijdreeksmomentumeffect en door gebruik te maken van de kenmerken van de voorspellingsmarkt tonen we aan dat de data consistent is met onderreactie en niet met een rationele premie voor variantie of scheefheid.

Contents

List of Figures

List of Tables

Preface

This dissertation is submitted with the aim of obtaining a PhD degree. PhD stands for *philosophiae doctor*, doctor in philosophy. Notice that it does not say doctor in economics or doctor in finance, it's doctor in philosophy. Although the abbreviation might seem archaic and inaccurate given how wide science has branched out since PhDs were first rewarded, I think it still captures the spirit of the undertaking remarkably well. Researchers in all domains are thinking about how the world works and are trying to understand it better. We are all philosophers or "wisdom lovers" in the original Greek meaning. To do justice to the term "PhD", I will take the liberty in this preface to philosophize on the broad subject of my research domain and reflect about what it means for society. The preface is purposefully written from a bird's-eye perspective to make it readable for a general audience. There is enough room for more detailed and technical expositions in later chapters.

"The ideas of economists and political philosophers, both when they are right and when they are wrong, are more powerful than is commonly understood. Indeed, the world is ruled by little else. Practical men, who believe themselves to be quite exempt from any intellectual influences, are usually slaves of some defunct economist. Madmen in authority, who hear voices in the air, are distilling their frenzy from some academic scribbler of a few years back. I am sure that the power of vested interests is vastly exaggerated compared with the gradual encroachment of ideas."

Keynes (1936, p. 383)

Markets

Many important decisions in our world are made via markets. Markets price our products, value our companies and determine our wages. These market outcomes have a large impact on our behavior and our choices in life, so it is vital to know whether markets function properly. The answer to this question is central to economics and has far-reaching practical implications. Indeed, most policy debates, for example on health care, education, housing, climate change or the stability of the financial system, can be boiled down to questions on how far we should go with the outsourcing of decision making to market forces.

In the last century, our ideas about markets have probably been influenced the most by Friedrich Hayek[1], one of the godfathers of neoliberalism and winner of the 1974 Nobel prize in economics. Hayek lived and worked in an era when the most important political and economic question was whether economies should be centrally planned or not. Having

[1] "Perhaps no person better represents the notion of the power of ideas in the twentieth century than does F.A. Hayek." (Caldwell, 2013, p. 33)

endured the hardships of the Great Depression in the 1930s, many intellectuals at the time, both on the left and the right[2], favored central planning or scientific planning as it was called. It seemed obvious that an economy where teams of brilliant bureaucrats meticulously manage the output of different industries would be more rational, more prosperous, and less volatile than a decentralized, chaotic economy where no one is in charge. “Planning is the grand panacea of our age" wrote Robbins (1937, p. 3) when describing the *zeitgeist*. Politicians also joined the hype, even in England where the Labour Party in 1942 explicitly advocated that “a planned society must replace the old competitive system” (Labour Party, 1942). Hayek was not a fan of central planning and government intervention. He argued, against the general consensus at the time, that we should not rely on central planning boards to manage our economic activities, but on markets (Hayek & Caldwell, 2014). His insights into the effectiveness of competitive markets as spontaneous, decentralized coordination mechanisms would become his main contribution to economics and these ideas went on to shape the world (Bowles, Kirman, & Sethi, 2017).

Hayek starts off with the observation that modern economies are incredibly complex and information is very decentralized. We live in a world with billions of consumers who all have different preferences for many different goods and services. Determining what the world should produce on any given day to satisfy everyone’s needs given that there are only a limited amount resources available is very hard. How should we decide who produces what, in which quantity and at what price? To make matters worse, much of the information that is needed to make such a calculation is of course not publicly available. The preferences that consumers have, but also the information suppliers possess, like the knowhow and inputs required to efficiently produce an iPhone, are scattered across many individuals. It would take even the most ambitious government agency years to gather only a fraction of the information that is necessary to plan and run a global economy (and the info would already be out of date before it could be used as input for a decision). “If you start from a belief that the most knowledgeable person on Earth does not have even 1 percent of the total knowledge on Earth, that shoots down […] economic central planning” (Sowell, 2017).

For example, the Covid-19 pandemic very rapidly and drastically increased the demand for face masks. In a centrally planned economy, there would exist a yearly quotum for the number of masks that have to be produced via a specified method. It would take a lot of formal communication between the doctors and nurses who notice the mask shortage, the central planner, the manufacturing plant and the suppliers of the plant to tackle the shortage. After the shortage is over, it would take a similar effort to ensure that resources are not wasted in making face masks nobody is using. In a market economy on the other hand, this communication happens automatically via the price system. When customers start buying large numbers of masks, suppliers of medical equipment increase their prices due to the

[2] This is why Hayek dedicated his most influential work “The Road to Serfdom” to “the socialists of all parties”, i.e. to the central planners of all parties.

higher demand. The higher price is an incentive for the firms that manufacture the masks to ramp up their production. Even firms that make related products could be lured into making face masks now their price is considerably higher. Higher prices also encourage entrepreneurs to find alternative production methods and could be a motivation for other institutions to economize on their face mask use or to shift to substitutes whenever possible. Warehouses that would still have a stock of face masks are incentivized to sell them. All these actions help to alleviate the shortage and are driven solely by the price system. The higher price is a call to action for an entire network of suppliers, consumers, entrepreneurs and warehouse managers. The individuals who are involved do not have to formally coordinate in any way, they do not even have to be aware of why the price increased in the first place. All they have to do is follow their profit incentive and use their local knowledge and skills[3].

In this view, the marvel of markets is that "spontaneous order" can emerge without any conscious central planning and without any omniscient entity. Questions like how many cars Germany should produce never even have to be asked, the market figures it out. If the demand for German cars would increase, prices would rise, which gives car makers the signal and the incentive to expand their production capacity. The same is true for the hundreds of manufacturers that produce the parts and materials that the cars are made of. They are part of a global supply chain that is completely coordinated by the price system. The economy runs itself.

This is Hayek's crucial insight into the functioning of markets. For him, markets are not merely platforms for the exchange of goods and services, but also powerful information systems. By buying and selling, households and companies send signals to the market. Markets absorb this information into their prices and these market prices in turn transmit signals to other market participants, like an advanced "telecommunications" system (Hayek, 1945, p. 527). A market economy runs much more smoothly than a centrally planned one as it makes much better use of dispersed information.

Take the stock price of Walmart as another example. Valuing such a behemoth that has thousands of suppliers, sells hundreds of thousands of products and employs millions of employees is extremely hard. It would take a government department months to gather the relevant information and come up with a reasonable static estimate. Yet, stock markets all around the world value thousands such companies every trading day, in real time. These prices are formed by a large number of individuals who each trade on what they know and they contribute their local information to the market in the process. In this perspective, markets are incredible tools for gathering and transmitting information, allocating resources and coordinating our economic activity. "I am convinced that if it were the result of deliberate

[3] The idea that markets can work very well even though individuals possess only local knowledge is sometimes called the "Hayek Hypothesis" (Forsythe, Nelson, Neumann, & Wright, 1992; Hurley & McDonough, 1995; V. L. Smith, 1982b).

human design […] this mechanism would have been acclaimed as one of the greatest triumphs of the human mind" (Hayek, 1945, p. 527).

To illustrate this point further, there are many rough but alluring analogies of self-organizing systems in the natural world. Beehives, ant colonies or fish schools all operate mesmerizingly efficiently without any central authority. Economists sometimes like to compare them to the market system when they get poetic (Surowiecki, 2005). A. Smith (1776) used the rather modest metaphor of the invisible hand to characterize the organizing principle of our economic activity. Given the apparent smoothness by which markets solve our extremely complex global economic coordination problems, it could be more appropriate to think in terms of the invisible quantum computer.

Hayek's views on the functioning of markets quickly gained traction within academia as he was very well connected. Hayek for example was the president of the Mont Pelerin Society, a think tank he co-founded with world-class academics including Milton Friedman, George Stigler, Frank Knight and Karl Popper. Today, 9 members of the Mont Pelerin Society have won a Nobel prize. His ideas not only became popular within academia, but also in policy circles. The view of markets as a collective superintelligence has become the ideal ideological ammunition for many generations of free-market politicians. Margaret Thatcher for example famously took out a copy of Hayek's *The constitution of liberty* (1960) during a policy meeting, smashed it on the table and declared "This is what we believe in!" (Blundell, 2008, p. 41). Ronald Reagan, also a Hayek fan, only half tongue-in-cheek claimed that "the nine most terrifying words in the English language are: I'm from the Government, and I'm here to help" (Reagan, 1986, p. 1) when discussing government intervention in markets[4].

The enthusiasm for free markets was further strengthened by the remarkable increase in material well-being in countries like China, Chile and former Soviet Republics that opened up their economies to market forces and lifted billions out of poverty by doing so. Hayek, who was still seen as a "right-wing nut" (Cassidy, 2013, p. 38) in the early eighties, became recognized as one of the most prominent thinkers of the 20th century. According to former Treasury Secretary Larry Summers: "the single most important thing to learn from an economics course today […] is the view that the invisible hand is more powerful than the hidden hand. Things will happen in well-organized efforts without direction, controls, plans. That's the consensus among economists. That's the Hayek legacy." (Yergin & Stanislaw, 1998, pp. 150-151).

The belief that the economy best runs itself led i.a. to radical deregulation of financial markets starting in the 1980s (Crotty, 2013). If markets are indeed a collective superintelligence, why

[4] Although Hayek is often quoted by free-market, small government enthusiasts, Hayek himself was not against government intervention per se. "In no system that could be rationally defended would the state just do nothing" (Hayek & Caldwell, 2014, p. 88). Hayek sees an important role for the government to construct an intelligently designed legal system which allows competition, to internalize externalities and to provide social services and limit working hours.

should we constrain them in any way? Gerard Debreu, a Nobel Prize winning economist, even declared that “the superiority of the liberal economy is incontestable and can be mathematically demonstrated” (Debreu, 1984). There was no lack of self-confidence at the highest levels of academia.

After the financial crisis of 2007-2008, the faith in free-market capitalism took a big hit. While chasing profits, deregulated markets grossly underestimated risks and brought the global financial system to the brink of implosion. In the US alone, 707 banks had to be bailed out and many others went bankrupt. The crisis led to nearly 10 million Americans losing their homes and the shockwaves were felt across the entire world (Andres, 2018). In Europe, the global financial crisis set the stage for the European debt crisis. Critical economists accused their colleagues to provide models that made policymakers believe markets could be self-regulated (Stiglitz, 2010). Even Alan Greenspan, one of the high priests of the deregulation movement admitted being in “a state of shocked disbelieve” (*The financial crisis and the role of federal regulators*, 2008) when the global financial crisis exposed large fault lines in his free-market ideology. Blindly outsourcing important decisions to market forces turned out not to work as smoothly as hoped. In the words of Nobel laureate Stiglitz (2002): “Adam Smith's invisible hand […] is invisible, at least in part, because it is not there”.

There is an old joke amongst economists about two dignitaries who are attending a military parade (Appelbaum, 2019). At the end of the parade, after the soldiers, tanks and missiles have passed, a small truck appears with a few chubby nerds in the back. “Who are they?”, the first dignitary asks the second? “Oh”, says the second, “they are economists, you wouldn’t believe the damage they can cause”.

A key problem with the Hayekian framework is that it could easily lead to detrimental feedback loops if people only rely on price signals to learn about the private information of others (Bikhchandani, Hirshleifer, & Welch, 1992; Bowles et al., 2017). Suppose for example that the stock price of Tesla starts rising rapidly. Investors who regard stock prices as informative could infer that there are others with private information about the growth of the industry who started buying the stock which pushes the price up. When information is costly, i.e. when it is hard to check whether the price increase was indeed driven by fundamental factors, it could be rational for investors just to assume it was and to start buying Tesla shares as well, which pushes up the price even further and can lead to a bubble. Similar herding behavior can for example also be seen when tourists choose a restaurant in a foreign city. When they do not know anything about the quality of the restaurants in the neighborhood, it is rational to follow the crowd and to go to the busiest restaurant around. The customers who are already present at the restaurant supposedly made their choice based on their private information. However, if everyone just follows the actions of others, we can quickly all end up at the wrong restaurant, or royally overfund certain industries (Bikhchandani, Hirshleifer, & Welch, 1998). To make an analogy to the natural world again, Hirshleifer (2020) compares such a process to the “death spiral” in ant colonies. These colonies organize themselves by a simple rule: each ant follows the ant in front of them. This goes remarkably well until an ant for some reason links up with an ant in the back. At this point, the pack starts running in a circle and the ants keep on going until they die of exhaustion.

The above discussion stresses that the way we think about markets is never inconsequential. Ideas of academics inspire politicians and get picked up by journalists and opinion leaders. Such ideas can have the power to shape our world, as already highlighted by the quote from Keynes at the beginning of this preface. Up to today, the belief that we best organize our society via markets is deeply entrenched in our political and economic conversation. Probably rightly so, as capitalism has proven over and over again to be a great engine for growth and prosperity. Neoliberal metaphors and ideas like the invisible hand, small government and *laissez-faire* have become fundamental cultural symbols (Abel & Kunz, 2018, p. 37).

However, as the global financial crisis and the eurozone crisis have shown, the debate on the merits of free market capitalism and its limits stays as relevant as ever. Certainly now more and more countries around the world are moving towards a free-market society (Lawson, 2019). Moreover, many countries also experiment with new asset markets for policy purposes. An important example is the EU-ETS, the world's biggest carbon market aimed at reducing greenhouse gasses via the trade of emission allowances (European Commission, 2021). Such an initiative is Hayekian in the sense that the carbon prices incentivize individual companies to leverage their local know-how to decrease their carbon footprint cost-effectively. Another example is the SCORE program run by DARPA (Defense Advanced Research Projects Agency), the R&D unit of the United States Department of Defense. SCORE stands for 'Systematizing Confidence in Open Research and Evidence' and the goal is to determine which scientific conclusions are replicable (Alipourfard et al., 2021). The program is a response to the 'replication crisis', i.e. the observation that the results of many empirical studies cannot be confirmed by other researchers. To evaluate the credibility of papers, DARPA runs a prediction market where traders can bet on whether papers are replicable or not. The project is a direct application of the information aggregation mechanism of markets. By pooling the predictions of many individuals together, the market outcomes will presumably be very accurate and be a good gauge of the robustness of the research.

All in all, it is crystal clear that our understanding of how markets work has greatly shaped our world in the past and will continue to do so in the future. The debate on the pros and cons of a market economy is vast. It includes many valid questions on productive efficiency, on inequality, economic development, externalities and so on. To limit our scope, we focus on an idea that is fundamental to our faith in markets: the efficient market hypothesis. To further narrow down, we study this hypothesis in prediction markets. These are platforms where assets are traded that have payoffs which are linked to event outcomes like for example a democrat winning the presidential election. As we discuss later, these markets turn out to be especially interesting research labs for testing the efficient market hypothesis.

Efficient markets

Fama (1970) operationalized the Hayekian idea that market prices are informative via the efficient market hypothesis. The efficient market hypothesis, in its most extreme form, states that market prices fully reflect all available information. Take the stock price of a pharmaceutical company for example. The efficient market hypothesis stipulates that everything there is to know about the company like the number of products it sells per year, its market share, the patents the company owns, the value of its office buildings, its inventories, projections about the growth of the company and its competitors, the probability that it will develop a new groundbreaking drug... will all be reflected in the stock price in an unbiased way based on the information set such that the company is "correctly" priced. Fundamentally, the efficient market hypothesis is an application of competitive equilibrium theory to asset prices (LeRoy, 1989). Traders chasing profits will keep on trading on information, as if led by an invisible hand Adam Smith would note, until the information is embedded into the prices and the market is efficient.

To the uninitiated, the efficient market hypothesis often sounds outrageously radical as financial markets are often depicted as irrational rollercoasters that seem far removed from any notion of efficiency (Malkiel, 2003). Keynes (1936, p. 159) for example famously criticized our reliance on financial markets: "When the capital development of a country becomes a by-product of the activities of a casino, the job is likely to be ill-done". However, an enormous empirical literature leads economists to conclude that the efficient market hypothesis is "probably the best-tested proposition in all the social sciences" (Cochrane, 2011a, p. 36) and that "there is no other proposition in economics which has more solid empirical evidence supporting it" (Jensen, 1978, p. 95).

Anecdotal evidence of the brilliance of market forces emerges in work like that of Roll (1984), who shows that prices on futures markets can be used to improve the weather forecasts of the National Weather Service in the US. Roll (1984) specifically looks at orange juice futures, which are contracts designed to buy or sell orange juice somewhere in the future. These contracts are handy for farmers who grow oranges as they can already lock in a price for their products today, even though they will only harvest their crop in a few months. Orange production is especially susceptible to weather conditions. A fierce storm or a very cold winter can easily wipe out a substantial amount of the total supply of oranges. For example, in an especially cold February night in 1895, almost every orange tree in Florida froze to death leading to a staggering 97% decrease in orange production. Such bad weather decreases orange supply which increases the price of oranges. This means that when bad weather is forecasted, individuals will already anticipate that the orange price will rise in the next months, leading to higher prices on futures markets. Alternatively, if Hayek and Fama are right and prices are indeed informative, we could turn the story upside down and look at whether changes in futures prices contain meaningful information about the weather. Indeed, Roll (1984) finds that movements in futures prices are informative and can even be used to improve the weather forecasts of the National Weather Service. There appears to exist a "wisdom of the crowds" where the collective intelligence of traders outperforms a department of weather experts.

Another example is a paper by L. Moore and Juh (2006). They study derivative prices from the early 20th century, which is long before the formal theory for derivative pricing, the Black and Scholes (1973) model, was developed. They show that even without a theoretical benchmark available at the time, the market prices of the derivatives approximated their theoretical prices, which can be computed in hindsight, surprisingly well. Even without formal models, the market had a sophisticated intuition of the forces driving derivative prices.

A last intriguing illustration is the work by Maloney and Mulherin (2003). They investigate the reaction of the stock market to the Challenger disaster, the space shuttle that exploded in 1986, just 73 seconds after liftoff. As a spacecraft is an incredibly complex piece of engineering, it is hard to find out which components malfunctioned, and which suppliers could be responsible. Indeed, it took an expert commission more than 4 months to identify Morton Thiokol, the supplier of the booster rockets, as the culprit. The stock market however seems to have made the right judgement in just minutes as the price and liquidity of Morton Thiokol stock dropped dramatically after the explosion while the stock price of the other suppliers all quickly rebounded after initial drops[5].

At other times however, the collective intelligence of markets seems nowhere to be found. Staying in the space exploration realm, the stock market apparently did not do a good job in assigning responsibility for the Columbia space shuttle crash in 2003 (Surowiecki, 2005). Another good example is the research by Huberman and Regev (2001). They report a 330% rise in the stock price of a biotech company right after the *New York Times* published an article on a recent breakthrough of the firm in developing a cancer drug. This in itself is not surprising of course, we expect stock prices to react rapidly to new information in an efficient market. However, the information was not genuinely new, it had already been published in *Nature* and in popular newspapers months earlier. It appears that the market reacted very strongly and permanently to news that was no news at all, which is hard to fit in the efficient market framework. Or take the work of Hirshleifer and Shumway (2003) who find that sunshine is strongly and significantly correlated with stock returns. This is again hard to rhyme with efficient markets. It could imply a certain arbitrariness of stock prices; markets would allocate resources differently on sunny days compared to cloudy days.

The market efficiency debate has been raging for decades. During this period, most of the empirical work used stock prices. The stock market has been the Large Hadron Collider of finance if you will. This is not surprising given the central role of stock markets in a capitalist

[5] Although leading academics like to cite the Challenger case as an example of market efficiency (see for example Hanson (2013) or Lo (2019)) it is worth pointing out that comparing the speed of the market with that of the expert commission is not totally fair as the later also played a political and legal role. However, this does not take away the fact that the market was very quick to correctly identify the problem.

economy and the availability of high quality, long term stock price datasets. However, testing efficiency in stock markets is a notoriously fishy undertaking. In an ideal world, a researcher can assess market efficiency by comparing market prices of assets with their true values. The problem is that for stocks, true values are never directly observable. Researchers can resort to models that generate theoretical prices and compare these prices to the market prices. However, when discrepancies between the market prices and the theoretical prices arise, it is not clear whether the market prices are wrong or whether the model that generates theoretical prices is wrong. Even when market prices coincide with theoretical prices this is not necessarily good news as there is an off chance that the market and the model are equally wrong. This fundamental untestability of efficiency in financial markets is called the joint hypothesis problem (Campbell, Lo, & MacKinlay, 1997). Relatedly, Summers (1985) quite brilliantly mocked financial economists by accusing them to engage in "ketchup economics". "They have shown that two quart bottles of ketchup invariably sell for twice as much as one quart bottles of ketchup except for deviations traceable to transactions costs […]. Indeed, most ketchup economists regard the efficiency of the ketchup market as the best established fact in empirical economics" (Summers, 1985, p. 634). Although it is a caricature, studying relative prices and arbitrage relations between financial assets is still core to the efficiency debate today.

A potential way out of the joint-hypothesis gridlock is to study market efficiency in a laboratory setting. The American economist Vernon Smith won the Nobel Prize in 2002 for his pioneering work in this field. In such experimental asset markets, participants (often students) buy and sell artificial assets and get a small monetary reward if they perform well. The large advantage of such studies is that the setting can be entirely controlled. The researchers can for example determine what the true value of the asset is or what information the participants receive. As a result, these studies can much more closely investigate to what extent the market price, resulting from the trading between the participants, resembles the true value of the assets. The empirical work shows that market prices indeed often closely approximate their theoretical values (e.g. Forsythe, Palfrey, and Plott (1982)), but also that bubbles and crashes appear to be common (V. L. Smith, Suchanek, & Williams, 1988). The large drawback of experiments is that they are just experiments. There are a lot of external validity concerns which make it hard to determine to what extent the results from these studies are relevant for the real world.

Prediction markets, the focus of this dissertation, could provide another original contribution to the efficiency debate. As noted, these are real world markets where assets are traded whose values depend on the outcomes of uncertain future events like who will win the Super Bowl or which party will win the elections. Instead of focusing on the untestable question of whether the price of a bottle of ketchup is correct, we can unleash the market mechanism on questions that are easily verifiable, like how many bottles of ketchup will be sold in the next quarter. If the market indeed efficiently aggregates information that is scattered across market participants, it should provide meaningful answers to those questions as well.

Moreover, in contrast to experimental asset markets, these markets are populated by real traders with real money at risk and have a long history which means there is a lot of data that is easily available for researchers.

Furthermore, if prediction markets indeed turn out to tap into our collective superintelligence, they can be very useful policy tools. In 2008, 22 top academics including 5 Nobel Prize winners[6] underlined the advantages of prediction markets in a piece titled "The Promise of Prediction Markets" in the leading journal *Science* (Arrow et al., 2008). They urged lawmakers to provide a more lenient legal framework for these platforms as they could improve our decision making in many domains. They point to for example the work by Polgreen, Nelson, Neumann, and Weinstein (2007) who show that prediction markets accurately forecast seasonal influenza by aggregating information that is scattered across physicians, nurses, patients, pharmacists and microbiologists.

The *leitmotif* of this dissertation is that prediction markets are not merely a relatively obscure, standalone research program. Prediction markets have many research-friendly characteristics that can help give new insights into the efficiency debate, and by extension, the broader economics literature. In prediction markets, efficiency can be assessed much more directly than in mainstream financial markets as the insurmountable joint hypothesis problem is easily bypassed because the true values are readily observable. As a result, prediction markets can allow us to gaze straight into the information aggregation mechanism that lies at the heart of financial economics.

[6] Kenneth Arrow, Paul Milgrom, Thomas Schelling, Robert Shiller and Vernon Smith (at the time of writing).

Chapter 1: Definitions, Context and Related Literature

1.1 Definition

In this dissertation, we study prediction markets[7]. Arrow et al. (2008, p. 877) define prediction markets very broadly as: "forums for trading contracts that yield payments based on the outcome of uncertain events". Similarly, Snowberg, Wolfers, and Zitzewitz (2005, p. 367) define prediction market assets as "simply gambles on uncertain future events". Other definitions are more goal oriented. According to Manski (2006, p. 425) "prediction markets are futures markets in which prices are used to predict future events", while Tziralis and Tatsiopoulos (2012) define prediction markets as "markets that are designed and run for the primary purpose of mining and aggregating information scattered among traders and subsequently using this information in the form of market values in order to make predictions about specific future events."

We will define prediction markets as "markets that trade in state-contingent claims of which the states are not related to the performance of financial assets". Our definition is more precise than that of Arrow et al. (2008) or Snowberg et al. (2005), which are so broad they could also include traditional stock options, and we lose the goal-oriented focus of the definitions of Manski (2006) and Tziralis and Tatsiopoulos (2012). By relaxing the purpose of prediction markets, we can include an insightful, substantial and longstanding literature on betting markets, which have assets that are mechanically not different from the assets in prediction markets in the sense of Manski (2006) and Tziralis and Tatsiopoulos (2012)[8].

In their simplest form, prediction market securities have a payoff R_i for outcome i ($i = 1, \dots, n$) which depends on the stochastic variable X such that:

$$R_i(X) = \begin{cases} 1 & X = i \\ 0 & X \neq i, \end{cases} \tag{1}$$

which is just a fancy way of stating that a prediction asset pays out 1 if the underlying event outcome substantiates and zero otherwise. This type of binary options is generally referred to as Arrow-Debreu securities or state-contingent claims in the financial economics literature (Pennacchi, 2008). If the prediction asset has a price p_i, its return is given by:

[7] Other names include information markets, decision markets, event futures or idea futures.

[8] One could argue that betting on sports games is intrinsically enjoyable and therefore different from trading prediction assets. However, nothing prevents individuals from using prediction assets for entertainment purposes either. As noted by Wolfers and Zitzewitz (2004, p. 108) "most contingent commodity markets involve some mix of risk sharing, fun and information transmission, so these distinctions are not impermeable".

$$r_i(X) = \begin{cases} \frac{1}{p_i} - 1 & X = i \\ -1 & X \neq i. \end{cases} \quad (2)$$

The expected return weighs the returns by their probabilities:

$$E[r_i(X)] = \pi_i \left(\frac{1}{p_i} - 1 \right) + (1 - \pi_i)(-1) = \frac{\pi_i}{p_i} - 1 \quad (3)$$

where π_i is the true but generally unknown probability of outcome i. Note that this last expression has a very intuitive interpretation: as long as the market price coincides with the true probability, the expected return on the asset is zero (in absence of transaction costs).

In the sports betting literature, which we treat as a natural subset of the prediction market literature, potential gains are often expressed as odds. In this dissertation, we adhere to the decimal odds convention. These odds, o_i, represent the potential payoff per unit bet. This is conveniently the reciprocal of the price, which is the amount paid per unit of potential payoff: $o_i = 1/p_i$. If we express prices as odds, the expected return becomes:

$$E[r_i(X)] = \pi_i o_i - 1. \quad (4)$$

1.2 Theoretical legitimation

The underlying idea that gets economists excited about these markets is that market prices are supposed to be informative in the Hayekian sense. Moreover, in an efficient market, prediction markets would cleverly aggregate all information that exists about a certain event. Take the US presidential elections for example. There are many individuals who can make an educated guess of who is going to win (campaign staffers, journalists, opinion pollsters, pundits, voters…). They all possess some information on for example local voter behavior, demographics, the economic situation of the swing states, insights into campaign finances and super PAC funding, conversations with focus groups of voters, sentiment during rallies… In a prediction market, everyone who holds information gets an incentive to trade on the information until it is fully reflected in the price. If the market is indeed efficient and we can assume risk neutrality (which is often plausible given the small amount of money individuals trade on prediction markets)[9], the market price is our best guess of the event outcome

[9] Gürkaynak, Wolfers, Carroll, and Szeidl (2005) confirm that risk premia are so small in the prediction market context that the can be ignored. In the empirical work of this dissertation, we focus on

(Snowberg, Wolfers, & Zitzewitz, 2013). Furthermore, no other available information, like opinion polls, or a combination thereof can improve this forecast. Although market efficiency and risk neutrality are strong assumptions that will never be perfectly true, much of the empirical work is consistent with these assumptions.

As prediction markets incentivize individuals to bring information to the market and exploit mispricing, we can reasonably assume that they are not predictably wrong:

$$E(outcome_{t+1} - forecast_t \mid information_t) = 0, \tag{5}$$

where the expected market forecast error, conditional on the information set at time t, is zero. Applied to our prediction markets we have:

$$E\left(R_{i,t+1} - p_{i,t} \mid \Omega_t\right) = 0, \tag{6}$$

where $R_{i,t+1}$ is the payoff of the asset on outcome i at $t+1$, $p_{i,t}$ is the price of the asset at time t and Ω_t is the total information set at time t. In linear representation we have:

$$R_{i,t+1} = \alpha + \beta \times p_{i,t} + \varepsilon_i, \tag{7}$$

where we can test the joint null hypothesis that $\alpha = 0$ and $\beta = 1$, ε_i is an error term, uncorrelated with the forecasts and all other variables in the information set. Alternatively, we can specify

$$R_{i,t+1} - p_{i,t} = \alpha + \beta \times X_t + \varepsilon_i, \tag{8}$$

where we regress the forecast error on any available information X_t from the information set at time t and test the joint null that $\alpha = 0$ and $\beta = 0$. This last regression is a generalization of (7) where the coefficient of the market forecast is restricted to 1. This setup is the prediction market equivalent of the efficiency and rational expectations tests often carried out in mainstream economics. Examples there include regressing future spot prices on forward prices to test if forward prices are unbiased predictors (while assuming no risk premium) or testing inflation expectations (see for example Levich (1989) and Andolfatto, Hendry, and Moran (2008) respectively).

Testing whether the forecast error is unbiased is equivalent to testing whether expected returns are zero. We can scale (6) by the price (given that the price is not zero) such that:

$$E\left(\frac{R_{i,t+1} - p_{i,t}}{p_{i,t}} \mid \Omega_t\right) = E(r_{i,t+1} \mid \Omega_t) = 0. \tag{9}$$

contracts on sports games where this assumption is realistic given the short-term nature of the contracts and the lack of non-diversifiable risk. For other contracts like assets on a presidential election, the assumption is less straightforward, but nevertheless consistently accepted in the literature.

Although the regressions provide a framework for understanding and testing the unbiasedness of market forecasts, they do not provide a priori indications that this would be the case. In what follows in this section, we discuss a theoretical model that, under respective assumptions, suggests a direct link between market forecasts and outcomes.

As noted earlier, prediction assets are essentially Arrow-Debreu securities, i.e. state contingent claims that have a payoff of 1 if the state substantiates and zero otherwise. Suppose we have a world where only two outcomes are possible: $i = 1, 2$ (for example for a football game where $i = 1$ indicates the home team wins the final game of a tournament, $i = 2$ means the away team wins). The payoff vectors are trivial: (1,0) and (0,1) for the prediction asset on outcome 1 and 2 respectively. Starting from the primal asset pricing equation, which goes back to at least Beja (1971), we can quite easily show that under no arbitrage, the equilibrium forecast will indeed be the expected outcome. If m_{t+1} is the stochastic discount factor, we have:

$$p_{i,t} = E\left(m_{t+1} R_{i,t+1} | \, \Omega_t\right). \qquad (10)$$

We can reasonably assume that the stochastic discount factor, which represents both the time value of money and the systematic risk, equals one. We can write

$$p_{i,t} = E\left(R_{i,t+1} | \, \Omega_t\right) = \pi_{i,t}, \qquad (11)$$

where $\pi_{i,t}$ is the probability that state i substantiates. Consequently, there is a direct link between equilibrium prices and probabilities.

1.3 Efficiency in prediction markets

In the prediction market literature, two main efficiency benchmarks are commonly used (Thaler & Ziemba, 1988).

1) Positive expected returns do not exist:

$$E(r_{i,t+1} | \, \Omega_t) \leq 0 \;\forall i. \qquad (12)$$

This is the most straightforward definition and can be tested either in the aggregate by checking that expected returns after transaction costs[10], conditional on some information set, are smaller or equal to zero, or in more specific cases, by checking arbitrage opportunities are not possible.

[10] In the discussion above, transaction costs were assumed to be zero. In reality, platforms charge a fee for their services. More information on these transaction costs follows in the section on market microstructures.

2) Expected returns are equal to the negative commission:

$$E\left(r_{i,t+1} | \Omega_t\right) = E\left(r_{j,t+1} | \Omega_t\right) = -c \ \leq 0 \ \forall i,j. \tag{13}$$

Even if positive expected returns would not be possible, some gamblers could still be able to beat the average trader consistently when trading on some information set (i.e. have less negative returns). In this second and stronger efficiency definition, such abnormal profits are not possible, all expected losses are equal to the commission charged by the platform.

Much of the empirical work is organized around investigating the predictive power of different information sets Ω_t. Popular choices for Ω_t include for example current and past prices, expert opinions or other variables related to the underlying event (like for example team quality in sports prediction markets).

1.4 Taxonomy of prediction asset contracts

As noted above, Arrow-Debreu type prediction assets can reveal the probability of an uncertain outcome. This type of contract is studied in chapters 3 and 4 of this dissertation. However, as highlighted in Table 1, by changing the contract design, prediction assets can reveal other market expectations.

Table 1: taxonomy of contracts, adapted from Wolfers and Zitzewitz (2004).

Contract	**Example**	**Payoffs**	**Reveals market expectation of**
Arrow-Debreu (sometimes also called "winner-take-all" contracts".)	Outcome i: Real Madrid wins the champions league.	Price is p_i and the contract pays 1 if outcome i occurs, 0 otherwise. (In sports betting, the prices are often expressed in terms of odds, which are the inverses of the prices.)	Probability that outcome i occurs: π_i
Spread	The Green Bay Packers defeat the New England Patriots by more than y^* points.	Contract costs 1 and pays 2 if $Y > y^*$, 0 otherwise.	Median value of Y
Index	Joe Biden will win y^*% of the popular vote	Contract pays y^*	Mean value of Y

A second popular contract type are spread bets, which are studied in chapter 2 of this dissertation. In spread contracts, agents bet on whether a stochastic variable will be higher or lower than a market established benchmark. For example, whether a presidential candidate will win more than y^*% of the popular vote. For these contracts, the prices and payoffs are fixed, but the benchmark y^* is adjusted by the trading that takes place. When the contract costs 1 and pays off 2 to winners and 0 to losers, y^* will reflect the median value of Y. This because the contract is only fair in this setup whenever the probability of winning the bet equals the probability of losing the bet.

A last common contract type are index bets. Such an asset pays for example 1 for every 10000 unemployed Americans in the next jobs report. The price for an index asset will reflect the expected value of the underlying variable.

Note that by combining multiple contracts on the same underlying variable, prediction markets can be used to estimate its probability distribution (Wolfers & Zitzewitz, 2004). For example, if we want to know the probability distribution of the numbers of goals Real Madrid will score in their next game, we could set up a market in which a series of Arrow-Debreu assets are traded where the underlying events are the number of goals scored (e.g. contract 1: Real Madrid will score 0 goals, contract 2: Real Madrid will score 1 goal…).

Similarly, suppose we want to know the 60th percentile of the distribution of next quarter's GDP growth. We could design a contract that costs 60 and pays 100 whenever $Y < y^*$. This contract is indeed only fair if y^* is the 60th percentile as the expected return is then $-60 + 100 \times p(Y < y^*) = 0$.

In the context of index contracts, we could set up a market where the payoff is a function of next quarter's GDP growth and secondly, a market where the payoff is the square of the growth. Market prices will then estimate both $E(Y)$ and $E(Y^2)$ from which we can compute the variance of Y as $Var(Y) = E(Y^2) - E(Y)^2$. Higher order moments like skewness or kurtosis could similarly be computed by adding even higher order contracts.

Another way to extend the informational content of prediction markets is by trading assets that are conditional upon a second event. For example, consider a contract that pays 1 if the Democrats win the next presidential elections and zero otherwise. As discussed above, the market price of this contract reveals $Pr(Democrats\ win)$. While insightful, it does not directly guide the party delegates in their choice of the candidate with the highest win probability. In this context, it can be interesting to design a second contract that also pays 1 if the Democrats win, but only when Joe Biden is the Democratic candidate. The contract is called off when Joe Biden does not become the Democratic candidate and all traders are refunded. Such a contract reveals the probability of the Democrats winning, conditional on Joe Biden being the candidate: $Pr(Democrats\ win | Joe\ Biden\ candidacy)$. Such contracts can be useful to evaluate policy choices in general (Hanson, 2013). For example, to estimate the effects of a tax decrease, a conditional contract could be traded on the question "What will next year's GDP growth be given that the tax decrease is implemented?" and be compared with the unconditional GDP growth forecast.

1.5 Taxonomy of market microstructures

Prediction markets exist in diverse environments with distinct market microstructures. While the multiplicity of microstructures could seem redundant at first, it is an advantage for researchers. Each system has its own characteristics which allows for different questions to be addressed. In this section we introduce some of the most common systems and highlight their advantages from a research perspective. In our taxonomy we first distinguish between market marker, and non-market maker systems before narrowing down further.

1.5.1 Non-market maker systems

In the non-market maker systems, platform operators play a passive role. That is, they facilitate price discovery and trading by matching buyers with sellers, but do not interfere with the pricing process or become a counterparty in a trade. In these systems, platform operators do not take on any event risk, i.e. their profits are not dependent on the event outcomes.

1.5.1.1 Exchanges

Like traditional financial assets, many prediction assets are traded on exchanges that operate via a continuous double auction. On these exchanges buyers and sellers can place limit orders or accept limit orders that have been placed. The prices shown to traders are the highest bid price and the lowest ask price. Exchanges have a few advantages over bookmaker systems. First, as traders on exchanges effectively trade against each other, they can just as easily long or short outcomes. Second, they can unwind their position at any time at the respective market price. And third, exchanges do not skew prices or odds in their favor, as bookmakers do, but charge a commission on wins. This commission is generally smaller than the (implicit) commission paid via bookmakers.

Some notable non-profit prediction market exchanges are Iowa Electronic Markets, operated by the University of Iowa, PredictIt, operated by the Victoria University of Wellington New Zealand and the Sauder School of Business Prediction Markets affiliated with the University of British Columbia. These platforms are run for research and educational purposes and mainly offer prediction assets on political events. Large for-profit exchanges include Betfair, Betdaq and Smarkets. These platforms are mainly focused on sports prediction assets, but also offer a wide variety of political, and other prediction assets. For example, during the COVID-19 lockdown, assets were traded where the underlying was related to the date English schools would reopen. Another interesting example is the Hollywood Stock Exchange, which offers prediction assets related to the movie industry like who is going to win an Oscar and what the box office of a certain movie will be.

1.5.1.2 Pari-mutuel betting (pool betting)

Pari-mutuel betting is mostly associated with horse racing[11]. Per race, the operator pools all the money bet and deducts a percentage, the track take[12]. This track take covers expenses, taxes and profits and can be quite substantial. To give an indication, bets organized by the Hong Kong Jockey Club have a track take of 17.5% for the most common types of bets, and a take of 25% for more exotic bets as of 2020 (Hong Kong Jockey Club, 2020). After the race is run, the pool minus track take is distributed among the winners in proportion to their bets. In this system, the odds o_i are entirely determined by the relative proportions bet on each horse:[13]

$$o_i = \frac{B(1-t) - b_i}{b_i} + 1, \tag{14}$$

where B is the total amount bet on all outcomes, t is the track take and b_i is the amount bet on outcome i. The market implied probabilities Q_i are given by:

$$Q_i = \frac{b_i}{B}. \tag{15}$$

By substituting (13) into (12) we have:

$$o_i = \frac{(1-t) - Q_i}{Q_i} + 1. \tag{16}$$

We can rewrite this last equation as:

$$Q_i o_i - 1 = -t. \tag{17}$$

Notice that if the true probability π_i equals the market implied probability Q_i, the left-hand side of the above equation is just the expected return definition (4) and thus

$$E[r_i(X)] = -t. \tag{18}$$

We again reach the general conclusion that positive expected returns are only possible when the market implied probabilities are not equal to the true probabilities, i.e. the market misestimates the true probabilities. If these probabilities are equal, traders expect to lose the track take (Bolton & Chapman, 1986).

---

[11] Some prediction markets on macroeconomic variables are also run via this system (Gürkaynak et al., 2005).

[12] There is another, more implicit transaction costs gamblers pay: the breakage, meaning that the odds are rounded down to the nearest 5 or 10 cents.

[13] To keep consistent notation, we depart from the convention to quote pari-mutuel odds as fractional odds: $o_i^{pari-mutuel} = o_i - 1$.

In pari-mutuel betting, the odds change mechanically whenever new bets are made (see (14)). This means that the final odds are not known until gambling has ceased, which implies gamblers are never certain at which odds they are making their bets. To accommodate the desire of the gamblers to have at least a ballpark figure of the final odds, the current odds are continuously computed and shown on the 'tote board'. The total betting window usually lasts about 20 to 30 minutes, which is much shorter than the trading windows in prediction markets in general.

Often, different types of bets are offered on horse tracks including win, place and show bets. To place a win bet, a gambler selects a single horse and only gets a payout if the selected horse wins. A place bet pays off when the selected horse finishes first or second, and a show bet pays off when the selected horse finishes in the top three. These bets are kept in separate pools creating a distinct win, place and show pool respectively. Additionally, "exotic" bets are often offered, which depend on the performance of multiple horses. An exacta for example pays off when the first and second horse, in correct order, are selected. A quinella is a similar bet, but here the order is not important. Other examples include trifectas, superfectas and daily doubles. As the bets are kept in separate pools per type, arbitrage opportunities can arise when combining bets from different pools. (We will not go into this literature in this dissertation, we refer interested readers to e.g. Asch, Malkiel, and Quandt (1984), Edelman and O'Brian (2004), Gramm, McKinney, and Owens (2012), Hausch and Ziemba (1990b) and Willis (1964)).

1.5.2 Market maker systems

Market makers, in contrast to non-market makers, actively set the conditions on which agents can trade[14] and become the counterparty for all traders. As a result, they can take on event risk. Market makers can suffer large losses when they misestimate the probabilities of the underlying events.

There is quite a bit of debate on the objective of market makers in prediction markets. Traditionally market makers were understood to be balancing their books. By doing so they make sure that the event risk is virtually eliminated so they can profit from the commission they charge. However, more recently, this interpretation of market marker behavior is questioned. Levitt (2004) for example finds evidence that bookmakers do not try to balance the books for every individual game. Rather, they earn a profit in the long run as they are better at predicting outcomes than the average trader. Furthermore, it appears that market makers systematically misprice prediction assets to exploit behavioral biases of traders. By doing so they increase their profits by 20% to 30% over a price-setting policy that balances supply with demand. This strategic market maker behavior can be inconvenient from a research perspective. We only indirectly observe the behavior and expectations of the market

[14] Jaffe and Winkler (1976) discuss the similarities between market makers in financial and prediction markets and their relationship with investors.

participants via possibly skewed unilaterally determined market maker prices. However, comfortingly, competitive forces and smart traders will prevent market makers to skew prices too heavily. Furthermore, the strategic behavior of the market makers is a fruitful topic of research in itself. It can teach us more about how market makers operate and interact with one another, but it can also give insight into how participants judge uncertain outcomes and for which outcomes they are systematically overpaying. In this sense strategic market maker behavior is a mold of the behavioral biases of market participants.

Like market makers in financial markets, market makers in prediction markets pool their exposures in some cases to hedge event risk. For example, suppose a British and a Spanish bookmaker are offering bets on a game between Manchester United and Real Madrid. If gamblers support their home teams the British (Spanish) bookmaker will receive relatively too many bets on Manchester United (Real Madrid). Instead of the bookmakers taking a position in the game, they can lay off their excessive bets with each other (Strumpf, 2003).

1.5.2.1 Fixed odds betting

In fixed odds betting, market makers specify odds, o_i, on event outcomes i at which they accept bets. In this system bookmakers charge an implicit commission by skewing the odds in their favor. This is a disadvantage from a research perspective as all prices are distorted by construction. Suppose that π_i is the objective probability of outcome i. Fair odds o_i^* would be

$$o_i^* = \frac{1}{\pi_i}. \tag{19}$$

However, in order to make a profit, bookmakers subtract a commission τ_i from the fair odds. This distorts the link between the odds and the probabilities:

$$o_i = o_i^* - \tau_i < \frac{1}{\pi_i}. \tag{20}$$

As τ_i is generally unknown, it is not directly possible to derive the true probabilities from the offered odds and the commission charged on an individual outcome cannot be directly observed. In the literature, the commission over all outcomes is commonly defined as:

$$c = \sum_i {o_i}^{-1} - 1, \tag{21}$$

which is conveniently computable from the posted odds. This commission differs drastically between bookmakers, seasons and events as shown in chapter 3 of this dissertation.

Researchers often want to deduct outcome probabilities from which the skew imposed by the commission is removed. These market implied probabilities Q_i are often approximated via:

$$Q_i = \frac{{o_i}^{-1}}{\sum_i {o_i}^{-1}}, \tag{22}$$

where we make the implicit assumption that the commission is spread out in proportion to the true probabilities across outcomes.

Bookmakers typically announce their odds a few days before the event. Posted odds can change over time because of i.a. event-related news. Whenever a gambler makes a bet however, the current odds are locked in. Subsequent odds changes only affect the gamblers who enter later. In contrast to pari-mutuel betting, a gambler knows all the conditions of the bet when it is made[15].

1.5.2.2 Spread betting

In spread betting on sports events, traders do not simply bet on which team is going to win or lose like in fixed odds betting. In spread betting, traders bet on whether a team is going to win by more, or lose by less, than the point spread, PS_i. The point spread is "the market's expectation of the number of points by which the favorite will outscore the underdog."(Golec & Tamarkin, 1991, p. 313)

Spread betting is most commonly used in American football and basketball betting. Bookmakers set this point spread in proportion to the relative team qualities and it is announced a few days before the game. The goal of the spread is to equalize the probabilities of both teams to be involved in a winning bet. Suppose a very strong team is playing against a very weak team. A simple bet on which team will win the game will be heavily in favor of the stronger team. However, with a spread, the bookmaker can level the playing field by requiring not only that the stronger team wins, but that it wins by for example at least a 14-point difference. In the case that the actual difference in points is equal to the point spread, i.e. a "push", the wager is refunded. Bookmakers can remove the possibility of a push by working with non-integer point spreads.

If the spread indeed equalizes the win probabilities of bets on either team, the fair odds would be 2. However, payout happens according to the 11 for 10 rule. This means that an $11 winning bet only yields a profit of $10. This is below the fair payout, which allows the bookmaker to make a profit. If the wagers are perfectly balanced between the two teams, the bookmaker pays out $21 for every $22 it receives. Because of the 11 for 10 rule, a gambler who wants to break even has to win 52.4% of his bets. This can be seen by solving

$$f \times 10 - (1 - f) \times 11 = 0 \tag{23}$$

to f, the fraction of winning bets. The rule can also be interpreted as paying a $1 fee to make a bet at fair odds.

The spread, or handicap, can change over time because of i.a. game-related news or large volumes placed on one of the teams. However, whenever a gambler makes a bet, the current

[15] A notable exception to this is betting with a market maker at "starting price", which are the odds at which a sizable bet could have been made with a bookmaker just before the end of betting. For a discussion see Dowie (1976).

point spread is locked in. Subsequent spread changes only affect the gamblers who enter later. In contrast to pari-mutuel betting, a gambler knows all the conditions of the bet when it is made.

The popularity of spread betting is sometimes explained by the increased thrill of betting on a score difference compared to betting on the outcome. Alternatively, under some circumstances it could be more profitable for a bookmaker to offer spread bets than to offer fixed odds bets (Bassett Jr, 1981; Woodland & Woodland, 1991).

Spread betting brings about the methodological advantage that the probability of winning a bet can be modelled via a binomial distribution where the success rate is 50%. In large samples, the binomial distribution can be conveniently approximated by the normal distribution, which is useful to evaluate the statistical performance of trading strategies. A related advantage is that all bets have the same risk and return characteristics. This eliminates the need to make assumptions about the shape of the utility curve in market efficiency tests (Dana & Knetter, 1994).

1.6 Prediction markets for research purposes

Prediction markets are alluring research labs for many fields including economics, psychology and political science. In this section we summarize the main advantages of prediction markets from a financial economics perspective but many of the advantages spill over to other domains as well. To make the discussion more digestible, the advantages are grouped into three categories: products, market & pricing, and research.

1.6.1 Product advantages

1.6.1.1 Elementary assets

Prediction assets are typically simple binary options that have a positive payoff if the underlying event takes place and zero otherwise. This payoff structure is very easy to understand for all parties involved which potentially eases efficiency. In a lab setting Carlin, Kogan, and Lowery (2013) for example show that lower asset complexity leads to higher efficiency. Furthermore, these assets are essentially Arrow-Debreu securities, which are the building blocks of asset pricing, solidifying a strong theoretical link (Pennacchi, 2008).

1.6.1.2 True values are exogenous and revealed

As discussed in the preface, a major advantage of prediction markets is the exogenous revelation of the true values. The event outcomes are known ex post and are independent of the behavior of traders. This circumvents the dreaded joint hypothesis problem as researchers can systematically compare market prices with the true values. “Economists have given great attention to stock markets in their efforts to test the concept of market efficiency, yet wagering markets are, in one key respect, better suited for testing efficiency and rationality. The advantage of wagering markets is that each asset (bet) has a well-defined termination point at which its value becomes certain” (Thaler & Ziemba, 1988, pp. 161-162).

1.6.1.3 Short maturities

Prediction assets typically have short maturities. Tetlock (2004) finds that 58.4% of contracts traded on a major prediction market platform last less than a day while only 2.2% last more than 5 months. This relatively short time span allows individuals to quickly evaluate their investment decisions and could enhance learning (Thaler & Ziemba, 1988). In experimental research for example, Forsythe et al. (1982) stress the importance of replication for asset prices to converge to a rational expectations equilibrium. Furthermore, short maturities largely remove the necessity to incorporate the time value of money in the analyses. Lastly, the short duration of the contracts might increase arbitrage forces on prediction markets (Tetlock, 2004). In traditional financial markets, noise trader risk can discourage arbitrageurs as asset prices could diverge from their fundamental values for a long time (De Long, Shleifer, Summers, & Waldmann, 1990). In contrast, in prediction markets assets quickly reach their fundamental values which can embolden arbitrageurs to take more aggressive positions.

1.6.1.4 Homogeneous products

Prediction market assets on the same event but offered by different suppliers are virtually perfectly homogeneous. For example, a gambler who wants to make a bet on a game between Liverpool and Manchester United can choose between many different online bookmakers. All the offerings are virtually identical except for the product price. This is useful for industrial organization studies as there is no need to control for differences in product features or quality.

1.6.2 Market & pricing advantages

1.6.2.1 Resemblance to traditional financial markets

Prediction markets and traditional financial markets are both competitive speculative markets. In both markets, a large number of participants collectively determine prices, and therefore also the returns, of assets who's future payoffs are uncertain[16] (Ali, 1979). In both markets, the participants are incentivized by profit opportunities to gather information. Trough trading, this information is reflected into prices. In both markets, trades are subject to transaction costs, participants can easily enter and can get professional advice. Prediction markets, like derivatives trading and active asset management, are zero sum in nature (before commission) (Levitt, 2004).

[16] Note that we mean uncertainty in the Knightian (1921) sense, i.e. the probability distributions are unknown. Because of this, prediction markets and financial markets are categorically different from games of chance like roulette where only risk is involved, i.e. the outcome is not known, but the probability distribution is. "Unlike a casino gambler, a good horseplayer does not count only on luck. There is a reward to gathering information to improve one's probability estimates, and the bettors with superior information and ability to analyze it will be more successful than the rest." (Figlewski, 1979, p. 78)

1.6.2.2 The pricing of prediction assets is idiosyncratic

Prices of prediction assets do not comove with aggregate risks, they are idiosyncratic (Moskowitz, 2021). There is for example no market factor driving the pricing of all bets on football games played on May 3rd nor is there a correlation between the pricing of two different horse races[17] (Snyder, 1978). There are no systematic risk premia as there is supposedly no correlation between marginal utility of wealth and the payoffs of the assets[18]. This element, together with the revelation of true values allows researchers to separate between rational and behavioral asset pricing theories that could be simultaneously at work in traditional financial markets. Without risk premia, market prices can be directly compared to the future value of the asset at termination and deviations can be attributed to behavioral biases. Relatedly, the sums that are traded on prediction markets are often small, which makes it reasonable to argue investors are risk neutral to the idiosyncratic risk involved (Wolfers & Zitzewitz, 2004).

Asset prices should equal the expected discounted cashflows. Much of the contemporary finance research focusses on the "discounted" part (Cochrane, 2011b). However, as there is supposedly no discounting needed for prediction assets (no time value of money and no risk premia), these assets allow us to focus directly on the "expected" part.

1.6.2.3 Relatively small information set

A prediction market is a "market-in-miniature" (Hausch & Ziemba, 1990a, p. 61). The information set applicable to the pricing of prediction assets is much more limited compared to that of traditional assets. For example, the information set that is relevant to estimate the probability that Liverpool will win a soccer game against Real Madrid includes elements like past team performance, individual skills of the players, injuries and weather forecasts. Although it can quickly become a high dimensional dataset, it is multiple orders of magnitude smaller than the information set relevant to pricing the stock of a multinational. This relatively small information set can positively impact prediction market efficiency as the attention of traders is limited (Hirshleifer, Lim, & Teoh, 2009; Simon, 1971). Relatedly, the outcomes of the underlying events are more stable and predictable compared to the events on traditional financial markets (Benter, 2004). For example, there are thousands of soccer games played every week so a very good understanding of the probability distributions of the possible outcomes can be developed and modeled statistically. Lastly, because the information set is relatively small, changes in the information set are often clear and transparent. Furthermore, as much of the info is dispersed publicly, for example a presidential candidate blundering on live television, the changes are not subject to information leakage. This allows us to precisely study the market reaction to changes in the information set. For example, Croxson and Reade

---

[17] Snyder (1978) acknowledges these insights were suggested by a referee, long live the anonymous referees.

[18] Again, this is most clear for sports bets, however it is commonly assumed to hold for other prediction assets as well in the literature.

(2014) study how market prices on a betting exchange move when goals are scored during football games. They find that the odds update swiftly and fully to reflect the new information.

1.6.3 Research advantages

1.6.3.1 Out of sample

A large proportion of the empirical finance literature focusses on stocks, very often even only US equities. As thousands of scholars have been looking into this fairly small dataset for decades, models will be found that fit the data very well, but that do not generalize out of sample (Lo & MacKinlay, 1990). These datamining concerns can be alleviated by out of sample tests. As many prediction assets are not at all related with financial or economic variables, these markets are a fertile environment for testing hypotheses. Take behavioral hypotheses for example. A common criticism to these is that "allowing for irrationality opens a Pandora's box of ad hoc stories that will have little out-of-sample predictive power" (Daniel, Hirshleifer, & Subrahmanyam, 1998, p. 1841). A general behavioral theory should be able to explain cognitive glitches hindering agents from making optimal decisions under uncertainty, irrespective of whether prediction assets or capital market securities are involved. Finding similar patterns in seemingly unconnected markets could expose fundamental behavioral symmetries.

1.6.3.2 Natural field research

There exists a mature strand of literature that tests market performance in lab settings (see for example V. L. Smith (1982a) for a very early overview). However, the gain in controllability that experiments offer is at least partly offset by external validity concerns (Levitt & List, 2007). In prediction markets on the other hand, traders can be observed in their natural habitat, without being aware they are observed and with real money at risk. At the same time, experiment-like characteristics as short maturities, simple assets and known and exogenous true values are maintained. Furthermore, the existence of different prediction market microstructures makes it possible to test a large number of different hypotheses.

1.6.3.3 Data availability[19]

For a wide range of prediction assets, event outcomes and their market prices have been meticulously recorded. This is especially true for sports bets where high quality datasets, often spanning multiple decades and continents, are easily available to researchers. The data availability and high event frequency is the reason that many prediction market researchers

---

[19] Maybe this advantage is more relevant for psychologists and political scientists as financial economists already have access to more quantitative data than is probably good for them. The predictive modeling literature also gratefully uses prediction market datasets to test and machine learning algorithms (Geurkink, Boone, Verstockt, & Bourgois, 2021; Horvat & Job, 2020; Huang & Li, 2021; Hubáček, Šourek, & Železný, 2019; Maymin, 2019; Stübinger, Mangold, & Knoll, 2020).

work with sports bets. A presidential election happens once every 4 or 5 years in most countries, while dozens of sports games take place every weekend.

1.7 Prediction markets for policy purposes

The outcomes of prediction or betting markets have been used to gauge the probability of uncertain outcomes for centuries. Figlewski (1979, p. 77) remarks that betting markets are "dating back thousands of years. (One may wonder whether the comparatively recent development of trading in corporate equity will prove to be as durable an institution.)". In this section, we will sketch an overview of the domains where prediction markets are often used. We will discuss the performance of prediction markets and their potential pitfalls.

Note that prediction markets have a number of characteristics that can make them more attractive compared to other prediction methods like opinion polls or expert opinions. First, individuals on prediction markets are self-selected and incentivized to put their money where their mouth is, which enhances the revelation of truthful information (Hanson, 1999). Traders who make accurate forecasts get compensated while noise traders get penalized. Second, opinions are weighted by the size of the bet giving a louder voice to the individuals who are relatively certain of their forecast. Over time, the relatively successful traders will see their influence increase while unsuccessful individuals see their funds decrease or drop out of the market entirely. Third, prediction markets offer dynamic, real-time forecasts that quickly incorporate new information whereas surveys are taken at a much lower frequency. Fourth, in a prediction market, the entire population of traders does not have to be rational or free of behavioral biases to be efficient. Indeed, Forsythe et al. (1992) for example show that many individuals in the Iowa Electronic Markets suffer from consensus bias, i.e. Democrats (Republicans) overestimate the probability that a Democrat (Republican) will win the election. However, as it is the marginal trader who determines the prices rather than the average trader, market prices can be efficient even though many individuals are biased. Lastly, when forecasting elections, unrepresentative samples would invalidate polls, this is not the case for prediction markets as on these markets, traders are incentivized to make predictions about the entire population. For example, Berg and Rietz (2006) show that an overwhelming majority of traders on the Iowa Electronic Markets are white, college educated males. Nevertheless, this unrepresentative sample of the population succeeds at forecasting elections with a high accuracy.

1.7.1 Sports prediction markets

Sports betting markets are by far the most popular, most liquid and most studied prediction markets. The empirical work dates back to psychologist Griffith (1949) who compared market implied win probabilities (Q_i) from pari-mutuel horse racing odds with their empirical counterparts (π_i). Although the efficient market vocabulary was not yet developed, his approach still reverberates in the contemporary literature. His conclusions were twofold and in hindsight, it is quite remarkable how much these conclusions, from an early five-page paper, align with the conclusions from the rest of the literature. First, in general the market seems to be doing a very good job in assessing the probabilities of the horses winning the race. Second, the occurrence of low probability events (longshots) is overestimated while the occurrence of high probability events is underestimated (favorites). This phenomenon was later called the "favorite-longshot bias" and has become one of the main stylized facts emerging from the prediction market literature (more on this favorite-longshot bias in section 1.8). In related work were Hoerl and Fallin (1974) merely compare the ranking of horses with their true race finish orders they conclude that these are "amazingly consistent" (p. 229) which "demonstrates that individuals with incentive can on the average successfully discriminate small differences in items" (p. 230).

Pankoff (1968) noticed the similarities between stock markets and prediction markets. Inspired by the work of Fama (1965) and Mandelbrot (1966), he explicitly introduced market efficiency concepts in the prediction markets literature. To test for efficiency, he regresses the actual score difference in NFL games on the bookmaker's predicted score differences. The joint null that both the intercept is zero and the coefficient of the predicted scores is 1 cannot be rejected as expected in an efficient market. The stylized fact emerging from the related work is that efficiency can seldomly be clearly rejected (Angelini & De Angelis, 2019; Croxson & Reade, 2014; Elaad, Reade, & Singleton, 2020; Gandar, Zuber, O'Brien, & Russo, 1988; Golec & Tamarkin, 1991; Gray & Gray, 1997; Oorlog, 1995; Pope & Peel, 1989; Sauer, Brajer, Ferris, & Marr, 1988; Zuber, Gandar, & Bowers, 1985).

A concern that is raised is that the standard statistical tests do not have enough power to distinguish between hypotheses. Gandar et al. (1988) refer to Summers (1986), who shows that conventional methods for testing market efficiency in traditional financial markets are rather weak. Summers simulates prices under nonrational expectations and shows that the standard tests are too weak to pick up the irrationality. As much research fails to reject the hypothesis of efficiency, it is too often wrongly interpreted as evidence in favor in accepting the hypothesis. This issue is accentuated by papers claiming to find exploitable inefficiencies by designing some kind of trading strategy (Crafts, 1985; Dixon & Coles, 1997; Shank, 2019; Vergin, 1977). The second chapter of this dissertation is dedicated to this point and we will put forward evidence that many of the claimed inefficiencies are probably the result of overfitting, i.e. finding patterns in a sample that do not generalize to the population.

Next, there is a literature dedicated to the question of whether individuals can outperform sports prediction markets. Servan-Schreiber, Wolfers, Pennock, and Galebach (2004) study the relative performance of two prediction markets (TradeSports & NewsFutures) against 1947 individuals when making predictions on NFL games. When all the individuals and the two prediction markets are ranked from best to worst in terms of forecasting accuracy, NewsFutures ranked 6th and TradeSports ranked 8th out of almost 2000 contenders by the end of the 2003-2004 NFL season. This highlights that individuals rarely outperform a prediction market, they need extraordinary luck or skill to do so. In a similar vein, Boulier and Stekler (2003) and Song, Boulier, and Stekler (2007, p. 412) find that the forecasts of the betting market are "substantially superior" to those of newspaper and television journalists.

Interestingly, A. Brown and Reade (2019) find that individuals can outperform the prediction market forecasts when their predictions are aggregated. A. Brown and Reade (2019) find a wisdom of the crowd effect when they combine the decisions of many thousands of individual bettors into a single forecast.

1.7.2 (Geo)political prediction markets

Rhode and Strumpf (2004) discuss the rich history of presidential betting markets in the US between 1868 and 1940. In an era before scientific polling, these prediction markets offered valuable insights into the popularity of presidential candidates and their chances of winning the presidency. Furthermore, these markets where extremely popular as for some periods betting on political outcomes exceeded stock and bond trading. These political prediction markets where closely watched by policy makers and leading newspapers would report on the prices in these markets on a daily basis. Even then, the usefulness of markets to incentivize information discovery and aggregation was already widely understood: the "old axiom in the financial district [is] that Wall Street betting odds are never wrong" (New York Times, 1924, p. E1). The importance of these prediction markets declined drastically after World War 2 due to the rise of opinion polls and questions about the ethics and legality of betting.

The empirical work on the accuracy of political prediction markets shows that prediction markets generally outperform election polls. For example, Berg, Nelson, and Rietz (2008) study US elections between 1988 and 2004 and compare 964 polls with market prices on the Iowa Electronic Market. In 74% of the cases, the market prices where closer to the eventual outcome than the polls which is statistically highly significant. Furthermore, the average absolute error for the 964 polls was 3.37% compared to just 1.82% for the market.

The promise of prediction markets has also not escaped the attention of the intelligence community. In the US, the Defense Advanced Research Projects Agency (DARPA) launched a prediction market initiative in 2003 (Wolfers & Zitzewitz, 2004). DARPA was founded in 1957 after the Russians launched Sputnik, with the goal of working on breakthrough innovations so that the US would never be technologically surprised again by an adversary. The goal of the prediction market project was to get good estimates of geopolitically important information like the economic or military power of certain countries or the probability of specific conflicts. The contracts traded in the prediction market might involve questions as "Will dictator X from country Y be overthrown before date Z" or "Will China's military spending grow by more than X% by year Y". However, the "Policy Analysis Market", as the project was known, quickly got cancelled after political uproar and a media storm where critics accused DARPA to be creating the opportunity to bet on future terrorist attacks. Criticism that by the way was largely unfounded as the large majority of the events underlying the contracts were general geopolitical trends rather than specific terrorist attacks (Hanson, 2005).

Anecdotal evidence shows that prediction markets on geopolitical risks could have worked well. On 2 May 2011, Osama bin Laden was killed during a US Navy SEAL operation. In the months before, a contract on whether bin Laden would be captured or killed before 31 December traded on intrade.com, a prediction market platform. The probability of the underlying event hovered around 7% for a long time, before it quickly rose to 99% on the day of Bin Ladens death (Snowberg et al., 2013). Interestingly, this determined market movement happed before any of the mainstream media outlets had reported about the death of bin Laden. This illustrates both the speed at which markets incorporate new information and the effectiveness of the incentives markets give to individuals to bring information to the market.

A more elaborate discussion of how governments could use prediction markets to improve policy making can be found for example in Hahn and Tetlock (2005) or Hanson (2013).

1.7.3 Corporate prediction markets

Corporates have also been experimenting with internal prediction markets to make forecasts. These internal markets try to aggregate information that is scattered across different roles, departments, and branches. Furthermore, they can allow employees to anonymously trade on their information and opinions while they have skin in the game. Opinions which they might not openly share in a meeting for political reasons or "shooting-the-messenger" concerns. The outcomes of these corporate prediction markets can meaningfully influence decision making. Ford Motor Company for example noted in a 2011 press release that they will not include a bike rack and an in-car vacuum cleaner in new car models based on a prediction market outcome (Cowgill & Zitzewitz, 2015).

In the relevant empirical work, Plott and Chen (2002) report that an internal prediction market run at Hewlett-Packard (HP) predicted product sales more accurately than the official HP forecast. Similar results were obtained at Siemens where a prediction market accurately forecasted an unforeseen delay in a software project (Ortner, 1998). More recently, Google and Ford Motor Company have been experimenting with corporate prediction markets as well (Cowgill & Zitzewitz, 2015). At Ford, the prediction market had a mean-squared error that was 25% lower than the expert panel the firm traditionally used to forecast weekly sales. In the corporate prediction market run at Google, Cowgill, Wolfers, and Zitzewitz (2009) document an optimism bias in an otherwise reasonably efficient market. This could be due to a genuine optimism bias or due to employees being overly positive about the outcomes of their own projects to influence their manager's views of their work. However, the optimism bias fades away. The market becomes more efficient over time as traders get more experienced and self-selection pushes unprofitable participants out of the market.

Other examples include prediction markets at Nokia (Hankins & Lee, 2011), Intel (Gillen, Plott, & Shum, 2013) and Yahoo (Mangold et al., 2005)[20]. Many of the authors in the corporate

[20] Other companies that have implemented internal prediction markets include Abbott Labs, Arcelor Mittal, Best Buy, Boeing, Chevron, Chrysler, CNBC, Corning, Electronic Arts, Eli Lilly, Frito Lay, General Electric, General Mills, Hallmark, InterContinental Hotels, J&J, Lockheed Martin, Masterfoods, Microsoft, Missile Defense Agency, Misys, MITRE Corp., Motorola, NASA, Nucor, Overstock.com,

prediction market literature note that it is quite remarkable that these markets perform well although they have much less participants than public prediction markets. Indeed, a trader population of a few dozen employees already leads to fairly accurate predictions. The relative lack of liquidity does not appear to be a serious issue, which is in line with the results from experimental research where asset markets appear to aggregate information very well, even with a limited number of traders (Sunder, 1995).

1.7.4 Economic prediction markets

In the domain of economic forecasting, Goldman Sachs and Deutsche Bank ran parimutuel markets on "economic derivatives" i.e. prediction markets on macro-economic variables. These prediction market forecasts for a range of economic indicators were consistently slightly more accurate than an expert panel. More specifically, the average forecast error of the market was 5.5% lower than that of the experts but this reduction is not statistically significant (Gürkaynak et al., 2005; Snowberg et al., 2013). Furthermore, in horse race regressions like in Fair and Shiller (1990) where both the expert and market forecasts are regressed on the actual outcome, the market forecasts are highly significant and have a coefficient close to unity while the expert forecasts are not statistically significant and even have a coefficient below zero. Moreover, in a test where the forecast errors of the market are regressed on a number of variables like the past forecast errors, the null of market efficiency could not be rejected while the null of efficiency of the expert forecasts is rejected (Gürkaynak et al., 2005). Relatedly, Teschner, Stathel, and Weinhardt (2011) ran a prediction market in Germany in cooperation with the newspaper *Handelsblatt* to forecast German GDP, inflation, investments, exports and unemployment figures. They find that the prediction market estimates are marginally better than Bloomberg forecasts although the differences are not statistically significant. Furthermore, they show that their markets are at least weakly efficient as past price changes do not predict future changes.

1.8 Limitations of prediction markets

The empirical work on prediction markets is generally consistent with the view that these markets are informationally efficient and outperform alternative forecasting methods. However, there are some caveats that have to be considered when interpreting prediction market prices. In the next sections, we discuss some potential limitations related to participants, biases, cost benefit ratio and incentives of prediction markets.

1.8.1 Participants and biases

A potential shortcoming of prediction markets is that they could be plagued by similar behavioral biases as traditional financial markets (Hirshleifer, 2015). For example, prediction markets typically perform badly when forecasting very high and very low probability events,

---

PayPal, Pfizer, Procter & Gamble, Qualcomm, Rite Solutions, SanDisk, Sony, Starwood, TNT, T-Mobile and the WD-40 Company (Cowgill & Zitzewitz, 2015; Snowberg et al., 2013). However, as many of the questions in these internal prediction markets are sensitive there has been little reporting on the mechanics and outcomes of these markets.

a phenomenon that is called the favorite-longshot bias. Griffith (1949) for example already demonstrates that market implied probabilities less closely approximate empirical probabilities for very (un)likely events. More precisely, the returns on low probability events are too low while the returns on low probability events are too high. The existence of a favorite-longshot bias has been documented in dozens of studies in i.a. horse racing (L. V. Williams & Paton, 1997), soccer (Angelini & De Angelis, 2019), college basketball and college football (Berkowitz, Depken II, & Gandar, 2017) but also in political betting markets (Snowberg et al., 2013). One of the main explanations of this behavior is that individuals are attracted to skewness i.e., assets with lottery-like characteristics (Golec & Tamarkin, 1998). This tendency was already described in very early experimental research by Preston and Baratta (1948) and is also observed on stock markets (Annaert, De Ceuster, & Verstegen, 2013). As a result, Snowberg et al. (2013, p. 13) urge to "use extreme caution when interpreting results based on contracts that imply a risk-neutral probability between 0 and 10%, or 90% and 100%".

A related concern is the number of noise traders on prediction markets. It might seem counterintuitive, but markets need liquidity from noise traders to function smoothly. If markets are only populated by informed, rational traders whose only concern is expected return, the market unravels due to the no trade theorem (Milgrom & Stokey, 1982). This is because rational traders will not accept to enter into a trade with a counterparty if they know that the counterparty is also knowledgeable and rational. The counterparty only makes the offer because they have private information and the trade will only make themselves better off. Uninformed traders incentivize informed traders to bring information to the market.

Efficient price discovery could also be hindered by speculative bubbles i.e., market prices getting uncoupled from fundamental values (Haw Allensworth, 2009). It is a stylized fact in the experimental asset market literature that even in lab settings, where fundamental values can be known by design, bubbles and crashes take place (see Palan (2013) for a review). In the prediction market literature, Wolfers and Zitzewitz (2004) put forward the hypothesis that a speculative bubble occurred in contracts on Hilary Clinton becoming the Democratic presidential nominee in 2003. The price of these contracts was very elevated although Clinton publicly declared she was not a candidate. Another example is the work by Berlemann and Vöpel (2012) that claims to have identified a bubble in a sports prediction contract on which country would win the 2010 World Cup. The price of the contract that paid 1 if Germany would win the World Cup quickly rose from 0.19 to 0.30 in a matter of hours before reverting to 0.16 during a period when apparently no new information was released. Interestingly, in contrast to experimental asset markets where bubble formation appears to be a stylized fact, the evidence in prediction markets of bubbles is very scarce and anecdotal.

1.8.2 Unfavorable cost/benefit ratio

Another consideration is that although prediction markets generally outperform other methods like polls or statistical models, the relative advantage is often pretty small[21]. Accordingly, Goel, Reeves, Watts, and Pennock (2010) somewhat temper the prediction market enthusiasm by urging policy makers to assess both the relative costs and the relative benefits of prediction markets over traditional models before rushing to conclusions. In their analyses for example, the Las Vegas football betting market is only 3% more accurate than a trivial statistical model and only 1% more accurate than a panel of experts. Similar results are found in baseball and the prediction of the box-office revenues of movies. In every case the prediction market performs best but differences are so small that they will not be meaningful in many forecasting exercises. Gadanecz, Moessner, and Upper (2007) reach a similar conclusion with respect to forecasting economic variables. More generally, these reflections point to strongly diminishing returns to information and suggest that there are stringent limits to the extent that uncertain events can be predicted in the first place (Goel et al., 2010).

Similarly, (Yagudin, Sempere, & Lifland, 2021) cite high costs as a potential reason why corporate prediction markets are often discontinued. They give the example of an internal prediction market at Google where traders cumulatively spent almost 10000 hours on the platform, which results in an opportunity cost of labor of almost $2 million given the high wages of Google employees. Generally, when the event that is being forecasted by the prediction market is not of vital importance, the additional cost of running a prediction market can quickly outnumber the potential benefits of improved accuracy.

1.8.3 Incentives

The incentives prediction markets can offer to informed traders to join the platform can also be limited because of legal reason (Wolfers & Zitzewitz, 2006). The Iowa Electronic Markets for example are legally required to limit trades to $500 per account. Other platforms sometimes even opt to only use play money instead of real money to bypass regulation. Are these limited incentives strong enough to convince traders that spending time on prediction markets is worthwhile? Surprisingly, Servan-Schreiber et al. (2004) find that markets that use play money are as accurate as real money markets. Relatedly, Rosenbloom and Notz (2006)

---

[21] Relatedly, although prediction markets are often praised for their power to forecast presidential elections better than opinion polls, Erikson and Wlezien (2008) argue that their performance cannot be directly compared, which many researchers do. Prediction markets deliberately assess the final outcome of the election whereas polls merely record voting preferences of the participants on the day the poll was taken. When they statistically correct for the difference, opinion polls seem to outperform markets. “The implication is that markets are slower to recognize election winners than what can be learned by applying a reasonable understanding of polling history to the interpretation of current polls” (Erikson & Wlezien, 2008, p. 194).

find that both play and real money markets perform well. However, they find that real money markets statistically outperform play money markets for non-sports related contracts.

A final concern is that prediction market prices could be manipulated. For example, a presidential candidate who is trailing by a large margin could be interested in skewing the prediction market forecast upward to mobilize voters who would not come out to vote if they have the feeling that the race is already run. Hansen, Schmidt, and Strobel (2004) report evidence of a German political party (FDP) that tried to influence its predicted vote share. The party was struggling to reach the 5% voting threshold required to be represented in parliament and wanted to create a bandwagon effect and positive media coverage with an increasing prediction market price. However, although the potential for manipulation is often mentioned as a possible drawback of prediction markets, manipulation attempts are rarely successful. “None of them had much of a discernible effect on prices, except during a short transition phase” (Wolfers & Zitzewitz, 2004, p. 119). Hanson and Oprea (2009) argue that manipulators add incentives for informed traders to enter the market by increasing the expected returns. As a result, the market mechanism counteracts manipulation attempts. In the model of Hanson and Oprea (2009), the presence of manipulators can even increase prediction market accuracy as it incentivizes other traders to become even better informed.

We refer the interested reader to Yagudin et al. (2021) for a more elaborate discussion on the limits of prediction markets in a corporate setting. Yagudin et al. (2021) discuss i.e. the technological, cultural and semantic challenges that arise when implementing a prediction market.

1.9 Structure of this dissertation

Chapters 2 to 4 of this dissertation contain 3 empirical studies which are discussed in more detail below. In these chapters, we examine market efficiency from different angles and in different microstructures. We do not solely focus on the efficient market hypothesis, but also think about related questions on meta-science, industrial organization and trader psychology. We limit ourselves to studying sports prediction markets, in line with much of the earlier literature, as these are by far the most mature and liquid markets, contract maturities are very short and data is easily and excessively available.

Chapter 2: Efficient Spread Betting Markets: A Literature Review
This chapter is published in the Journal of Sports Economics:

> Vandenbruaene, J., De Ceuster, M., & Annaert, J. (2022). Efficient Spread Betting Markets: A Literature Review. *Journal of Sports Economics*, 1-43 DOI: 10.1177/15270025211071042

In this chapter, we study market efficiency in spread betting on American sports games (basketball and American football). In particular, we review the literature that claims to have found profitable trading strategies. We find that although the strategies appear statistically significantly profitable in isolation, their profitability vanishes when the correct joint-hypotheses frameworks are applied. Next to putting these studies in a new light, we also urge for more meta-scientific reflexes in the related literature.

Chapter 3: How the Internet Can Shape Markets: the Case of Sports Betting

In this chapter, we look at transaction costs in English Premier League fixed-odds betting. The transaction costs give insight into the pricing strategies of bookmakers and the evolution of the industry over time. We show that the price setting is consistent with the canonical industrial organization model of Varian (1980). Furthermore, we show that the favorite-longshot bias, one of the main stylized facts in the betting literature, dissipated over time. This is consistent with the cost explanation of the favorite-longshot bias put forward by Hurley and McDonough (1995).

Chapter 4: Does Time Series Momentum Also Exist Outside Traditional Financial Markets? Near-Laboratory Evidence From Sports Betting

In this chapter, we investigate whether time-series momentum, an anomaly present in traditional financial markets, also exists in sports betting. We analyze 17380 soccer games worldwide and find a strong presence of momentum in pre-game betting odds. Moreover, we leverage the near-lab characteristics of betting markets to assign the effect to behavioral underreaction. This could be an out-of-sample indication that the underreaction often documented in financial markets indeed has a behavioral root.

Chapter 2: Efficient Spread Betting Markets: A Literature Review

Abstract

Are simple trading strategies profitable? It is a question that has been on the minds of academics and practitioners for decades. In this paper, we review the longstanding literature on trading strategies in spread betting (also known as handicap betting), a popular sports betting microstructure. We review over 600 strategy implementations and find that market efficiency and systematic misperceptions are not mutually exclusive per se. Predictable glitches occur, but they are too small to be profitably exploited which is consistent with efficient markets. Furthermore, while controlling for data mining issues is becoming mainstream in finance, it has not yet made its way into this literature. We provide evidence that the hurdle rate of $|z| > 3$ which has been put forward in the broader finance literature should also be used in betting market research.

JEL Classification

C12, G14, G40

Keywords

Market efficiency, sports betting, spread betting, handicap betting, literature review

In this chapter we review the longstanding literature on simple rule-of-thumb or mechanical strategies in sports betting. The quest for profitable trading strategies receives substantial attention in the broader finance literature. Practitioners are interested in finding methods for accumulating wealth, while academics are interested in the informational efficiency implications of profitable strategies (or both). In an efficient market, asset prices summarize all available information such that simple trading rules cannot lead to risk-adjusted excess returns (Fama, 1970). The existence of persistently profitable strategies could, for example, expose significant behavioral biases which can have resource allocation implications. In financial markets, strategies that consist of sorting assets on price-to-fundamentals ratios (value) or on their past performance (momentum) are generally profitable (Asness, Moskowitz, & Pedersen, 2013). However, it is not clear whether these are genuine market inefficiencies or rational risk compensations. Sports betting markets, due to their design simplicity, provide more direct tests of market efficiency.

Sports betting markets have a long history[22] in both economics and psychology research as they are essentially “simple financial markets” (Sauer, 1998, p. 2021). In contrast to earlier, more general literature reviews (Sauer, 1998; Thaler & Ziemba, 1988) we focus on easily implementable mechanical strategies. We zoom in on such strategies as they provide a more direct test of market efficiency compared to tests based on regressions or on the statistical modeling of underlying game variables, which are also common in the literature. Furthermore, we concentrate on point spread betting (also known as handicap betting), the market microstructure where all bets have a winning probability of close to 50% by design. This setting has the methodological advantage that all the assets have identical risk-return characteristics (Dana & Knetter, 1994). Furthermore, as the returns of bets across different games are independent, the returns are iid[23]. With the risk explanation crossed out, persistently profitable trading rules that are easily implemented and based on public information are direct evidence of market inefficiencies.

We review more than 40 years of literature and over 600 strategy implementations and find evidence of statistically significant market inefficiencies. For example, the market quite persistently misestimates the probability that underdogs will beat the spread. Leveraging this information increases returns above that of a naïve, random trading strategy. At first sight, we also find economically significant market inefficiencies. However, the sports betting

[22] Betting arguably even lies at the origin of modern probability theory. Mathematicians Blaise Pascal and Pierre de Fermat developed fundamental probability concepts while discussing a game of chance (Devlin, 2010). Furthermore, via the solutions of the St. Petersburg paradox, which involves a theoretical lottery, many core economic concepts like utility functions and expected utility maximization were introduced (Bernoulli, 1954; V. L. Smith, 1971)

[23] Note that this is generally not the case in other betting market microstructures like pari-mutuel or fixed odds betting. In these markets, assets with very different risk-return profiles coexist. This induces a need to adjust for the risk-return differences between the assets as agents generally seem to prefer lottery-like assets with a low probability of a high return over assets with a high probability of a low return (Bird & McCrae, 1987). This empirical regularity is called the “favorite-longshot bias”.

literature is plagued with type 1 error inconsistencies i.e. there are many examples of papers claiming to find inefficiencies that were later rebuked by out of sample tests. It is common practice to test a battery of strategies based on some easily observable variables for a wide range of parameter values while only vaguely referring to data mining issues. Furthermore, in the papers we review, statistical methods that control for the number of hypotheses tested were never used. The hurdle rates designed for single hypothesis testing (like $|z| > 1.96$) are routinely used in a multiple testing exercise. Our analyses based on three multiple testing adjustments (Bonferroni; Holm; Benjamini, Hochberg, and Yekutieli) indicate that a hurdle rate of $|z| > 3$, which was put forward by Harvey, Liu, and Zhu (2016) for research in equity markets and Benjamin et al. (2018) for research communities in general, should also be the hurdle rate for betting market research. Under this stricter hurdle rate, none of the reviewed strategies were significantly profitable after transaction costs, which is consistent with an efficient sports betting market. Lastly, we observe a strong inverse relationship between the profitability of a strategy and its sample size. This observation is again in line with an efficient market where inefficiencies are chance results.

The usual disclaimer for literature surveys applies. We summarize, interpret and discuss many important results, but this review is by no means a complete catalog of all papers that have been written on the subject. The rest of this chapter is structured as follows. In section 2.1 we discuss the usefulness of sports betting as a research lab for finance. Section 2.2 introduces the point spread betting market microstructure. Section 2.3 discusses the methodology used to benchmark statistical and economic significance given the large number of tests conducted in the literature. Section 2.4 zooms in on individual strategies while section 2.5 zooms out and puts the results in context. A further discussion can be found in section 2.6 and section 2.7 concludes.

2.1 Sports betting as a research lab for finance

Empirical work in sports betting markets dates back to Griffith (1949). Since then, countless researchers have embraced the methodological advantages of sports betting markets to test their hypotheses. The links between sports betting markets and traditional financial markets like the stock market are clear. Both are competitive speculative markets in which a large number of participants collectively determine the prices of assets whose future payoffs are uncertain (Ali, 1979). Moreover, sports betting, like trading derivatives and active asset management, is a zero sum game (before commissions) (Levitt, 2004). However, sports betting markets have several features that make them interesting research labs in general and specifically allow for notably clean efficiency tests.

a) The assets are very simple. Sports bets are typically binary options that have a single positive payoff if the underlying event takes place. This payoff structure is very easy to understand for all parties involved which can ease efficiency. In a lab setting, Carlin et al. (2013) for example show that lower asset complexity leads to higher efficiency.
b) The assets have very short maturities of days to hours or even minutes. This relatively short time span allows individuals to quickly evaluate their investment decisions and

can enhance learning (Thaler & Ziemba, 1988). In experimental research, Forsythe et al. (1982) stress the importance of replication for asset prices to converge to a rational expectations equilibrium. Furthermore, the very short maturities virtually remove any necessity to incorporate the time value of money in analyses.

c) The assets' true values are exogenously revealed. The event outcomes are known ex post and are independent of the behavior of traders. This circumvents the dreaded joint hypothesis problem[24] as researchers can systematically compare market prices of assets with their true values (Campbell et al., 1997; Thaler & Ziemba, 1988).
d) The expected payoff of a sports bet at a particular point in time is idiosyncratic and does not comove with aggregate risk factors (Moskowitz, 2021; Snyder, 1978). This is very different from capital market assets where the returns are correlated with each other and the stochastic discount factor.
e) The information set relevant to the pricing of sports bets is much smaller compared to that of a multinational company which can enhance efficiency as the attention span of traders is limited (Hirshleifer et al., 2009; Simon, 1971).
f) The sports betting landscape consists of very different market microstructures (point spread betting, pari-mutuel betting, fixed odds betting…) in virtually any sport. This element coupled with the depth of historical data that is available provides researchers a wealth of natural experiments (for recent examples, see Berkowitz, Depken II, and Gandar (2015), A. Brown (2014), Croxson and Reade (2014) or Mills and Salaga (2018)).
g) Many of the above features can be replicated in a lab setting, but the gain in controllability that experiments offer is at least partially offset by external validity concerns (Levitt & List, 2007). In betting markets agents can be studied in their natural habitat, without being aware that they are observed and with real money at risk.

As a result, researchers have gratefully used these "market[s]-in-miniature" (Hausch & Ziemba, 1990a, p. 61) in many topics including the market's forecasting abilities (Asch, Malkiel, & Quandt, 1982; Griffith, 1949) arbitrage relations (Franck, Verbeek, & Nüesch, 2013; Marshall, 2009) testing prospect theory (Snowberg & Wolfers, 2010) or asset price clustering (A. Brown & Yang, 2016).

[24] Testing efficiency in stock markets is a notoriously fishy undertaking. Market prices can never be compared with the true value of stocks as the latter are never revealed. Researchers can resort to models that generate theoretical prices and compare these to market prices. However, when discrepancies arise, it is not clear whether the market prices are wrong or whether the model that generates theoretical prices is wrong, or both. This fundamental untestability of efficiency in stock markets is called the joint hypothesis problem.

2.2 Point spread betting market microstructure

In spread betting, agents bet on whether a team is going to win by more or lose by less than the point spread. Point spreads are set by bookmakers i.e. market makers[25] who are the counterparty to all gamblers. The point spreads are set in proportion to the relative team qualities. This equalizes the probability of winning a bet on either team. As an example, suppose a very strong team plays against a very weak team. A simple bet on which team will win the game will heavily favor the stronger team. However, with a spread, the bookmaker can level the playing field by requiring not only that the stronger team wins, but that it wins by, for example, at least a 14-point difference. Bookmakers typically first announce their point spread a few days before the game (i.e. the opening spread). The spread can change because of i.a. game-related news or large volumes placed on one of the teams, right until the game is about to start (i.e. the closing spread). However, whenever a gambler makes a bet, the point spread quoted on the moment the bet is made is locked in. Subsequent spread changes only affect the gamblers who enter later. In contrast to pari-mutuel betting, a gambler knows all the conditions of the bet when it is made.

Point spread betting is arguably the most popular betting microstructure in the United States. Spread betting is mostly associated with American football and basketball where it is common to score many points in a game. The betting industry grafted onto these two sports is enormous. In 2020 for example, more than 10% of adult Americans indicated they would bet on the Super Bowl, the most important football game of the year, and in 2019, 20% of adult Americans indicated they would bet on March Madness, the NCAA men's basketball tournament (American Gaming Association, 2019, 2020). The popularity of spread betting is sometimes explained by the increased thrill of betting on a score difference compared to betting on the outcome. Alternatively, under some circumstances it could be more profitable for a bookmaker to offer spread bets than to offer fixed odds bets (Bassett Jr, 1981). Research in these point spread betting markets dates back to Pankoff (1968) who explicitly introduced the efficiency jargon in the betting literature, inspired by his contemporaries Fama (1965) and Mandelbrot (1966).

[25] Jaffe and Winkler (1976) discuss the similarities between market makers in financial and sports betting markets and their relationship with investors. Furthermore, it is important to appreciate that the risk bookmakers take is categorically different from that of other gambling establishments like casinos. While the latter exploit the law of large numbers to secure asymptotically certain profits, bookmakers can suffer large losses when they systematically misestimate game outcome probabilities or bettor behavior. While the outcome probabilities for a casino game like roulette are common knowledge, pricing a sports bet is much harder. This introduces incentives for sports bettors to gather information as they do not just rely on luck (like their roulette colleagues), but also on their ability to correctly estimate game outcome probabilities (Figlewski, 1979). (Or at least it appears. In an efficient market, the marginal profits to analyzing information are again zero.)

If the spread indeed equalizes the win probabilities of bets on either team the fair odds would be 2. However, bookmakers are not in the business for the fun of it, so they charge a fee for their services just like market makers in traditional financial markets. Payout happens according to the 11 for 10 rule[26]. This means that an \$11 winning bet only yields a profit of \$10. This is below the fair payout, which allows the bookmaker to make a profit. A gambler who wants to break even must achieve a win fraction of at least $\frac{11}{21}$, approximately 52.4% (or alternatively, lose less than $1-\frac{11}{21}$, approximately 47.6%). This can be seen by solving

$$f \times 10 - (1-f) \times 11 = 0 \tag{1}$$

to f, the fraction of winning bets.

If the total amount bet is perfectly balanced between the two teams, the bookmaker pays out \$21 for every \$22 it receives. Traditionally, bookmakers were understood to focus on achieving such a balance. By doing so, they take no risk as they can pay off the winning bets by the losing bets and collect a commission along the way. In this view, excessive volume placed on one team will induce the bookmakers to adjust the point spread in order to incentivize gamblers to bet on the other team. As a result, if the spread differs from the market's consensus, market forces will push the spread towards the equilibrium value. This also means that the point spread is not necessarily an unbiased predictor of the margin of victory. If the market's expectations are biased, bookmakers will anticipate and purposely bias the point spread to equalize the volumes bet on both sides to avoid having to take an active position in the game. As a result, the point spread will be a forecast of the market's expectation of the game outcome instead of the game outcome itself. More recent research however shows that bookmakers are not trying to nullify their risk in every game. Bookmakers can earn more when there are more losers whose stakes can be collected than winners who have to be paid. There is empirical evidence that bookmakers indeed maximize their profits by offering slightly biased lines, i.e. more than 50% of the volume on one side and take active positions in the game outcome as a result (Levitt, 2004; Paul & Weinbach, 2011; Strumpf, 2003). There is of course a limit to how far bookmakers can go with such practices as witty gamblers will quickly exploit flagrant profit opportunities.

It is worth mentioning that a strand of the spread betting literature examines the totals market where gamblers bet on the total number of points scored by the two teams combined. The efficiency of this market is beyond our scope. We refer interested readers to Paul and

[26] Levitt (2004) notes that a "major puzzle in this industry is the rarity of price competition, i.e., the vig is almost universally 10%". This point is further explored by Sandford and Shea (2013). They attribute it to the first mover disadvantage bookmakers have when setting their lines. Gamblers can consequently make their bets with information bookmakers did not have when they set their lines. More recently however, bookmaker competition starts to bring down the commission charged (Berkowitz, Depken II, & Gandar, 2018). Papers in which the authors state the 11 for 10 commission structure did not apply (which were very few) were not included in this review to keep the hurdle rate constant.

Weinbach (2002), Paul, Weinbach, and Wilson (2004) and DiFilippo, Krieger, Davis, and Fodor (2014).

2.3 Methodology

Spread betting brings about the methodological advantage that the probability of winning a series of bets can be modelled via a binomial distribution where successive outcomes are independent. In large samples, the binomial distribution can be conveniently approximated by the normal distribution. Two benchmarks are commonly used to evaluate the performance of trading strategies.

a) **Statistical efficiency:** the win fraction is indistinguishable from randomness (50%). Under the null, the point spreads reflect all information such that the expected return of every bet is equal.
b) **Economic efficiency:** the win fraction is not significantly higher than 52.4% (or lower than 47.6%). Under this null hypothesis expected returns do not have to be equal, but differences cannot be so large that profit opportunities arise.

The advantage of using these two benchmarks jointly is that both exploitable and unexploitable inefficiencies can be identified.

The benchmarks result in the following hypotheses and test statistics (Woodland & Woodland, 1997).

Hypothesis 1: the trading strategy is statistically efficient:

$$H_{0,1}: \pi = 0.5$$

$$H_{a,1}: \pi \neq 0.5,$$

where π is the win fraction. The test statistic is

$$Z_1 = \frac{(\hat{\pi} - 0.5)}{\sqrt{\frac{(0.5)(1-0.5)}{n}}},$$

where $\hat{\pi}$ is the empirical win fraction and n the number of bets.

Hypothesis 2: the trading strategy is economically efficient:

$$H_{0,2}: \pi = \frac{11}{21}$$

$$H_{a,2}: \pi > \frac{11}{21},$$

with a similar test statistic:

$$Z_2 = \frac{\left(\hat{\pi} - \frac{11}{21}\right)}{\sqrt{\frac{\frac{11}{21}\left(1 - \frac{11}{21}\right)}{n}}}.$$

In the discussion of the trading strategies, we will only report this second test statistic when the strategy is profitable (empirical win fraction larger than $\frac{11}{21}$), and we can reject the null of randomness (at the 5% significance level). For strategies with winning percentages significantly below 50%, we use the benchmark of $1 - \frac{11}{21} \approx 47.6\%$.

To further streamline the exhibition, we only present the *z*-statistics defined above in the analyses. Some older papers lack significance tests or use other methods including the test proposed by Tryfos, Casey, Cook, Leger, and Pylypiak (1984), which was shown to be slightly biased by Woodland and Woodland (1997), or use a likelihood ratio test (Even & Noble, 1992). In these cases, the above *z*-statistics are computed if the required data are provided. Furthermore, to save space and avoid data mining issues we only present consolidated results on the longest time period available in each paper and leave out the year by year analyses. Moreover, a small number of articles deploy strategies for which it is not clear that they can be implemented (strategies that rely on closing line information or strategies that assume more favorable point spreads could be obtained by setting up a betting syndicate that exploits price differences between different regions). These strategies are not included in this overview. Lastly, ties are excluded from the analyses as it is common bookmaker policy to simply refund bets in these scenarios (or avoid ties by non-integer point spreads).

2.3.1 Multiple Testing

We initially benchmark the test statistics against the common single hypothesis test values of $|z| > 1.96$ and $|z| > 1.64$ for the two- and one-tailed tests respectively. The strategy implementations with *z*-statistics that exceed these critical values are deemed statistically significant in the original studies. However, the trading strategy literature in spread betting is a textbook example of a situation where corrections for multiple testing are crucial to limit flagrant *p*-hacking. Scholars test hundreds of possible strategies, often without any theoretical underpinning. When enough strategies are tried out, significant results will be found even if the null is true, by construction of the hypothesis test (type 1 error). In this paper, we review 628 strategies, so the risk of many type 1 errors is very real. Moreover, when researchers find an interesting strategy, they often start digging in the periphery. As a result, many slight alterations of the same profitable strategy are proposed. Alternatively, some promising strategies are tested multiple times in similar or overlapping datasets (for example, first between 1970-1985, and in a later follow-up study between 1970-1995). Some implementations are so similar that the returns are almost identical and the *z*-statistics very highly correlated. An example from the reviewed strategies includes betting on home underdogs when the spread is 8.5 in the NBA between 1995-2002 and betting on home underdogs when the spread is 9 in the NBA between 1995-2002. If we count both these strategies, we are essentially double counting profitable strategies and will vastly overestimate the true degree of inefficiency. In contrast, strategies that are not deemed profitable are often not published (publication bias) and not further dissected which artificially suppresses the number of unprofitable strategies. As a result, we get a lopsided literature that is tilted in favor of profitability. A testable consequence of such a scenario is that we find pockets of profitability centered around a few strategies in a few samples that do not generalize out of their samples and find too few unprofitable strategies.

We try to alleviate the concern related to the number of proposed test by using multiple testing methodologies (we rely on Harvey et al. (2016) who propose a multiple testing framework for finance in general). The issue of correlated *z*-statistics is trickier, we propose a pragmatic approach that limits the overlap between strategies.

2.3.2 Taming the family-wise error rate and the false discovery rate

The significance level α controls the type 1 error rate in a single hypothesis test and is usually set to 5%. When multiple tests are carried out, α should be adjusted. If not, the probability of making a type 1 error, i.e. the family-wise error rate, quickly approaches 100%.

The most common approach to limiting the family-wise error rate to the usual 5% level is the Bonferroni adjustment which shrinks the original α by the number of tests:

$$\alpha^{Bonferroni} = \frac{\alpha}{N},$$

where N is the number of tested hypotheses. The objective of the Bonferroni adjustment is somewhat extreme (controlling the probability of making a single type 1 error), which results in harsh hurdle rates when the number of tests increases. In our case it would amount to rejecting the null of all implementations with *z*-scores above 3.95 for the two-sided tests as we have a sample of 628 hypotheses. Note that we implicitly assume here that all tests that were conducted are included in our sample. This is clearly not realistic but still a useful exercise as the results can be thought of as a lower bound for the hurdle rate. As an example, the hurdle rate would rise to 4.11 if we were to assume that we are only observing half of all conducted tests.

Another well-known method to control the family-wise error rate is Holm's adjustment, which sequentially tests all *p*-values against a dynamic benchmark. The algorithm consists of a few steps:

1) Order the *p*-values from small to large: $p_{(1)} \leq p_{(2)} \leq \cdots \leq p_{(i)} \leq \cdots \leq p_{(N)}$.
2) For each *p*-value (starting from the smallest), check if $p_i < \frac{\alpha}{N+1-i} = \alpha_i^{Holm}$.
3) Reject the respective null if the inequality holds. If the inequality does not hold, do not reject the respective hypothesis and all other hypotheses with larger *p*-values.

Holm's method is dynamic, i.e. the index number i in the denominator makes the hurdle rate different for every hypothesis, in contrast to the Bonferroni method. Note that for $i = 1$, $\alpha_i^{Holm} = \alpha^{Bonferroni}$. For $i = 2$, $\alpha_i^{Holm} > \alpha^{Bonferroni}$, making Holm's adjustment less stringent, leading to more rejections and all rejections via Bonferroni are also rejected via Holm.

A last method we deploy is the Benjamini, Hochberg, and Yekutieli (BHY) adjustment that is algorithmically somewhat similar to Holm's adjustment. In contrast to the previous two adjustments, BHY's targets the false discovery rate, i.e. the expected proportion of false discoveries, and makes sure it stays below α. It consists of the following steps:

a) Order the *p*-values from small to large: $p_{(1)} \leq p_{(2)} \leq \cdots \leq p_{(i)} \leq \cdots \leq p_{(N)}$.
b) Find the largest i such that: $p_i \leq \frac{i}{N \times c(N)} \alpha = \alpha_i^{BHY}$. (It can be shown that setting $c(N) = \sum_{j=1}^{N} \frac{1}{j}$ is suitable under arbitrary dependency among the test statistics.)
c) Reject the respective null hypotheses for p_k for $k = 1, \ldots, i$ and accept the other null hypotheses.

When we apply the multiple testing adjustments to our data, it is important to account for the correlation between z-statistics of the different hypotheses. Correlation among the *z*-statistics, which is certainly present, makes the multiple testing methods too stringent. Consider the extreme case where we test the same hypothesis 100 times. Instead of using the ordinary *p*-value hurdle rate of 5% which would be appropriate as we are essentially conducting a single hypothesis test, we would use $\frac{5\%}{100}$ under the Bonferroni adjustment, which is of course far too conservative.

In the remainder of this section, we apply the multiple testing adjustments to our data set to determine the appropriate critical values. As mentioned previously, the data set contains 628 strategy implementations, but many of these implementations test the same strategy (like for example betting on the home team). Moreover, the samples in which these different implementations of the same strategy are tested often overlap. To determine the appropriate critical values for the *z*-statistics given this dependence, we make a subsample of 85 strategy implementations. In this subsample, the dependence is removed to a large extent as we only include one implementation of each strategy per tournament (so we remove implementations of the same strategy in different periods). If the patterns in betting markets follow a stationary process, tests in different time periods measure the same phenomenon anyway. We do include implementations of the same strategy when tested in a different tournament because of the large institutional differences. If the strategy is tested for multiple parameter values, we still only include one implementation as implementations for subsequent parameter values are often highly correlated. We make two subsamples of 85 strategy implementations via the process described above, one with the smallest and one with the largest *z*-values per strategy and tournament. We compute the appropriate *z*-score benchmarks in both subsamples, which are shown in Table 2.

Table 2: Z-score hurdle rates under different testing methods (naïve single testing, the Bonferroni adjustment, Holm's adjustment and Benjamini, Hochberg, and Yekutieli's adjustment). For the dynamic methods (Holm and BHY) the hurdle rate represents the z-score the first insignificant strategy should achieve in order to reject its null.

Testing method	**Min one sided**	**Max one sided**	**Min two sided**	**Max two sided**
Naïve single test	1.64	1.64	1.96	1.96
Bonferroni adjustment	3.24	3.24	3.44	3.44
Holm's adjustment	3.24	3.23	3.44	3.43
BHY adjustment	3.68	3.31	3.85	3.50

Table 2 provides us a number of different critical values which we can use as a benchmark. Note that all multiple test benchmarks are at least 3. Also note that the multiple test benchmarks are relatively robust to changes in the number of tested hypotheses. Even if we make the widely unrealistic assumption that only 20 hypothesis tests were ever carried out, the Bonferroni hurdle rate would already be above 3. As a result, using a multiple test hurdle rate of $|z| > 3$ is very reasonable (although it increases the probability of type 2 errors). This choice is both consistent with the analysis from table 2 and with the previous proposals of i.a. Harvey et al. (2016) and Benjamin et al. (2018).

2.4 Mechanical trading strategies

To make the trading strategy zoo more manageable we fit the reviewed strategies in the taxonomy shown below.

a) **Game characteristics**
   1) **Home team** (betting on the home team)
   2) **Underdog** (betting on the underdog, i.e. the team that receives a head start via the spread)
   3) **Home underdog** (betting on the home underdog)
   4) **Home favorite** (betting on the home favorite where the favorite is the team that receives a disadvantage via the point spread)
   5) **Familiarity** (for example, betting on a team that plays on the surface it is used to)
   6) **Fatigue** (for example, betting on a team on a road trip)
   7) **Attention & Importance** (for example, betting on the home team in a playoff game)
   8) **Absences** (for example, betting on a team with an absent top player)

b) **Past performance**
   1) **Performance against the spread** (for example, betting on teams that beat the spread last game)
   2) **Performance not against the spread** (for example, betting on teams that won their last 3 games)
   3) **Spread movements** (for example, bet on teams that became larger underdogs between opening and closing spread)

We first distinguish between strategies based on current game characteristics (like the location of the game) and past performance (for example whether a team won the last game or not). The first four items of this first category are individual strategies as they are so common in the literature (home team, underdog, home underdog, home favorite). The next four are container items for strategies related to familiarity, fatigue, attention & importance and absences. The second category (past performance) contains the large family of momentum and contrarian strategies. We subset this category by performance against the spread, performance not against the spread and spread movements. The performance against the spread strategies are especially interesting as they take both the game outcomes and expectations (the spread) into account.

We collected the relevant papers by querying for "spread betting" and either "efficiency" or "efficient" in the EBSCO discovery service. This resulted in a sample of 157 papers. We removed all papers that did not implement strategies and further expanded our sample by reviewing the bibliographies of the relevant papers and the papers that cited the examined papers (backward/forward snowballing). This resulted in a final sample of 46 papers. The last query was carried out in October of 2020. In what follows we highlight the most important results per strategy family. To make the discussion digestible we will often refer to the appendix where interested readers can find additional strategies. In the tables, we will highlight strategy implementations that reject the null under the single hypothesis benchmark in bold and the strategies that reject the null under the multiple test benchmark in italics.

2.4.1 Strategies based on game characteristics

2.4.1.1 home team

Consistently betting on home teams is one of the simplest and most tested mechanical trading rules. The well-known home-field advantage posits that home teams win more often than visiting teams. In NLF games between 1981 and 2004 for example, the home team outscored the visiting team by 3 points on average (Borghesi, 2007a). Factors that drive this effect include familiarity with the venue, crowd behavior, travel fatigue and referee biases (see Jamieson (2010) for a review). If the market does not adequately reflect this home-field advantage into prices, inefficiencies can occur. Table 3 summarizes the studies that implement the strategy of consistently betting on home teams. The data sets include NFL games, NBA games, college football games, college basketball games and Australian Football League games between 1973 and 2017. Overall, the market correctly discounts the home-field advantage. The empirical win fractions are not consistently above or below 50% and for only two strategies that were profitable in sample, the null of randomness could be rejected at the single hypothesis benchmark (and even in different directions). Furthermore, the null of unprofitability is never rejected.

Table 3: Overview of papers implementing "bet on home team" strategy.

Authors	**Data set**	$\hat{\pi}$	Z_1	Z_2
Lacey (1990)	NFL 1984-1986	0.476	-1.234	
Golec and Tamarkin (1991)	NFL 1973-1987	0.515	1.709	
Golec and Tamarkin (1991)	College football 1973-1987	0.498	-0.251	
Oorlog (1995)	NBA 1989-1991	0.486	-1.312	
Gray and Gray (1997)	NFL 1976-1994	0.504	0.383	
Vergin (1998)	NFL 1984-1995	0.489	-1.087	
Vergin and Sosik (1999)	NFL 1981-1996	0.499	-0.153	
Gandar, Zuber, and Lamb (2001)	NBA 1981-1997	0.495	-1.202	
Kochman and Goodwin (2004)	NFL 1998-2002	0.500	0.026	
Kochman and Goodwin (2004)	Preseason NFL 1998-2002	0.438	**-2.121**	-1.315
Boulier, Stekler, and Amundson (2006)	NFL 1994-2000	0.513	0.854	
Borghesi (2007b)	NFL 1981-2000	0.502	0.324	
Y. T. Sung and Tainsky (2014)	NFL 2002-2009	0.485	-1.307	
Paul, Weinbach, and Wilson (2014)	NFL 2007-2011	0.482	-1.262	
Sinkey and Logan (2014)	College football 1985-2008	0.511	**2.749**	
Humphreys, Paul, and Weinbach (2014)	College basketball 2007-2008	0.495	-0.565	
Coleman (2017)	College football 2004-2011	0.504	0.509	
Shank (2018)	NFL 2009-2017	0.489	-0.951	
Schnyzer and Hizgilov (2018)	Australian Football League 2001-20016	0.533	**2.949**	0.818
Shank (2019)	NFL 2003-2016	0.487	-1.500	

2.4.1.2 underdog

Another trading rule that requires almost no information is betting on the favorite or underdog. These are the teams that received a disadvantage or an advantage via the point spread respectively. Investigating this strategy seems meaningful as in other environments like pari-mutuel betting, it is a stylized fact that returns on favorites are much higher than returns on underdogs (Snowberg & Wolfers, 2010). However, it is worth repeating that in spread betting there are no real favorites or underdogs at the level of the bet. All bets have virtually the same risk-return characteristics, which is not at all the case in pari-mutuel betting or fixed-odds betting where you can regularly make bets at odds of 20 to 1 or more for example. Still, at the game level, the market could misestimate the winning probability of an underdog which can give rise to inefficiencies.

Table 4 summarizes the papers implementing the "bet on underdog" strategy. There appears to be an outspoken bias in favor of underdogs. Of the 22 implementations of this strategy, 20 find that underdogs win more than 50% of the time against the spread. Moreover, the null of randomness is rejected in 9 cases at the single test benchmark and once at the multiple test benchmark. Unprofitability is only once rejected at the single test benchmark.

The market appears to systematically underestimate the quality of underdogs such that the return of a strategy that bets on these teams will be higher than that of a naïve, random strategy. However, the bias is generally too small to be profitably exploited. A possible explanation of the tendency to underbet underdogs is that it is more fun to bet on and root for the team that is likely to win (Paul & Weinbach, 2005a). As the best performing teams receive most media attention, it could also be the case that the volume of news coverage biases gamblers into overestimating the favorites.

Table 4: Overview of papers implementing "bet on underdog" strategy.

Authors	**Data set**	$\hat{\pi}$	Z_1	Z_2
Vergin and Scriabin (1978)	NFL 1969-1974	0.515	0.968	
Tryfos et al. (1984)	NFL 1969-1981	0.526	**2.563**	0.223
Golec and Tamarkin (1991)	NFL 1973-1987	0.524	**2.742**	0.068
Golec and Tamarkin (1991)	College football 1973-1987	0.504	0.678	
Oorlog (1995)	NBA 1989-1991	0.501	0.087	
Gray and Gray (1997)	NFL 1976-1994	0.526	***3.303***	0.276
Paul, Weinbach, and Weinbach (2003)	College football 1976-2000	0.503	0.695	
Kochman and Goodwin (2004)	NFL 1998-2002	0.531	**2.394**	0.545
Kochman and Goodwin (2004)	Preseason NFL 1998-2002	0.581	**2.704**	**1.913**
Paul and Weinbach (2005a)	NBA 1995-2002	0.501	0.261	
Paul and Weinbach (2005b)	College basketball 1996-2004	0.496	-1.255	
Borghesi, Paul, and Weinbach (2009)	NFL 1981-2004	0.518	**2.687**	
Borghesi et al. (2009)	College football 1982-2004	0.510	**2.198**	
Borghesi et al. (2009)	AFL 1998-2006	0.538	**2.413**	0.914
Y. T. Sung and Tainsky (2014)	NFL 2002-2009	0.507	0.631	
Paul, Weinbach, and Wilson (2014)	NFL 2007-2011	0.505	0.344	
Sinkey and Logan (2014)	College football 1985-2008	0.508	1.875	
Humphreys et al. (2014)	College basketball 2007-2008	0.490	-1.215	
(Davis & Krieger, 2017)	NFL 1995-2014	0.503	0.460	
(Davis & Krieger, 2017)	Preseason NFL 1995-2014	0.524	1.656	
(Davis & Krieger, 2017)	NBA 2005-2014	0.501	0.221	
(Davis & Krieger, 2017)	Preseason NBA 2005-2014	0.542	**2.421**	1.045

Motivated by better than even winning probabilities of the unconditional underdog strategy, many researchers implement underdog strategies conditional on some point spread PS. Predicting the score difference in a game between two very unevenly matched teams might be harder which could induce further biases (Vergin & Scriabin, 1978). The results of these conditional underdog strategies are similar to the unconditional underdog strategy. For 46 out of 51 implementations, the underdog beats the spread in more than 50% of the games. Furthermore, the null of randomness is rejected in 17 cases at the single test benchmark and twice at the multiple test benchmark. The null of unprofitability is never rejected at the multiple test benchmark. The supporting tables can be found in the appendix subsection on underdogs.

2.4.1.3 home underdog

Meshing home team and underdog information results in the strategy that most systematically rejects the nulls of both randomness and unprofitability at the single test benchmark. In total, 45 home underdog implementations are reviewed (both unconditional shown in Table 5 and conditional on the point spread, shown in appendix). The null of randomness is rejected 22 times at the single test benchmark and 6 times at the multiple test benchmark. Betting on home underdogs was even significantly profitable (at the single test benchmark) in NFL games between 1973 and 1987. However, in more recent periods, the win fraction is below 50%. It seems that this inefficiency has faded over time, an observation also made by Gray and Gray (1997) when they study the returns of the strategy season by season. At the multiple testing benchmark, the strategy was never profitable.

Table 5: Overview of papers implementing "bet on home underdog" strategy

Authors	Data set	$\hat{\pi}$	Z_1	Z_2
Amoako-Adu, Marmer, and Yagil (1985)	NFL 1979-1981	0.599	**2.743**	**2.085**
Golec and Tamarkin (1991)	NFL 1973-1987	0.556	***3.743***	**2.156**
Golec and Tamarkin (1991)	College football 1973-1987	0.503	0.341	
Oorlog (1995)	NBA 1989-1991	0.479	-0.988	
Gray and Gray (1997)	NFL 1976-1994	0.546	***3.347***	1.627
Vergin and Sosik (1999)	NFL 1981-1996	0.525	1.613	
Gandar et al. (2001)	NBA 1981-1997	0.493	-0.945	
Paul et al. (2003)	College football 1976-2000	0.503	0.340	
Paul and Weinbach (2005b)	College basketball 1996-2004	0.497	-0.407	
Borghesi (2007b)	NFL 1981-2000	0.532	**2.341**	0.572
Borghesi et al. (2009)	NFL 1981-2004	0.530	**2.490**	0.521
Borghesi et al. (2009)	College football 1982-2004	0.522	**2.946**	
Borghesi et al. (2009)	AFL 1998-2006	0.522	0.728	
Y. T. Sung and Tainsky (2014)	NFL 2002-2009	0.512	0.591	
Paul, Weinbach, and Wilson (2014)	NFL 2007-2011	0.481	-0.786	
Sinkey and Logan (2014)	College football 1985-2008	0.519	**2.696**	
Humphreys et al. (2014)	College basketball 2007-2008	0.465	**-1.989**	-0.626
Shank (2018)	NFL 2009-2017	0.490	-0.516	
Shank (2019)	NFL 2003-2016	0.485	-0.992	

Several explanations for the home underdog bias have been proposed. First, large home underdogs are the worst teams in the league and bettors may be hesitant to bet on such low-quality teams. Second, when away favorites are leading by a comfortable margin, they might relax their performance and substitute their best players off the field to avoid injury and fatigue. As a result, the favorite wins the game, but does not cover the spread. This effect is arguably larger for away teams as home crowds will be disappointed if their team does not fully commit or if key players stop playing early. Some even go further and hint that this observation might be consistent with point shaving: corruption where players maximize their utility by both winning the game while at the same time receiving a bribe to fail to cover the spread (Wolfers, 2006). Ashman, Bowman, and Lambrinos (2010) further add that bad teams have little opportunities in a season to get recognized. They might be extra motivated when they get to play a big team at home to prove what they are worth, leading to an unexpectedly good performance.

For completeness, we include the papers implementing the "bet on home favorites" strategy in appendix. The null of randomness is never rejected.

2.4.1.4 familiarity

In this section we review strategies that try to exploit differences in familiarity with game circumstances between the teams.

Boulier et al. (2006) try to exploit differences in playing field surfaces in NFL games (turf versus grass). The authors propose the strategy of betting on the home team when it hosts a visiting team that is used to playing on a different surface. The strategy is profitable in sample and randomness is rejected at the single test benchmark, but not at the multiple test benchmark as shown in row 1 of Table 6.

Borghesi (2007a) investigates whether temperature information can be profitably exploited. To control for the temperature teams are familiar with, the author constructs a temperature acclimatization advantage variable. For example, if a team from Miami plays an away game in New England in December, it is clearly less familiar with the game conditions. One of the strategies he proposes is betting on home teams in the coldest quartile of game day temperatures conditional on the acclimatization advantage, shown in table rows 2 to 5 of Table 6. It appears that the market does not fully incorporate the acclimatization difficulties that occur on the coldest game days when the acclimatization difficulties for the visiting team are the largest. The null of randomness of this strategy implemented for NFL games between 1981-2004 is decisively rejected at the single test benchmark, but again not at the multiple test benchmark. The author also tests the converse strategy of betting on home teams in the warmest quartile of games conditional on the acclimatization advantage, but no statistically significant results are found (see in appendix under the familiarity subsection).

Familiarity with the climate is further investigated by Kuester and Sanders (2011). They find that betting on teams from arid regions when they host teams from humid regions is profitable and the null of unprofitability is even rejected at the single test benchmark, but again not at the multiple test benchmark as shown in row 6 of Table 6. Just like the strategy of Borghesi (2007a) discussed above, the performance of this strategy is also not symmetric. For the converse strategy of betting on teams from humid regions when they host teams from arid regions, the null of randomness is not rejected at the single test benchmark. The difference could be explained by aridity being more performance adverse and harder to acclimatize to. Aridity is also correlated with for example altitude, which has a large impact on the oxygen uptake of the athlete. (A few extensions of the strategy can be found in appendix).

We finish this section with Shank (2019), who studies the performance of home teams in divisional NFL games. The NFL schedule stipulates that each team plays its divisional rivals twice a year, while they are only guaranteed to play teams outside their division once every three or four years. As a result, the familiarity with divisional rivals' coaches, players, tactics etc., is much higher, to which the market can misreact. Indeed, home teams only cover in 47% of divisional games between 2003-2016, rejecting the null of randomness at the single test benchmark, but not at the multiple test benchmark as shown in row 8 of Table 6.

Table 6: Overview of papers implementing "bet conditional on familiarity characteristics" strategies.

	Authors	**Data set**	**Conditioning**	$\hat{\pi}$	Z_1	Z_2
1	Boulier et al. (2006)	NFL 1994-2000	Bet on home team when it hosts a team that plays on a different surface (grass/turf) in its own venue	0.534	**1.970**	0.593
2	Borghesi (2007a)	NFL 1981-2004	Q1 acclimatization advantage (highest)	0.565	**2.416**	1.530
3			Q2 acclimatization advantage	0.540	1.413	
4			Q3 acclimatization advantage	0.515	0.571	
5			Q4 acclimatization advantage (lowest)	0.547	1.565	
6	Kuester and Sanders (2011)	College football 2000-2006	Bet on teams from arid regions when they host teams from humid regions	0.566	**2.653**	**1.704**
7			Bet on teams from humid regions when they host teams from arid regions	0.498	-0.440	
8	Shank (2019)	NFL 2003-2016	Bet on divisional game home team	0.469	**-2.170**	-0.481

2.4.1.5 fatigue

Fatigue is another major factor that can impact game performance and must be accounted for in the point spread. Lacey (1990) and Vergin (1998) devise strategies where the location of a team's previous game proxies for fatigue. Home teams that played at home in the previous game are supposedly well rested while away teams that also played away in their previous game traveled more. The strategies are shown in the first four rows of Table 7 and are never profitable nor significantly different from randomness.

Y. T. Sung and Tainsky (2014) investigate whether the bye-week induces inefficiencies. An NFL season consists of 17 weeks where each team plays only 16 games. This means that every team gets one week off each season (between the fourth and tenth week of the season). The bye-week gives players and staff the time to unwind and rest. The authors build their hypotheses on the strand of the psychology literature that established a positive relationship between days off and subsequent performance. As a result, if the betting market does not accurately estimate the value of the bye-week, the performance of teams that took a week off might be underestimated. They propose a battery of strategies of which we highlight a selection in rows 5 to 8 of Table 7 (rest shown in appendix). Interestingly, the null of unprofitability is rejected for two strategies at the single test benchmark: betting on the favorite after it had a bye-week and betting on the away favorite after it had a bye-week. This last strategy has an empirical win fraction of over 73%. Betting on underdogs after their bye-weeks was never profitable. The authors hypothesize that rest affects strong teams and weaker teams differently. However, if we benchmark the results at the more appropriate multiple test critical values, none of the strategies are statistically significant.

Ashman et al. (2010) test whether player fatigue is correctly priced in NBA point spreads. As NBA teams often face games on consecutive nights, fatigue is more of an issue for basketball players compared to athletes in other sports. Table 7 rows 9 to 11 show the result of betting on the home team when respectively the visiting team, both teams and home team had back-to-back games. Randomness is rejected for home teams playing back-to-back games at the multiple test benchmark. Apparently, the betting market does not fully recognize that fatigue at least partially cancels out the home field advantage.

Ashman et al. (2010) further dissect the results from this last strategy conditional on the number of days of rest the away team had shown in rows 12 to 20 of Table 7. Furthermore, in rows 13, 16 and 19 the sample is limited to games where the home team traveled one or two time zones eastwards between their back-to-back games. Underperformance arising from eastward travel is in line with Jehue, Street, and Huizenga (1993) who find that West Coast teams perform badly when they travel to the east. Home team underperformance is statistically significant at the multiple test benchmark when the visiting team rested for one or two games. Inspired by the above results, Ashman et al. (2010) further condition the strategies on other game related information (shown in appendix in the fatigue subsection). The results are qualitatively similar.

We end this section with Schnyzer and Hizgilov (2018) who specifically focus on the effect of jet lag induced fatigue. They study the Australian Football League, which has the methodological advantage that many games take place on neutral grounds. Jet lag has been shown to worsen the performance of athletes, but the question of course is whether the betting market efficiently incorporates this information into prices (Jehue et al., 1993). Results of several strategies conditional on the jet lag of the visiting team are shown in Table 7 rows 21 to 28. Interestingly, there appears to be no jet lag effect on neutral grounds, but there appear to be jet lag effects on the non-neutral grounds (relative to the single test benchmark only). The authors argue that the jet lag effect on non-neutral grounds is just a home team bias.

Table 7: Overview of papers implementing "bet conditional on fatigue characteristics" strategies.

	Authors	Data set	Conditioning	$\hat{\pi}$	Z_1	Z_2
1	Lacey (1990)	NFL 1984-1986	Bet on home team after previous home game	0.478	-0.596	
2			Bet on away team after previous away game	0.516	0.435	
3	Vergin (1998)	NFL 1984-1995	Bet on home team after previous home game	0.498	-0.110	
4			Bet on away team after previous away game	0.513	0.692	
5	Y. T. Sung and Tainsky (2014)	NFL 2002-2009	Bet on favorites that had a bye-week	0.625	**2.915**	**2.363**
6			Bet on underdogs that had a bye-week	0.445	-1.144	
7			Bet on away favorites that had a bye-week	0.732	**2.967**	**2.663**
8			Bet on home favorites that had a bye-week	0.448	-0.855	
9	Ashman et al. (2010)	NBA 1990-2009	Bet on home team in the second game of back-to-back games for the away team	0.506	0.873	
10			Bet on home team in the second game of back-to-back games for both teams	0.499	-0.048	
11			Bet on home team in the second game of back-to-back games for the home team	0.459	***-3.086***	-1.312
Rows 12-20 show the strategy "bet on home teams in the 2nd game of back-to-back games for the home team" conditional on the days of rest the away team had (0, 1 or 2, >2) and on whether the home team travelled one or two time zones to the east between their back-to-back games (E) or not (No E).						
12			0	0.499	-0.048	
13			0 E	0.491	-0.330	
14			0 No E	0.501	0.107	
15			1 or 2	0.455	***-3.162***	-1.508
16			1 or 2 E	0.424	**-2.442**	**-1.684**
17			1 or 2 No E	0.463	**-2.298**	-0.827
18			>2	0.486	-0.374	
19			>2 E	0.310	**-2.469**	**-2.163**
20			>2 No E	0.540	0.940	
Rows 21-28 show the strategy "bet on visiting team" conditional on the jet lag (time difference).						
21	Schnyzer and Hizgilov (2018)	Australian Football League 2001-2016	Gain 2+ hours	0.458	-1.414	
22			Gain 1 hour	0.571	1.000	
23			No change	0.454	**-2.547**	-1.221
24			Lose 1 hour	0.533	0.516	
25			Lose 2+ hours	0.431	**-2.256**	-1.477
26			Gain 2+ hours or lose 2+ hours	0.445	**-2.585**	-1.462
27			Lose 2+ hours neutral field	0.545	0.302	
28			Gain 2+ hours or lose 2+ hours neutral field	0.583	0.577	

2.4.1.6 attention & Importance

Another common input to trading strategies is the importance of a game and the attention it receives. Amoako-Adu et al. (1985) and Vergin and Sosik (1999) propose the strategy of betting on all home teams in Monday night NFL games. For a long time, Monday night games were the only games broadcasted in prime time leading to substantial media and fan attention. These spotlights can be a strong incentive for teams to perform better and these games tend to attract more casual bettors. As depicted in Table 8, Monday night home team bets had a win fraction of 68% between 1979-1981 and 60% between 1976-1996, leading to statistically significant profits in both periods at the single test benchmark, but not at the multiple test benchmark (although the null of randomness is rejected at the multiple test benchmark). Shank (2018) extends the strategy as it is now common practice to also broadcast games in prime time on Thursday, Saturday and Sunday night. Although the return of betting on home teams in prime-time games is higher compared to that of the regular Sunday games (compare row 5 and 6 of Table 8), the null of randomness is not rejected.

Table 8: Overview of papers implementing "bet conditional on attention characteristics" strategies.

	Authors	**Data set**	**Conditioning**	$\hat{\pi}$	Z_1	Z_2
1	Amoako-Adu et al. (1985)	NFL 1979-1981	Home team Monday night	0.682	**2.412**	**2.099**
2	Vergin and Sosik (1999)	NFL 1976-1996	Home team Monday night	0.608	***3.837***	**2.997**
3			Home team Monday night underdog	0.667	***3.464***	**2.973**
4			Home team Monday night favorite	0.571	1.801	
5	Shank (2018)	NFL 2009-2017	Home team Prime time	0.520	0.852	
6			Home team Not prime time	0.481	-1.504	
7	Vergin and Sosik (1999)	NFL 1976-1996	Home team Playoff	0.586	**2.304**	**1.665**
8	Gandar et al. (2001)	NFL 1997-1999	Home team Playoff	0.446	-0.930	
9		NBA 1981-1997	Home team Playoff	0.511	0.728	
10	Borghesi (2007b)	NFL 1981-2000	Home team Playoff	0.592	**2.507**	**1.863**
11	Vergin and Sosik (1999)	NFL 1976-1996	Home team Playoff underdog	0.737	**2.065**	**1.859**
12	Gandar et al. (2001)	NFL 1997-1999	Home team Playoff underdog	0.412	-0.728	
13		NBA 1981-1997	Home team Playoff underdog	0.543	1.093	
14	Borghesi (2007b)	NFL 1981-2000	Home team Playoff underdog	0.778	**2.357**	**2.157**
15	Vergin and Sosik (1999)	NFL 1976-1996	Home team Playoff favorite	0.577	1.794	
16	Gandar et al. (2001)	NFL 1997-1999	Home team Playoff favorite	0.456	-0.662	
17		NBA 1981-1997	Home team Playoff favorite	0.505	0.297	
18	Hickman (2020)	March Madness 1996-2019	Bet on the higher-seeded team	0.494	-0.466	
19			Bet on higher-seeded team in round 1	0.497	-0.144	
20			Bet on higher-seeded team in round 2	0.507	0.255	
21			Bet on higher-seeded team in rounds 3-6	0.473	-0.979	
22	Lacey (1990)	NFL 1984-1986	Bet on teams on the week before a divisional game	0.513	0.482	
23			Bet on teams on the week after a divisional game	0.506	0.224	

In a similar vein, Vergin and Sosik (1999), Gandar et al. (2001) and Borghesi (2007b) test betting on playoff games for the NFL and the NBA. These games, like the prime-time games, receive considerably more attention and attract a large amount of casual, less informed, bettors. Furthermore, these games often involve teams that rarely play against each other and can take place on a neutral location, which could complicate the pricing. Lastly, the stakes are especially high in these games as losing teams are eliminated. The strategy of Vergin and Sosik (1999) and Borghesi (2007b) to bet on home teams in playoff games rejects the null of unprofitability at the single test benchmark (but randomness is not rejected at the multiple test benchmark). The similar strategy of betting on home underdogs in playoff games is also significantly profitable at the single test benchmark and has an astounding empirical win fraction of over 70%. Surprised by these results, Gandar et al. (2001) revisit the strategies in a large sample of NBA games, in NFL games beyond the sample used by Vergin and Sosik (1999) and in MLB games (these results are not shown due to the different microstructure of baseball betting). None of these datasets contain evidence that betting on playoff home games is significantly different from randomness, "the bias found by Vergin and Sosik was short-lived" (Gandar et al., 2001, p. 451).

Hickman (2020) focusses on NCAA basketball "March Madness" games. March Madness is a 6-round single-elimination postseason tournament with 64 teams. As these games are played on neutral courts the home/away distinction cannot be made. However, all teams in the tournament are divided into seeds where seed 1 represents the best teams and seed 16 the worst, based on the opinion of a selection committee. Hickman (2020) tests whether consistently betting on the higher-seeded team in March Madness games results in profits as shown in Table 8 (rows 18 to 21). The null of randomness can never be rejected, additional results are shown in appendix, none of them are statistically significant.

We end this section by looking at potential psychological factors that make teams perform differently in the weeks leading up to, or after an important game. Lacey (1990) investigates whether strategies that bet on teams in games before or after divisional games result in profits. Teams might underperform in the week preceding a divisional game as they are already preparing for the divisional game (looking past their opponents) and after a divisional game as a result of the aftermath of a big win or loss. However, as summarized in Table 8 (row 22 and 23), the profits of these strategies do not differ from randomness.

2.4.1.7 absences

Dare, Dennis, and Paul (2015) investigate betting market efficiency in the NBA when players are absent because of for example injury, sickness, suspension or personal reasons. Table 9 shows the results of the strategy of betting on the team with most absences. To further refine the strategy, it is also tested conditional on the value of the payer(s) that is (are) absent, indicated by the Approximate Value (AV) index, which is proportional to the quality of the player (see the paper for more information on how to compute this metric). The results in Table 9 show that betting on teams that miss players wins more than half of the time and the null of randomness can be rejected in 1 case (only at the single test benchmark). The analysis is repeated for home teams and away teams (tables shown in appendix). Home teams with absent players consistently cover more than 50%, although randomness is never rejected. For away teams the evidence is mixed.

Colquitt, Godwin, and Shortridge (2007) investigate the role of coaching changes on betting markets. Inspired by the literature on CEO turnover and subsequent stock price behavior, they investigate whether betting markets react efficiently when a team changes its coach. As shown in Table 9 row 7, when an inexperienced coach takes over, the betting market underestimates the team's ability as these teams cover 63% of the time, which is statistically profitable (again, only at the single test benchmark). The effect fades quickly in the next games. When an experienced coach takes over, randomness is never rejected. Further results with respect to the runup to a coaching change are reported in appendix, the null of randomness is never rejected.

Table 9: Overview of papers implementing "bet conditional on absences" strategies.

	Authors	**Data set**	**Conditioning**	$\hat{\pi}$	Z_1	Z_2
Rows 1-5 show the strategy "bet on team with most absences" conditional on the AV. (AV is the total Approximate Value of the players who are absent. The higher the AV, the more valuable players are missing. When both teams have absences, the AV is the difference between the values of players missing for each team).						
1	Dare et al. (2015)	NBA 1996-2005	unconditional	0.513	1.925	
2			AV≥ 5	0.517	**2.105**	
3			AV ≥ 10	0.511	0.905	
4			AV ≥ 15	0.520	0.887	
5			AV ≥ 20	0.525	0.665	
Rows 6-11 show the strategy "bet on team with a new coach" conditional on the time after the change (1-3 to 7-9 games) and on whether the new coach has previous experience as a NBA head coach (EX) or not (N EX).						
6	Colquitt et al. (2007)	NBA 1988-2002	1-3, EX	0.547	0.808	
7			1-3, N EX	0.631	**2.400**	**1.966**
8			4-6, EX	0.525	0.447	
9			4-6, N EX	0.565	1.193	
10			7-9, EX	0.481	-0.333	
11			7-9, N EX	0.481	-0.333	

2.4.2 Strategies based on past performance

In this section we summarize the large family of both momentum and contrarian trading rules. Momentum strategies extrapolate past performance into the future, while contrarian strategies do just the opposite. These strategies are especially interesting as they are also intensely studied in the mainstream finance literature. Momentum especially is considered to be the "premier anomaly" (Fama & French, 2008, p. 1653). Stocks that have outperformed in the past 3 to 12 months continue to outperform in the near future. Such profitable momentum strategies are awkward as they seem to imply that markets are not even weakly efficient. To make momentum profits compatible with the neoclassical rational framework risks would have to increase after good past performance, which is counterintuitive (Lewellen, 2002) although several attempts for risk based explanations have been made (Galariotis, 2013; Johnson, 2002; Li, 2018). Contrarian strategies found their way into the broader finance literature via seminal work of De Bondt and Thaler (1985). Stocks that have performed relatively well in a 2 to 5-year period relatively underperform the following years and vice versa. This phenomenon is most readily explained by investor overreaction to news which is corrected in the long run but others point to varying risk-premia (K. Chan, 1988; Fama & French, 1996). As the risk-return profile of all spread bets is equal by construction as discussed previously, any risk explanation can be quickly ruled out in our context. If profitable momentum betting strategies would be found, we could more confidently point to behavioral explanations.

2.4.2.1 performance against the spread

A first straightforward strategy consists of betting on the team that beat the spread by the largest average amount in the last *k* weeks. This is the team that outperformed the most, relative to the expectations. If momentum (contrarian) patterns would exist, we would expect this team to overperform (underperform) in the future. As shown in Table 10 rows 1-4, the evidence is mixed. In the early days (1969-1981) the momentum strategy was generally profitable in sample (although randomness was never rejected) while later periods are consistent with a profitable contrarian strategy but the null of unprofitability is only once rejected at the single test benchmark.

A variation on the same theme is not just betting on the one team that outperformed the most, but on all teams that are on win streaks against the spread as shown in Table 10 rows 5-10. The evidence here is mostly consistent with contrarian strategies and the null of randomness is rejected for 1 implementation at the multiple test benchmark.

Table 10: Overview of papers implementing "bet conditional on performance against the spread" strategies.

	Authors	**Data set**	**k=1**	**k=2**	**k=3**	**k=4**	**k=5**	**k=6**
Rows 1-4 show the strategy "each week, bet on team that beat the spread by the largest average amount last k weeks".								
1	Vergin and Scriabin (1978)	NFL 1969-1974	$\hat{\pi}$: 0.526 Z_1: 0.453	$\hat{\pi}$: 0.569 Z_1: 1.179	$\hat{\pi}$: 0.538 Z_1: 0.620	$\hat{\pi}$: 0.627 Z_1: 1.953	$\hat{\pi}$: 0.528 Z_1: 0.412	$\hat{\pi}$: 0.521 Z_1: 0.289
2	Tryfos et al. (1984)	NFL 1969-1981	$\hat{\pi}$: 0.517 Z_1: 0.455	$\hat{\pi}$: 0.541 Z_1: 1.031	$\hat{\pi}$: 0.531 Z_1: 0.742	$\hat{\pi}$: 0.526 Z_1: 0.602	$\hat{\pi}$: 0.537 Z_1: 0.812	$\hat{\pi}$: 0.523 Z_1: 0.475
3	Gandar et al. (1988)	NFL 1980-1985	$\hat{\pi}$: 0.488 Z_1: -0.221	$\hat{\pi}$: 0.416 Z_1: -1.481	$\hat{\pi}$: 0.380 Z_1: **-2.018** Z_2: -1.618	$\hat{\pi}$: 0.409 Z_1: -1.477	$\hat{\pi}$: 0.508 Z_1: 0.130	
4	Vergin (2001)	NFL 1981-1995	$\hat{\pi}$: 0.493 Z_1: -0.204	$\hat{\pi}$: 0.418 Z_1: **-2.328** Z_2: **-1.654**	$\hat{\pi}$: 0.418 Z_1: **-2.255** Z_2: -1.602	$\hat{\pi}$: 0.468 Z_1: -0.841	$\hat{\pi}$: 0.457 Z_1: -1.100	
Rows 5-10 show the strategy "bet on teams that beat the spread last k games".								
5	Lacey (1990)	NFL 1984-1986	$\hat{\pi}$: 0.508 Z_1: 0.398	$\hat{\pi}$: 0.422 Z_1: **-2.795** Z_2: **-1.945**				
6	Oorlog (1995)	NBA 1989-1991	$\hat{\pi}$: 0.500 Z_1: -0.031					
7	Vergin (1998)	NFL 1984-1995	$\hat{\pi}$: 0.500 Z_1: -0.020	$\hat{\pi}$: 0.482 Z_1: -1.109				
8	Paul and Weinbach (2005a)	NBA 1995-2002		$\hat{\pi}$: 0.460 Z_1: ***-4.359*** Z_2: **-1.738**		$\hat{\pi}$: 0.462 Z_1: **-1.981** Z_2: -0.775		
9	Paul, Weinbach, and Humphreys (2011)	NBA 2003-2009		$\hat{\pi}$: 0.499 Z_1: -0.117		$\hat{\pi}$: 0.508 Z_1: 0.500		
10	Sinkey and Logan (2014)	College football 1985-2008		$\hat{\pi}$: 0.500 Z_1: -0.067				

The appendix contains 145 additional implementations based on performance against the spread including Camerer (1989) and Paul, Weinbach, and Humphreys (2014) who further refine the strategies shown in Table 10 by also looking at the performance of the opponents in the last games. The null of randomness is never rejected at the multiple test benchmark.

4.2.2 performance not against the spread

Next to momentum and contrarian strategies against the point spread, it is also common to define strategies relative to past game performance not against the point spread as shown in Table 11. Rows 1-4 implement the strategy of betting on the team that beat its opponents by the largest average amount last weeks. However, for none of the 22 strategies the null of randomness is rejected at the multiple test level.

Fodor, DiFilippo, Krieger, and Davis (2013) implement a longer-term contrarian strategy that exploits the sticky preferences of gamblers. They find that teams that qualified for the playoffs in the prior season are overrated by the market in the first game of the following season. Between seasons, teams can drastically change their lineup, coaches and tactics, which can have a large impact on their subsequent performance. However, gamblers' perceptions are still anchored to the successful campaign of the last season. These sticky preferences can be exploited by betting against teams that qualified for the playoffs last season when they face a team that did not qualify in the first week of the new season as shown in row 5 of Table 11. The strategy results in a win fraction of over 64% and the null of unprofitability is rejected at the single test benchmark, but not at the multiple test benchmark. The effect vanishes as expected in the second week of a new season as gamblers update their beliefs. (The authors also show the results for the strategy beyond game 6, these results are left to the appendix).

In a follow-up study, Bennett (2019) analyzes these sticky preferences in the college football setting. More specifically, he tests whether betting against teams in the top of the Associated Press poll (a prestigious ranking) last season is profitable in the first game of a new season. The results are shown in row 6 of Table 11 and the null of unprofitability can be rejected for top 10 teams, but again only at the single test benchmark. The overvaluation does not exist for the lower ranked teams (11 to 25).

Relatedly, Davis, McElfresh, Krieger, and Fodor (2015) analyze how information of this first game of a new season is used to make decisions related to the second game of a new season. They analyze the performance of the underdog in the second game of a season conditional on the outcome of the first game as shown in Table 11 row 7. In 1 of 5 strategies, statistical efficiency is rejected at the multiple test benchmark. The authors hypothesize that the lack of information (only 1 game played) makes it especially hard to establish efficient point spreads.

Table 11: Overview of papers implementing "bet conditional on performance not against the spread" strategies.

	Authors	Data set	k=1	k=2	k=3	k=4	k=5	k=6
Rows 1-4 show the strategy "each week, bet on team that beat its opponents by the largest average amount last k weeks".								
1	Vergin and Scriabin (1978)	NFL 1969-1974	$\hat{\pi}$: 0.462 Z_1: -0.679	$\hat{\pi}$: 0.563 Z_1: 1.068	$\hat{\pi}$: 0.603 Z_1: 1.638	$\hat{\pi}$: 0.667 Z_1: **2.517** Z_2: **2.160**	$\hat{\pi}$: 0.608 Z_1: 1.540	$\hat{\pi}$: 0.523 Z_1: 0.302
2	Tryfos et al. (1984)	NFL 1969-1981	$\hat{\pi}$: 0.483 Z_1: -0.455	$\hat{\pi}$: 0.525 Z_1: 0.636	$\hat{\pi}$: 0.510 Z_1: 0.249	$\hat{\pi}$: 0.550 Z_1: 1.136	$\hat{\pi}$: 0.538 Z_1: 0.825	$\hat{\pi}$: 0.528 Z_1: 0.583
3	Gandar et al. (1988)	NFL 1980-1985	$\hat{\pi}$: 0.433 Z_1: -1.265	$\hat{\pi}$: 0.392 Z_1: -1.913	$\hat{\pi}$: 0.457 Z_1: -0.717	$\hat{\pi}$: 0.500 Z_1: 0.000	$\hat{\pi}$: 0.483 Z_1: -0.263	
4	Vergin (2001)	NFL 1981-1995	$\hat{\pi}$: 0.509 Z_1: 0.265	$\hat{\pi}$: 0.470 Z_1: -0.844	$\hat{\pi}$: 0.440 Z_1: -1.664	$\hat{\pi}$: 0.443 Z_1: -1.508	$\hat{\pi}$: 0.459 Z_1: -1.038	
Row 5 shows the strategy "bet against teams that qualified for the playoffs last season when they face a team that did not qualify in game k of the next season".								
5	Fodor et al. (2013)	NFL 2004-2012	$\hat{\pi}$: 0.644 Z_1: **2.213** Z_2: **1.850**	$\hat{\pi}$: 0.500 Z_1: 0.000	$\hat{\pi}$: 0.507 Z_1: 0.120	$\hat{\pi}$: 0.475 Z_1: -0.391	$\hat{\pi}$: 0.424 Z_1: -1.231	$\hat{\pi}$: 0.544 Z_1: 0.662
Row 6 shows the strategy "bet on teams in top of AP poll in first game of next season when playing against a team not in the top 25" strategy. The strategy is further conditioned on the top team being the favorite (F) and whether the top team is part of the top 25, top 10 or top 11-25.								
			Top 25	**Top 25 F**	**Top 10**	**Top 10 F**	**Top 11-25**	**Top 11-25 F**
6	Bennett (2019)	College football 2008-2016	$\hat{\pi}$: 0.425 Z_1: **-1.971** Z_2: -1.344	$\hat{\pi}$: 0.436 Z_1: -1.645	$\hat{\pi}$: 0.373 Z_1: **-2.305** Z_2: **-1.873**	$\hat{\pi}$: 0.385 Z_1: **-2.038** Z_2: -1.619	$\hat{\pi}$: 0.471 Z_1: -0.588	$\hat{\pi}$: 0.484 Z_1: -0.308
Row 7 shows the strategy "bet on underdog in the second game of a season conditional on the performance in the first game".								
			Both teams won first game	**Favorite won first game underdog lost**	**Favorite lost first game underdog won**	**Both teams lost first game**	**All games**	
7	Davis et al. (2015)	NFL 1997-2012	$\hat{\pi}$: 0.585 Z_1: 1.236	$\hat{\pi}$: 0.435 Z_1: -1.193	$\hat{\pi}$: 0.463 Z_1: -0.469	$\hat{\pi}$: 0.707 Z_1: ***3.151*** Z_2: **2.792**	$\hat{\pi}$: 0.540 Z_1: 1.234	
Rows 8-10 shows the strategy "bet against teams that won their previous game by k points or more".								
			k=10	**k=15**	**k=20**			
8	Lacey (1990)	NFL 1984-1986	$\hat{\pi}$: 0.539 Z_1: 1.335	$\hat{\pi}$: 0.561 Z_1: 1.714	$\hat{\pi}$: 0.590 Z_1: **1.992** Z_2: 1.467			
9	Vergin (1998)	NFL 1984-1995	$\hat{\pi}$: 0.526 Z_1: 1.809	$\hat{\pi}$: 0.527 Z_1: 1.492	$\hat{\pi}$: 0.524 Z_1: 1.051			
10	Vergin (2001)	NFL 1981-1995	$\hat{\pi}$: 0.522 Z_1: 1.674	$\hat{\pi}$: 0.526 Z_1: 1.584	$\hat{\pi}$: 0.536 Z_1: 1.769			

Lacey (1990), Vergin (1998) and Vergin (2001) test the contrarian strategy of betting against teams that won their previous game by a large margin. Results are shown in the final rows of Table 11. Eight out of nine strategies are profitable in sample although the null of randomness is only once rejected at the single test benchmark and never at the multiple test benchmark. Interestingly, the profitability of the strategies rises almost monotonically with the size of the win in the previous game. Results of the converse strategy of betting on teams that lost by a large margin are shown in appendix. The null of randomness is never rejected.

More strategies can be found in appendix. The null of randomness is never rejected at the multiple test benchmark in 47 additional implementations.

4.2.3 spread movements

We end the past performance discussion by reviewing the strategies that use movements of the point spread as their trading signal. Gandar et al. (1988) propose the strategy of systematically betting on the team that became more of an underdog between the opening and closing line, i.e. the team the market assigns diminishing winning probabilities to. Such a strategy would be profitable if market movements are mainly driven by investor irrationality instead of efficient reactions to news. This strategy of betting against the market is profitable as shown in Table 12 and the null of randomness is quite strongly rejected at the single test benchmark, but again not at the multiple test benchmark.

Gandar, Dare, Brown, and Zuber (1998) and Shank (2018) also implement strategies that look at the difference between opening and closing point spreads. Interestingly, they find that when the home team becomes more of a favorite, its chances of beating the spread go down as shown in Table 12. Conversely, when the home team becomes more of an underdog its chances go up. These are signals that the point spread might overreact to the arrival of news and that gamblers can exploit this by betting in the opposite direction. The null of randomness is rejected once at the multiple test benchmark in 20 line movement strategies shown in Table 12. The appendix contains 26 additional strategies. The null of randomness is never rejected at the multiple test benchmark.

Table 12: Overview of papers implementing “bet conditional on spread movements” strategies

	Authors	Data set	Conditioning	$\hat{\pi}$	Z_1	Z_2
1	Gandar et al. (1988)	NFL 1980-1985	Bet on the team that becomes less favored (more of an underdog) over the course of the week’s betting	0.549	**2.909**	1.503
The next rows show the strategy “bet on home team when the spread for the home team moves by k points”.						
2	Gandar et al. (1998)	NBA 1985-1994	k≤-4	0.464	-0.378	
3			k=-3	0.493	-0.117	
4			k=-2	0.429	**-2.583**	**-1.719**
5			k=-1	0.510	0.689	
6			k=0	0.510	0.886	
7			k=1	0.490	-0.692	
8			k=2	0.480	-0.719	
9			k=3	0.364	**-2.216**	**-1.831**
10			4≤k	0.579	0.688	
11	Shank (2018)	NFL 2009-2017	k≤-4	0.421	-1.192	
12			k≤-3	0.464	-0.895	
13			k≤-2	0.489	-0.388	
14			k≤-1	0.437	***-3.064***	**-1.901**
15			k≤0	0.479	-1.188	
16			k=0	0.495	-0.200	
17			0<k	0.497	-0.179	
18			1≤k	0.498	-0.082	
19			2≤k	0.475	-0.900	
20			3≤k	0.514	0.329	
21			4≤k	0.608	1.540	

2.5 Review

In this section we summarize the reviewed strategies (both the strategies discussed in the main text and those in appendix). Table 13 shows a high-level overview of the effectiveness of the 628 strategy implementations reviewed in this paper. Over 50% were profitable in their sample (i.e. the empirical win fraction fell outside the 47.6%-52.4% interval). Profitable strategies were found in every sport and every strategy family. The null of randomness could be rejected for 18% of the implementations at the single test benchmark, but only for 3% at the multiple test benchmark. The null of unprofitability was rejected for 7% of the strategies at the single test benchmark, but never at the multiple test benchmark. It is worth noting that 40% of all unprofitability rejections at the single test benchmark originate from just three papers (namely Ashman et al. (2010), Paul and Weinbach (2005a) and Vergin and Sosik (1999)). Most rejections both in relative and absolute terms occur in the home underdog category. Also note that the significance of the momentum and contrarian strategies that receive a lot of attention in the mainstream finance literature is very limited.

Figure 1 visualizes all strategy implementations. The top left panel plots the empirical win fractions and the absolute value of the Z_1 statistics. The vast majority of implementations are located in the bottom left or bottom right quadrants which represent the implementations whose track records are indistinguishable from randomness. The implementations in the top right corner are the most interesting. These strategies are profitable in sample and reject the null of randomness at the multiple test benchmark. The red dots represent the implementations that reject the null of unprofitability at the single test benchmark. As mentioned previously, the null of unprofitability was never rejected at the multiple test benchmark. The top right panel of Figure 1 is a funnel plot, a scatter diagram of the empirical win fractions and the square root of the sample size. Funnel plots are often used in meta-analyses to summarize estimates and detect publication bias (Stanley & Doucouliagos, 2010). The funnel is centered right at 50% and shows a clear relationship between profitability and sample size. The strategies tested in the largest samples are unprofitable, the smaller the sample, the higher the likelihood of finding a profitable strategy[27]. This is consistent with an efficient market where deviations from randomness are chance results.

Table 13: General overview of the effectiveness of the reviewed strategies. The second column shows the number of strategy implementations. The third column shows the number of profitable strategies ($\hat{\pi} > 0.524$ or $\hat{\pi} < 0.476$) while columns four and five show the rejections of the null of randomness and the null of unprofitability respectively, both at the single test benchmark and at the multiple test benchmark (between brackets).

Sample	**n**	**Profitable in sample**	**Randomness rejected z>|1.96| (z>|3|)**	**Unprofitability rejected z>|1.64| (z>|3|)**
Full	628	324	113 (17)	45 (0)
Sports				
AFL	4	3	3 (0)	1 (0)
Australian Football League	9	9	4 (0)	0 (0)
College basketball	15	5	2 (0)	0 (0)
College football	87	28	17 (0)	5 (0)
NBA	176	93	36 (8)	19 (0)
March Madness	35	13	0 (0)	0 (0)
NFL	290	164	45 (9)	19 (0)
Preseason NFL	7	5	3 (0)	1 (0)
Preseason NBA	5	4	3 (0)	0 (0)
Strategies				
Home team	20	3	3 (0)	0 (0)
Underdog	73	37	26 (3)	2 (0)
Home underdog	45	29	22 (6)	10 (0)
Home favorite	18	1	0 (0)	0 (0)
Familiarity	18	11	5 (0)	1 (0)
Fatigue	48	35	18 (3)	8 (0)
Attention & Importance	61	31	8 (2)	8 (0)
Absences	24	8	3 (0)	1 (0)
Past performance against the spread	177	85	12 (1)	5 (0)
Past performance not against the spread	95	58	11 (1)	6 (0)
Spread movements	49	26	5 (1)	4 (0)

[27] The symmetry of the funnel is often inspected in meta-studies to detect possible publication bias. If for example only negative effect sizes are published because a negative sign is more intuitive or in line with theory, the funnel will be asymmetric (for example, see Havranek, Irsova, and Zeynalova (2018) on the relationship between tuition fees and college enrollment). In our case symmetry is less of a concern as it is not the sign of the effect that indicates profitability, but the absolute deviation from 50%.

The bottom left and right panels of Figure 1 display histograms of the absolute values of the Z_1 and Z_2 statistics respectively. Interestingly, the number of strategies with test statistics between 0 and 1 is much lower compared to what we would expect under the null.

To conclude the analysis, some of the individual strategies might strongly challenge the notion of market efficiency in sports betting. However, when the evidence is placed in the broader context and we account for the large data mining exercise that has been conducted over the decades, the evidence is consistent with the null of an efficient market. A last argument in favor of market efficiency next to the size of the *z*-statistics is the unpredictability of their signs. There are many examples of *z*-statistics flipping sign when the exact same strategy is tested in another sample. Moreover, sometimes the null of randomness is even rejected in the two opposite directions (see for example the home team strategy). This of course creates a clairvoyance issue with respect to the sign as it is a priori not clear in which direction we should implement the strategy when we want to exploit any bias. This point echoes Fama (1998), who argues that biases in both directions are consistent with the efficient market hypothesis where anomalies are chance results.

2.6 Discussion

Market efficiency in betting markets has been studied for decennia but there is still no clear consensus. The efficiency literature is especially susceptible to data mining issues which stand in the way of more definitive conclusions. It is common practice to devise a battery of strategies based on some easily observable variables without (or only a vague) reference to the underlying logic or psychological mechanisms that would make such strategies a priori interesting to investigate. “What bias are we testing for today?” Sauer (2005, p. 418) somewhat ironically asks when discussing the staleness in the literature, to which we can easily add “which subsample should we investigate today?”.

A general problem for behavioral trading strategies is that “each strategy can be defended persuasively on reasonably plausible a *priori* grounds” (Tryfos et al., 1984, p. 129). Indeed, a momentum strategy betting on teams which have been performing well can sensibly be defended by underreaction. The market does not yet fully appreciate the recent increase in team quality, such that assets on this team can be bought at discount prices. However, the diametrically opposite contrarian strategy of betting on teams that have been performing badly could also sound reasonable if we embed it in a story where the market overreacts and underestimates the true ability of the team. This point echoes the common criticism to the behavioral project: “allowing for irrationality opens a Pandora’s box of ad hoc stories that will have little out-of-sample predictive power” (Daniel et al., 1998, p. 1841). If a sensible story can be made for any strategy it appears that they all deserve to be closely investigated, which induces data mining concerns.

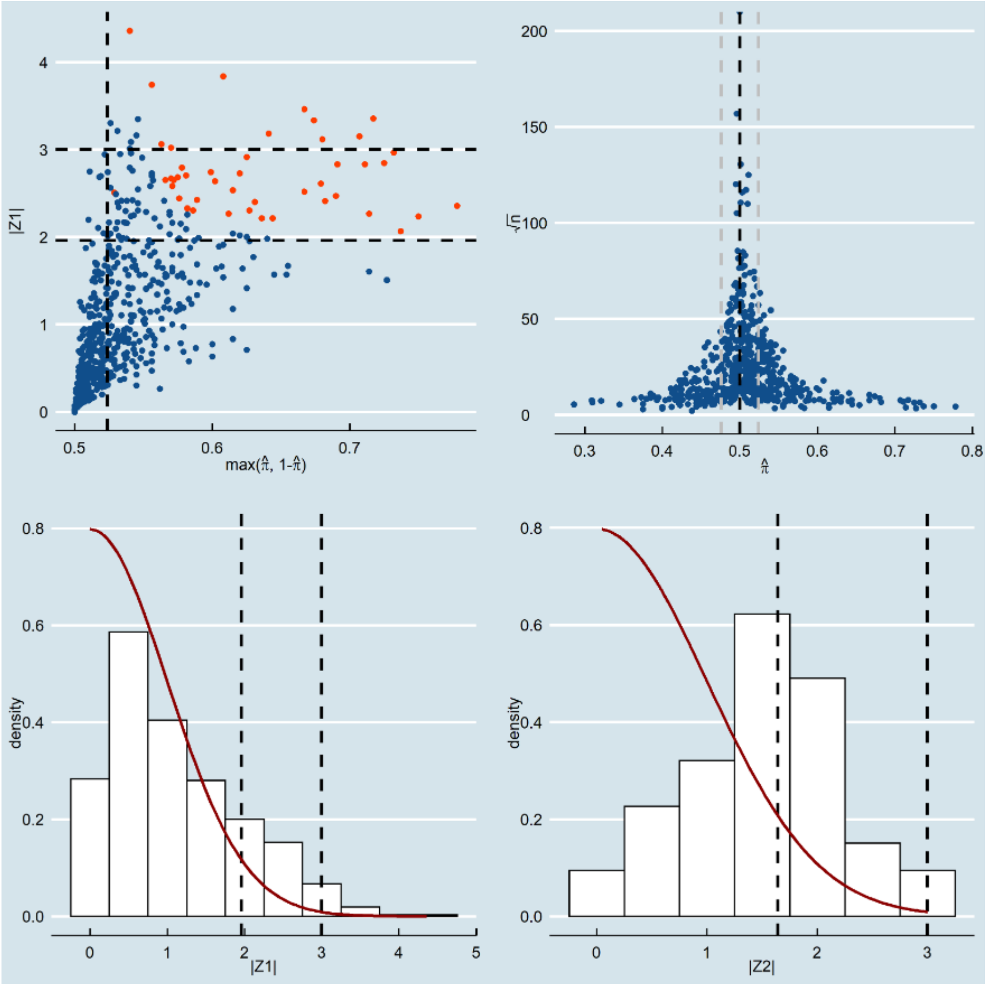

Figure 1: top left panel: scatterplot of the strategy implementations. The horizontal axis represents the empirical win fractions. The black dashed lines represent critical values (1.96 and 3 for the horizontal axis and 0.524 for the vertical axis). The red dots represent the strategy implementations that reject the null of unprofitability under the single test benchmark. Top right panel: funnel plot with the empirical win fractions on the horizontal axis and the square root of the sample size on the vertical axis. Bottom left and right panels show histograms of the absolute value of the Z1 statistics and the Z2 statistics respectively for the full sample of strategy implementations. The vertical blue lines show the critical values, i.e. 0.476 and 0.524 for in-sample profitability and 1.96 and 1.64 for the z-statistics. The red line is the folded standard normal distribution.

In defense of the anomaly dredging endeavors, efficiency requires that all information is properly discounted. Consistently testing any imaginable strategy in any subsample you can get your hands on seems warranted. Such practices can inductively expose unexpected behavioral glitches. However, in these cases, it is vital to properly subject the results to multiple testing methods. If not, we end up with a literature without a clear consensus and profitable strategies which are merely type 1 errors. An issue that is further amplified by the tendency of journals to publish significant results (Harvey, 2017). An interesting area for future research would be to test all proposed strategies both out of sample and post publication. For equities, McLean and Pontiff (2016) and Jacobs and Müller (2020) find that many claimed anomalies disappear over time.

Another interesting area for future research is the origin of the persistent biases. The most frequently used explanation is that the observed regularities are behavioral glitches. However, we should keep in mind that the observed perceived biases might just be rational, a point that is often overlooked in papers that claim to find inefficiencies. “Are we observing an inefficient market or simply one in which the tastes and preferences of the market participants lead to the observed results?” (Gabriel & Marsden, 1990, p. 885). If consumption benefits between betting on favorites and underdogs for example are large enough, rational utility maximizers will be bribed into giving up expected returns, a point that echoes the utility of gambling model (Conlisk, 1993; Humphreys, Paul, & Weinbach, 2013). Although the spread betting microstructure controls for risk-return differences that are expected to drive decision making in a mean-variance framework, agents could also derive consumption benefits from other asset characteristics. Distinguishing between misperceptions (biases) and non-risk-return related consumption benefits (which fit the rational framework) remains empirically difficult, but findings could spill over to the cross-section of expected stock returns (for example to explain the returns on glamour stocks). Consumption benefit differentials driving the decisions of agents in a spread betting context would of course be bad news for the cleanliness of this microstructure as an asset pricing lab. We would again be entangled in a joint-hypothesis problem in the attempt to construct a model that captures the non-risk-return related consumption benefits of the different assets.

2.7 Conclusion

In this review we examine over 600 betting strategies tested over 40 years. We operate in the spread betting context that has the nice characteristic that all assets have the same risk-return profile such that differences in returns between assets or strategies cannot be attributed to risk. Many of the reviewed strategies, when discussed individually, would point in the direction of severe market inefficiencies. However, placing these results in the bigger context takes the sting out.

We document a number of persistent biases, most notably the underdog bias, that could be levered to raise returns of a betting strategy above that of a naïve, random strategy. We find that 3% of the strategy implementations reject the null of randomness under the multiple test benchmark. However, these biases are too small to be profitably exploited. We find that 7% of strategies are significantly profitable under the common single hypothesis benchmark. This could lead researchers to conclude ample profit opportunities exist. However, when we factor in the large number of hypotheses tested over the last decades, we have to move the hurdle rate to at least $|z| > 3$ under which the null of unprofitability is never rejected. Furthermore, we observe a strong inverse relationship between the profitability of a strategy and its sample size, which is again in line with an efficient market where inefficiencies are chance results.

Both data mining and the publication bias most likely lead to more reports of statistically significant trading strategies than actually exist. It is reasonable to assume that our reported profitability rate is an upper bound. The fact that we find no significantly profitable strategies, even with a scientific process that could tilt the evidence in its direction, is a strong argument to not reject the null of market efficiency.

A counter argument that could be made is that successful traders never reveal their secrets. It might well be that the discoverers of highly profitable trading strategies choose to monetize their findings instead of publishing them in a journal. This might lead us to overestimate the true degree of efficiency. However, given the scrutiny betting strategies received over the last decades, it is not very likely that many profitable strategies would go unnoticed.

2.8 Appendix to chapter 2

Game characteristics

Underdog

Table 14 summarizes papers implementing the "bet on the underdog" strategy, conditional on some point spread *PS*. Note that there does not appear to exist a clean relation between the point spread cutoff value and the empirical win fraction, nor statistical significance. However, the tendency to underestimate conditional underdogs is clear as 46 out of 51 implementations indicate underdogs win more than 50% of the time.

Table 14: Overview of papers implementing "bet on underdog conditional on the point spread (PS)" strategy.

Authors	Data set	Conditioning	$\hat{\pi}$	Z_1	Z_2
Vergin and Scriabin (1978)	NFL 1969-1974	PS>5	0.546	**2.388**	1.153
		0<PS≤5	0.456	-1.655	
		5<PS≤10	0.543	1.735	
		10<PS≤15	0.530	0.882	
		15<PS	0.640	**1.980**	**1.645**
Tryfos et al. (1984)	NFL 1969-1981	PS>5	0.540	***3.010***	1.224
		0<PS≤5	0.506	0.410	
		5<PS≤10	0.536	**2.169**	0.709
		10<PS≤15	0.539	1.564	
		15<PS	0.589	1.687	
Gandar et al. (1988)	NFL 1980-1985	PS>5	0.502	0.091	
		0<PS≤5	0.543	**2.248**	0.999
		5<PS≤10	0.511	0.440	
		10<PS≤15	0.537	0.611	
		15<PS	0.375	-0.707	
Vergin (2001)	NFL 1969-1995	PS>5	0.531	***3.215***	0.713
Paul et al. (2003)	College football 1976-2000	7 < PS	0.503	0.467	
		28 < PS	0.538	**2.161**	0.817
Paul and Weinbach (2005b)	College basketball 1996-2004	10 < PS	0.506	0.930	
		20 < PS	0.529	1.652	
Borghesi et al. (2009)	NFL 1981-2004	7 < PS	0.525	1.813	
	College football 1982-2004	7 < PS	0.507	1.163	
	AFL 1998-2006	7 < PS	0.572	**2.653**	**1.776**
Humphreys et al. (2014)	College basketball 2007-2008	10 < PS	0.496	-0.279	
		12 < PS	0.507	0.364	
(Davis & Krieger, 2017)	NFL 1995-2014	PS>3	0.503	0.390	
		0<PS≤3	0.503	0.246	
		PS>5	0.504	0.423	
		0<PS≤5	0.502	0.235	
(Davis & Krieger, 2017)	Preseason NFL 1995-2014	PS>3	0.532	1.469	
		0<PS≤3	0.517	0.914	
		PS>5	0.577	**2.154**	1.492
		0<PS≤5	0.514	0.872	

Table 14 continued: Overview of papers implementing "bet on underdog conditional on the point spread (PS)" strategy.

(Davis & Krieger, 2017)	NBA 2005-2014	PS>3	0.499	-0.095	
		0<PS≤3	0.506	0.613	
		PS>5	0.498	-0.326	
		0<PS≤5	0.505	0.676	
	Preseason NBA 2005-2014	PS>3	0.556	**2.701**	1.557
		0<PS≤3	0.510	0.312	
		PS>5	0.569	**2.287**	1.498
		0<PS≤5	0.529	1.352	
Paul and Weinbach (2005a)	NBA 1995-2002	PS>8.5	0.509	0.800	
		PS>9	0.506	0.555	
		PS>9.5	0.510	0.770	
		PS>10	0.525	1.902	
		PS>10.5	0.530	**2.094**	0.458
		PS>11	0.532	**2.010**	0.494
		PS>11.5	0.537	**2.162**	0.776
		PS>12	0.541	**2.238**	0.941
		PS>12.5	0.556	**2.728**	1.561
		PS>13	0.552	**2.287**	1.247

Home underdog

Table 15: Overview of papers implementing "bet on home underdog" strategy.

Authors	**Data set**	**Conditioning**	$\hat{\pi}$	Z_1	Z_2
Paul et al. (2003)	College football 1976-2000	PS>7	0.517	1.607	
		PS>28	0.571	1.648	
Paul and Weinbach (2005b)	College basketball 1996-2004	PS>10	0.454	**-2.303**	-1.106
Borghesi (2007b)	NFL 1981-2000	PS>2	0.540	**2.792**	1.139
		PS>8	0.547	1.264	
Borghesi et al. (2009)	NFL 1981-2004	PS>7	0.524	0.790	
Borghesi et al. (2009)	College football 1982-2004	PS>7	0.521	**2.041**	
Borghesi et al. (2009)	AFL 1998-2006	PS>7	0.625	**2.000**	1.621
Humphreys et al. (2014)	College basketball 2007-2008	PS>10	0.433	-1.373	
		PS>12	0.417	-1.291	
Shank (2018)	NFL 2009-2017	PS>3	0.493	-0.323	
		PS>6	0.541	1.109	
		PS>10	0.630	1.769	
Paul and Weinbach (2005a)	NBA 1995-2002	PS>8.5	0.545	1.624	
		PS>9	0.555	1.788	
		PS>9.5	0.569	**2.032**	1.330
		PS>10	0.602	**2.639**	**2.028**
		PS>10.5	0.641	***3.182***	**2.646**
		PS>11	0.674	***3.336***	**2.883**
		PS>11.5	0.680	***3.118***	**2.708**
		PS>12	0.717	***3.357***	**2.991**
		PS>12.5	0.711	**2.832**	**2.516**
		PS>13	0.725	**2.846**	**2.548**
Ashman et al. (2010)	NBA 1990-2009	PS>11	0.571	**2.162**	1.437
		PS>12	0.620	**2.729**	**2.191**
Vergin and Sosik (1999)	NFL 1981-1996	PS=0	0.522	0.361	

Table 15 summarizes conditional home underdog strategies. The performance of the strategy in NBA games is striking. In these cases, there exists an almost monotonic relationship between performance and the point spread. For the largest home underdogs, empirical win fractions of over 70% are observed.

Home favorite

For completeness, we mention the results of the strategy of betting on home favorites in Table 16. The null of randomness is never rejected.

Table 16: Overview of papers implementing "bet on home favorite" strategy.

Authors	**Data set**	**Conditioning**	$\hat{\pi}$	Z_1	Z_2
Golec and Tamarkin (1991)	NFL 1973-1987	Unconditional	0.493	-0.642	
	College football 1973-1987	Unconditional	0.495	-0.595	
Oorlog (1995)	NBA 1989-1991	Unconditional	0.486	-1.097	
Vergin and Sosik (1999)	NFL 1981-1996	Unconditional	0.486	-1.358	
Gandar et al. (2001)	NBA 1981-1997	Unconditional	0.496	-0.841	
Y. T. Sung and Tainsky (2014)	NFL 2002-2009	Unconditional	0.484	-1.181	
Paul, Weinbach, and Wilson (2014)	NFL 2007-2011	Unconditional	0.483	-0.989	
Sinkey and Logan (2014)	College football 1985-2008	Unconditional	0.497	-0.513	
Humphreys et al. (2014)	College basketball 2007-2008	Unconditional	0.503	0.275	
Shank (2018)	NFL 2009-2017	Unconditional	0.487	-0.956	
Shank (2019)	NFL 2003-2016	Unconditional	0.488	-1.122	
Humphreys et al. (2014)	College basketball 2007-2008	PS≤ -12	0.486	-0.759	
		PS≤ -10	0.497	-0.164	
Shank (2018)	NFL 2009-2017	PS≤ -10	0.517	0.487	
		PS≤ -7	0.490	-0.445	
		PS≤ -4	0.494	-0.354	
		PS≤ -2	0.485	-0.997	
		PS = 0	0.549	0.700	

Familiarity

Borghesi (2007a) investigates whether temperature information can be profitably exploited. As a first exploration, he computes the empirical win fractions for the home team conditional on the temperature of the game as shown in the first four rows of Table 17. Interestingly, the home team covers significantly more than expected in the coldest games at the single test benchmark, but never at the multiple test benchmark. Rows 5 to 8 of Table 17 contain the strategy of betting on home games in the hottest quartile of game day temperatures conditional on the acclimatization advantage (the converse of the strategy discussed in the main text). Kuester and Sanders (2011) further investigate climate acclimatization challenges. For completeness, we include row 9 and 10 of Table 17 where the subsamples contain games between arid region teams or between humid region teams (so no acclimatization challenges). The strategies are not profitable in sample and the null of randomness is never rejected.

Table 17: Overview of papers implementing "bet on home team conditional on familiarity" strategies.

Authors	**Data set**	**Conditioning**	$\hat{\pi}$	Z_1	Z_2
Borghesi (2007a)	NFL 1981-2004	Q1 temperature (hottest)	0.473	-1.300	
		Q2 temperature	0.483	-1.565	
		Q3 temperature	0.501	0.053	
		Q4 temperature (coldest)	0.541	**2.960**	1.240
		Q1 temperature (hottest) and Q1 acclimatization advantage (highest)	0.475	-0.811	
		Q1 temperature (hottest) and Q2 acclimatization advantage	0.434	-1.543	
		Q1 temperature (hottest) and Q3 acclimatization advantage	0.518	0.381	
		Q1 temperature (hottest) and Q4 acclimatization advantage (lowest)	0.470	-0.492	
Kuester and Sanders (2011)	College football 2000-2006	Both teams from arid regions	0.498	-0.116	
		Both teams from humid regions	0.507	0.749	

Fatigue

Additional bye-week related strategies by Y. T. Sung and Tainsky (2014) and other tests of the "betting on the home team in the second game of back-to-back games when the away team had 1 or 2 days of rest" by Ashman et al. (2010) are shown in Table 18. Furthermore, Oorlog (1995) investigates whether betting on a team playing the last game of a road trip can be profitable. Inefficiencies could arise if the market misestimates the effect of road wear on team performance. Coleman (2017) tests whether betting on a favored home team in the latter half of the season when it hosts a visiting team that travelled one time zone to the east is profitable. This strategy seems promising based on his elaborate regression results. Although the null of randomness is rejected relative to the single test benchmark, the null of unprofitability is not.

Table 18: Overview of papers implementing "bet conditional on fatigue characteristics" strategies.

	Authors	Data set	Conditioning	$\hat{\pi}$	Z_1	Z_2
1	Y. T. Sung and Tainsky (2014)	NFL 2002-2009	home team after it had a bye-week	0.536	0.851	
2			away team after it had a bye-week	0.551	1.063	
3			home favorite after it had a bye-week	0.579	1.539	
4			home underdog after it had a bye-week	0.452	-0.617	
Rows 5-8 display the strategy "bet on the home team in the second game of back-to-back games for the home team when the visiting team had 1 or 2 days of rest" conditional on other info.						
5	Ashman et al. (2010)	NBA 1990-2009	both teams played away last game	0.470	-1.361	
6			the home team played away and away team played at home last game	0.443	**-2.649**	-1.536
7			he home team played at home while the away team played away last game	0.395	-1.889	
8			both teams played at home last game	0.500	0.000	
Rows 9-14 display the strategy "bet home team in back-to-back games when the visiting team is not having back-to-back games" conditional on whether it travelled one or two time zones to the east between back-to-back games (E) or not (No E) and other info.						
9			home team is an underdog	0.430	***-3.021***	**-1.998**
10			the home team is an underdog (E)	0.388	**-2.266**	**-1.785**
11			home team is an underdog (No E)	0.442	**-2.214**	-1.312
12			home team is a favorite	0.470	-1.798	
13			home team is a favorite (E)	0.411	**-2.426**	**-1.781**
14			home team is a favorite (No E)	0.485	-0.784	
Rows 15-18 display the strategy "bet on the home underdog in the second game of back-to-back games for the home team when the visiting team had 1 or 2 days of rest" conditional on other info.						
15			both teams played away last game	0.448	-1.412	
16			home team played away and the away team played at home last game	0.424	**-1.982**	-1.360
17			the home team played at home and the away team played away last game	0.286	**-2.268**	**-2.018**
18			both teams played at home last game	0.500	0.000	
19	Oorlog (1995)	NBA 1989-1991	bet on teams on the last game of a road trip	0.543	1.952	
20	Coleman (2017)	College football 2004-2013	bet on favored home teams in the latter half of the season when they host a visiting team that travelled one time zone to the east	0.554	**1.964**	1.092

Attention & Importance

Hickman (2020) also tests whether the market correctly estimates the quality of the teams per seed. The proposed strategy is to bet on a team when it plays a team from another seed. As shown in Table 19, in none of the 16 cases, the null of randomness is rejected. Furthermore, Hickman (2020) tests whether conference affiliation of the teams can be profitably exploited. A number of variations are shown in Table 19 but randomness can never be rejected.

Relatedly, E. Moore and Francisco (2019) investigate the performance of Power Five (P5)/Automatic Qualifying (AQ) college football teams when playing against a Football Championship Subdivision (FCS) team. The authors dissect the strategy by dividing the P5/AQ sample per conference. The P5/AQ sample includes the Southeastern Conference (SEC), the

Atlantic Coast Conference (ACC), the Big Ten, the Big Twelve and the Pacific 10/Pacific 12 and the Big East until 2012. It is worth noting that the SEC is considered the best conference in college football. Results conditional on the conference are shown in Table 19. Interestingly, the strategy of betting against SEC teams when they play against an FCS team rejects the null of unprofitability (only at the single test benchmark). The authors hypothesize that SEC teams might save their best players for next games when playing against FCS teams, or simply lack motivation.

Table 19: Overview of papers implementing "bet conditional on attention and importance characteristics" strategies.

Authors	**Data set**	**Conditioning**	$\hat{\pi}$	Z_1	Z_2
Hickman (2020)	March	bet on seed 1 teams (against a differently seeded team)	0.507	0.260	
	Madness	bet on seed 2 teams (against a differently seeded team)	0.444	-1.944	
	1996-2019	bet on seed 3 teams (against a differently seeded team)	0.520	0.656	
		bet on seed 4 teams (against a differently seeded team)	0.498	-0.064	
		bet on seed 5 teams (against a differently seeded team)	0.490	-0.283	
		bet on seed 6 teams (against a differently seeded team)	0.497	-0.073	
		bet on seed 7 teams (against a differently seeded team)	0.517	0.447	
		bet on seed 8 teams (against a differently seeded team)	0.536	0.936	
		bet on seed 9 teams (against a differently seeded team)	0.486	-0.329	
		bet on seed 10 teams (against a differently seeded team)	0.500	0.000	
		bet on seed 11 teams (against a differently seeded team)	0.532	0.801	
		bet on seed 12 teams (against a differently seeded team)	0.566	1.578	
		bet on seed 13 teams (against a differently seeded team)	0.462	-0.825	
		bet on seed 14 teams (against a differently seeded team)	0.425	-1.554	
		bet on seed 15 teams (against a differently seeded team)	0.529	0.594	
		bet on seed 16 teams (against a differently seeded team)	0.500	0.000	
		bet on higher seed when it comes from a major conference (ACC, Bog 10, Big 12, Big East, Pac-12, SEC)	0.510	0.558	
		bet on higher seed when lower seed comes from a major conference (ACC, Bog 10, Big 12, Big East, Pac-12, SEC)	0.491	-0.186	
		bet on higher seed when both teams come from a major conference (ACC, Bog 10, Big 12, Big East, Pac-12, SEC)	0.471	-1.318	
		bet on higher seed when both teams do not come from a major conference (ACC, Bog 10, Big 12, Big East, Pac-12, SEC)	0.504	0.086	
		bet on teams from the ACC conference (intra-conference games excluded)	0.447	-1.895	
		bet on teams from the Big 10 conference (intra-conference games excluded)	0.539	1.423	
		bet on teams from the Big 12 conference (intra-conference games excluded)	0.502	0.057	
		bet on teams from the Big East conference (intra-conference games excluded)	0.511	0.384	
		bet on teams from the Pac-12 conference (intra-conference games excluded)	0.525	0.768	
		bet on teams from the SEC conference (intra-conference games excluded)	0.507	0.236	
		bet on the higher-seeded team when PS≤ -20	0.478	-0.470	
		bet on the higher-seeded team when -20 <PS≤ -10	0.500	0.000	
		bet on the higher-seeded team when -10 <PS≤ -5	0.499	-0.045	
		bet on the higher-seeded team when -5 <PS≤ 0	0.480	-0.839	

Table 19 continued: Overview of papers implementing "bet conditional on attention and importance characteristics" strategies.

		bet on the higher-seeded team when PS>0	0.527	0.612	
E. Moore and Francisco (2019)	College football 2003-2018	bet on P5/AQ teams when playing an FCS team	0.499	-0.052	
		bet on ACC teams when playing an FCS team	0.568	1.279	
		bet on Big 10 teams when playing an FCS team	0.533	0.516	
		bet on Big 12 teams when playing an FCS team	0.473	-0.405	
		bet on Big East teams when playing an FCS team	0.455	-0.522	
		bet on PAC 10/12 teams when playing an FCS team	0.520	0.283	
		bet on SEC teams when playing an FCS team	0.385	**-2.535**	**-2.011**

Absences

Dare et al. (2015) further condition their strategy on home teams and away teams respectively as shown in Table 20.

To investigate how the market deals with potential rumors on coaching changes, Colquitt et al. (2007) also investigate the runup to the change. As shown in Table 20, there is no evidence betting markets are not efficient in the games leading up to a coaching change.

Table 20: Overview of papers implementing "bet conditional on absence characteristics" strategies.

Authors	Data set	Conditioning	$\hat{\pi}$	Z_1	Z_2
Dare et al. (2015)	NBA 1996-2005	bet on home teams with the most absences	0.511	1.122	
		bet on home teams with the most absences AV ≥ 5	0.509	0.798	
		bet on home teams with the most absences AV ≥ 10	0.517	0.949	
		bet on home teams with the most absences AV ≥ 15	0.558	1.830	
		bet on home teams with the most absences AV ≥ 20	0.543	0.778	
		bet on away teams with the most absences	0.485	-1.602	
		bet on away teams with the most absences AV ≥ 5	0.475	**-2.167**	-0.101
		bet on away teams with the most absences AV ≥ 10	0.494	-0.347	
		bet on away teams with the most absences AV ≥ 15	0.518	0.561	
		bet on away teams with the most absences AV ≥ 20	0.490	-0.198	
Colquitt et al. (2007)	NBA 1988-2002	bet on the team that will hire a new coach games 1-3 before change	0.500	0.000	
		bet on the team that will hire a new coach games 4-6 before change	0.481	-0.477	
		bet on the team that will hire a new coach games 7-9 before change	0.491	-0.236	

Performance against the spread

Camerer (1989) and Paul, Weinbach, and Humphreys (2014) further refine their strategies by conditioning on the performance of the opposing team against the spread. Table 21 shows the strategy of betting on teams on win streaks while Table 22 shows the opposite strategy of betting on teams that are on losing streaks. The evidence is mixed, i.e. the empirical win fractions are not consistently smaller or larger than 50%, furthermore, the null of randomness is only once rejected at the single test benchmark.

Table 21: Overview of papers implementing "bet on teams that are on a k game winning streak against the spread when playing a team on a shorter winning streak/losing streak against the spread" strategy.

Authors	**Data set**	**Conditioning**	$\hat{\pi}$	Z_1	Z_2
Camerer (1989)	NBA 1983-1986	k=1, shorter winning streak	0.520	0.532	
		k=2, shorter winning streak	0.510	0.381	
		k=3, shorter winning streak	0.466	-1.104	
		k=4, shorter winning streak	0.459	-1.031	
		k=5, shorter winning streak	0.461	-0.792	
		k=6, shorter winning streak	0.426	-1.152	
		k=7, shorter winning streak	0.476	-0.309	
		k=8, shorter winning streak	0.345	-1.671	
		k ≥ 9, shorter winning streak	0.421	-0.973	
		k=1, losing streak	0.520	0.532	
		k=2, losing streak	0.520	0.523	
		k=3, losing streak	0.452	-1.159	
		k=4, losing streak	0.529	0.542	
		k=5, losing streak	0.480	-0.283	
		k=6, losing streak	0.514	0.169	
		k=7, losing streak	0.556	0.471	
		k=8, losing streak	0.400	-0.632	
		k ≥ 9, losing streak	0.444	-0.471	
Paul, Weinbach, and Humphreys (2014)	NFL 2005-2010	k=1, shorter winning streak	0.495	-0.164	
		k=2, shorter winning streak	0.490	-0.277	
		k=3, shorter winning streak	0.495	-0.097	
		k=4, shorter winning streak	0.520	0.283	
		k=5, shorter winning streak	0.478	-0.295	
		k=1, losing streak	0.492	-0.282	
		k=2, losing streak	0.519	0.434	
		k=3, losing streak	0.393	-1.664	
		k=4, losing streak	0.500	0.000	
		k=5, losing streak	0.455	-0.426	

Table 22: Overview of papers implementing "bet on teams that are on a k game losing streak against the spread when playing a team on a shorter losing streak/winning streak against the spread" strategy.

Authors	**Data set**	**Conditioning**	$\hat{\pi}$	Z_1	Z_2
Camerer (1989)	NBA 1983-1986	k=1, shorter losing streak	0.532	0.836	
		k=2, shorter losing streak	0.519	0.667	
		k=3, shorter losing streak	0.520	0.635	
		k=4, shorter losing streak	0.538	0.955	
		k=5, shorter losing streak	0.449	-1.010	
		k=6, shorter losing streak	0.596	1.457	
		k=7, shorter losing streak	0.444	-0.667	
		k=8, shorter losing streak	0.750	**2.236**	**2.025**
		k ≥ 9, shorter losing streak	0.615	1.177	
		k=1, winning streak	0.532	0.836	
		k=2, winning streak	0.532	0.836	
		k=3, winning streak	0.519	0.440	
		k=4, winning streak	0.536	0.655	
		k=5, winning streak	0.500	0.000	
		k=6, winning streak	0.654	1.569	
		k=7, winning streak	0.400	-0.775	
		k=8, winning streak	0.727	1.508	
		k ≥ 9, winning streak	0.615	0.832	
Paul, Weinbach, and Humphreys (2014)	NFL 2005-2010	k=1, shorter losing streak	0.491	-0.325	
		k=2, shorter losing streak	0.551	1.569	
		k=3, shorter losing streak	0.550	1.044	
		k=4, shorter losing streak	0.558	0.832	
		k=5, shorter losing streak	0.654	1.569	
		k=1, winning streak	0.491	-0.333	
		k=2, winning streak	0.529	0.676	
		k=3, winning streak	0.536	0.535	
		k=4, winning streak	0.556	0.577	
		k=5, winning streak	0.714	1.604	

Table 23 contains additional performance against the spread strategies of which we highlight a few. Woodland and Woodland (2000) and Sinkey and Logan (2014) investigate whether profitable strategies can be found at the intersection between past performance against the spread and other game variables (favorite/underdog or home/away). Vergin (2001) tests whether the performance against the spread in a previous season contains useful information. Kochman, Goodwin, and Gilliam (2017) test whether teams that have a very lopsided record against the spread in the beginning of the season regress to the mean in terms of performance against the spread. More specifically they propose the strategy of betting on all teams that lost at least 4 out of the 5 first games against the spread and betting against all teams that won at least 4 of their first 5 games against the spread. The null of randomness is never rejected.

Table 23: Overview of papers implementing "bet conditional on performance against the spread characteristics" strategies.

Authors	**Data set**	**Conditioning**	$\hat{\pi}$	Z_1	Z_2
Vergin and Scriabin (1978)	NFL 1969-1974	bet on teams that had a winning point spread record the year before	0.506	0.322	
Vergin and Scriabin (1978)	NFL 1969-1974	bet against teams that had a losing point spread record the year before	0.495	-0.266	
Gandar et al. (1988)	NFL 1980-1985	bet the underdog against a favored team that, as a favorite in the previous week, covered the spread by at least 10 points	0.581	**2.089**	1.476
Lacey (1990)	NFL 1984-1986	bet on teams that failed to beat the spread last two games	0.425	**-2.683**	**-1.834**
Vergin (1998)	NFL 1984-1995	bet on teams that failed to beat the spread last two games	0.507	0.423	
Sinkey and Logan (2014)	College football 1985-2008	bet on teams that failed to beat the spread last two games	0.495	-0.687	
Oorlog (1995)	NBA 1989-1991	bet on teams that have a better win record against the spread for the season to date	0.512	1.111	
Oorlog (1995)	NBA 1989-1991	in the second half of the season, bet on the team with the better win record against the spread in the first half of the season	0.510	0.643	
Woodland and Woodland (2000)	NFL 1985-1997	bet against favorite teams that covered last game	0.507	0.545	
Woodland and Woodland (2000)	NFL 1985-1997	bet against favorite teams that covered at least 2 consecutive games	0.501	0.077	
Woodland and Woodland (2000)	NFL 1985-1997	bet against favorite teams that covered at least 3 consecutive games	0.498	-0.055	
Woodland and Woodland (2000)	NFL 1985-1997	bet against favorite teams that covered at least 4 consecutive games	0.535	0.884	
Woodland and Woodland (2000)	NFL 1985-1997	bet on underdog teams that failed to cover last game	0.525	1.864	
Woodland and Woodland (2000)	NFL 1985-1997	bet on underdog teams that failed to cover at least 2 consecutive games	0.523	1.238	
Woodland and Woodland (2000)	NFL 1985-1997	bet on underdog teams that failed to cover at least 3 consecutive games	0.515	0.579	
Woodland and Woodland (2000)	NFL 1985-1997	bet on underdog teams that failed to cover at least 4 consecutive games	0.503	0.074	
Woodland and Woodland (2000)	NFL 1985-1997	bet against favorite teams that won and covered last game	0.511	0.777	

Table 23 continued: Overview of papers implementing "bet conditional on performance against the spread characteristics" strategies.

Woodland and Woodland (2000)	NFL 1985-1997	bet against favorite teams that won and covered at least 2 consecutive games	0.521	0.994	
Woodland and Woodland (2000)	NFL 1985-1997	bet against favorite teams that won and covered at least 3 consecutive games	0.469	-0.927	
Woodland and Woodland (2000)	NFL 1985-1997	bet against favorite teams that won and covered at least 4 consecutive games	0.480	-0.400	
Woodland and Woodland (2000)	NFL 1985-1997	bet on underdog teams that lost or tied and failed to cover last game	0.527	1.844	
Woodland and Woodland (2000)	NFL 1985-1997	bet on underdog teams that lost or tied and failed to cover at least 2 consecutive games	0.526	1.233	
Woodland and Woodland (2000)	NFL 1985-1997	bet on underdog teams that lost or tied and failed to cover at least 3 consecutive games	0.504	0.128	
Woodland and Woodland (2000)	NFL 1985-1997	bet on underdog teams that lost or tied and failed to cover at least 4 consecutive games	0.495	-0.097	
Vergin (2001)	NFL 1981-1995	bet against teams that covered the spread by 10 points or more last game	0.513	0.930	
Vergin (2001)	NFL 1981-1995	bet against teams that covered the spread by 15 points or more last game	0.518	1.010	
Vergin (2001)	NFL 1981-1995	bet against teams that covered the spread by 20 points or more last game	0.526	1.107	
Vergin (2001)	NFL 1981-1995	bet on teams that failed to cover the spread by 10 points or more last game	0.511	0.825	
Vergin (2001)	NFL 1981-1995	bet on teams that failed to cover the spread by 15 points or more last game	0.531	1.704	
Vergin (2001)	NFL 1981-1995	bet on teams that failed to cover the spread by 20 points or more last game	0.525	1.067	
Vergin (2001)	NFL 1981-1995	bet against teams that had a net winning record against the spread of at least 4 games last season	0.537	1.630	
Vergin (2001)	NFL 1981-1995	bet against teams that had a net winning record against the spread of at least 5 games last season	0.542	1.660	
Vergin (2001)	NFL 1981-1995	bet against teams that had a net winning record against the spread of at least 6 games last season	0.517	0.484	
Vergin (2001)	NFL 1981-1995	bet against teams that had a net winning record against the spread of at least 7 games last season	0.567	1.555	
Vergin (2001)	NFL 1981-1995	bet against teams that had a net winning record against the spread of at least 8 games last season	0.558	0.762	
Vergin (2001)	NFL 1981-1995	bet on teams that had a net losing record against the spread of at least 4 games last season	0.497	-0.134	
Vergin (2001)	NFL 1981-1995	bet on teams that had a net losing record against the spread of at least 5 games last season	0.513	0.504	
Vergin (2001)	NFL 1981-1995	bet on teams that had a net losing record against the spread of at least 6 games last season	0.561	1.480	
Vergin (2001)	NFL 1981-1995	bet on teams that had a net losing record against the spread of at least 7 games last season	0.560	1.342	
Vergin (2001)	NFL 1981-1995	bet on teams that had a net losing record against the spread of at least 8 games last season	0.622	1.640	
Vergin (2001)	NFL 1981-1995	each week, bet on the team that lost against the spread by the largest average amount last week	0.505	0.137	
Vergin (2001)	NFL 1981-1995	each week, bet on the team that lost against the spread by the largest average amount last 2 weeks	0.500	0.000	
Vergin (2001)	NFL 1981-1995	each week, bet on the team that lost against the spread by the largest average amount last 3 weeks	0.531	0.866	

Table 23 continued: Overview of papers implementing "bet conditional on performance against the spread characteristics" strategies.

Vergin (2001)	NFL 1981-1995	each week, bet on the team that lost against the spread by the largest average amount last 4 weeks	0.445	-1.445	
Vergin (2001)	NFL 1981-1995	each week, bet on the team that lost against the spread by the largest average amount last 5 weeks	0.529	0.723	
Paul and Weinbach (2005a)	NBA 1995-2002	bet against teams that are not on a >2 game losing streak against the spread versus teams on >2 game losing streaks against the spread	0.514	1.530	
Paul and Weinbach (2005a)	NBA 1995-2002	bet against teams that are not on a >4 game losing streak against the spread versus teams on >4 game losing streaks against the spread	0.512	0.640	
Paul et al. (2011)	NBA 2003-2009	bet against teams on a 2-game loss streak against the spread	0.498	-0.235	
Paul et al. (2011)	NBA 2003-2009	bet against teams on a 4-game loss streak against the spread	0.513	0.848	
Sinkey and Logan (2014)	College football 1985-2008	bet on home teams that beat the spread last two games	0.516	1.680	
Sinkey and Logan (2014)	College football 1985-2008	bet on underdogs that beat the spread last two games	0.489	-1.209	
Sinkey and Logan (2014)	College football 1985-2008	bet on home favorites that beat the spread last two games	0.498	-0.139	
Sinkey and Logan (2014)	College football 1985-2008	bet on home underdogs that beat the spread last two games	0.524	1.414	
Sinkey and Logan (2014)	College football 1985-2008	bet on home teams that failed to beat the spread last two games	0.488	-1.300	
Sinkey and Logan (2014)	College football 1985-2008	bet on underdogs that failed to beat the spread last two games	0.483	-1.604	
Sinkey and Logan (2014)	College football 1985-2008	bet on home favorites that failed to beat the spread last two games	0.477	-1.753	
Sinkey and Logan (2014)	College football 1985-2008	bet on home underdogs that failed to beat the spread last two games	0.507	0.446	
Kochman et al. (2017)	College football 2015-2016	bet against teams that won at least 4 of the first five games against the spread	0.525	0.632	
Kochman et al. (2017)	College football 2015-2016	bet on teams that lost at least 4 of the first five games against the spread	0.533	0.851	
Shank (2018)	NFL 2009-2017	bet on the home team if it covered the spread last two games	0.481	-0.799	
Shank (2018)	NFL 2009-2017	bet on the home team if it failed to cover the spread last two games	0.530	1.342	
Shank (2018)	NFL 2009-2017	bet on the away team if it covered the spread last two games	0.495	-0.190	
Shank (2018)	NFL 2009-2017	bet on the away team if it failed to cover the spread last two games	0.460	-1.785	

Table 23 continued: Overview of papers implementing "bet conditional on performance against the spread characteristics" strategies.

Bennett (2020)	College football 2006-2018	for BCS/Power 5 teams, bet on teams that that exceeded the point spread by 20 points or more and betting against teams that fell short by 20 points or more in the previous game	0.501	0.052	
Bennett (2020)	College football 2006-2018	for non BCS/Power 5 teams, bet on teams that that exceeded the point spread by 20 points or more and betting against teams that fell short by 20 points or more in the previous game	0.534	**2.416**	0.702
Bennett (2020)	College football 2006-2018	for BCS/Power 5 teams and non-BCS/Power 5 teams that played a BCS/Power 5 team, bet on teams that that exceeded the point spread by 20 points or more and betting against teams that fell short by 20 points or more in the previous game	0.500	-0.024	
Bennett (2020)	College football 2006-2018	for non-BCS/Power 5 teams that played another non BCS/Power 5 team, bet on teams that that exceeded the point spread by 20 points or more and betting against teams that fell short by 20 points or more in the previous game	0.542	**2.759**	1.207

Bennett (2020) implements strategies that condition on last game performance of both teams. The rule is to bet on teams that did well against the spread in the previous game and to bet against teams that performed poorly against the spread when they play teams whose results were closer to the spread last game. The strategy is tested for different parameter values and shown in Table 24. The rows condition on the difference between the spread and the actual outcome of a team in its prior game. The columns indicate the result against the spread of its opponent in its own previous game. For example, in the cell with row header **≥35** and column header **<35**, the betting rule is implemented on teams where the difference between the outcome and point spread in the previous game was 35 points or more, while the difference for the opponent was smaller than 35 in its previous game. In only 1 of 22 tests, the null of randomness is rejected at the single test benchmark.

Table 24: Bennett (2020) in college football games between 2006-2018. Strategy implemented is "bet on teams that did well against the spread in the previous game and bet against teams that performed poorly against the spread when they play teams whose results were closer to the spread in their previous game" strategy. The rows condition on the difference between the spread and the actual outcome of a team in its prior game. The columns indicate the results against the spread of its opponent in their previous game.

	<35	**<30**	**<25**	**<20**	**<15**	**<10**	**<5**
≥35	$\hat{\boldsymbol{\pi}}$: 0.511 $\boldsymbol{Z_1}$: 0.465	$\hat{\boldsymbol{\pi}}$: 0.509 $\boldsymbol{Z_1}$: 0.381	$\hat{\boldsymbol{\pi}}$: 0.508 $\boldsymbol{Z_1}$: 0.340	$\hat{\boldsymbol{\pi}}$: 0.507 $\boldsymbol{Z_1}$: 0.258	$\hat{\boldsymbol{\pi}}$: 0.505 $\boldsymbol{Z_1}$: 0.172	$\hat{\boldsymbol{\pi}}$: 0.514 $\boldsymbol{Z_1}$: 0.412	$\hat{\boldsymbol{\pi}}$: 0.539 $\boldsymbol{Z_1}$: 0.839
≥30	/	$\hat{\boldsymbol{\pi}}$: 0.497 $\boldsymbol{Z_1}$: -0.166	$\hat{\boldsymbol{\pi}}$: 0.501 $\boldsymbol{Z_1}$: 0.068	$\hat{\boldsymbol{\pi}}$: 0.502 $\boldsymbol{Z_1}$: 0.109	$\hat{\boldsymbol{\pi}}$: 0.498 $\boldsymbol{Z_1}$: -0.119	$\hat{\boldsymbol{\pi}}$: 0.516 $\boldsymbol{Z_1}$: 0.697	$\hat{\boldsymbol{\pi}}$: 0.521 $\boldsymbol{Z_1}$: 0.686
≥25	/	/	$\hat{\boldsymbol{\pi}}$: 0.515 $\boldsymbol{Z_1}$: 1.202	$\hat{\boldsymbol{\pi}}$: 0.516 $\boldsymbol{Z_1}$: 1.244	$\hat{\boldsymbol{\pi}}$: 0.512 $\boldsymbol{Z_1}$: 0.850	$\hat{\boldsymbol{\pi}}$: 0.523 $\boldsymbol{Z_1}$: 1.422	$\hat{\boldsymbol{\pi}}$: 0.518 $\boldsymbol{Z_1}$: 0.828
≥20	/	/	/	$\hat{\boldsymbol{\pi}}$: 0.516 $\boldsymbol{Z_1}$: 1.693	$\hat{\boldsymbol{\pi}}$: 0.515 $\boldsymbol{Z_1}$: 1.489	$\hat{\boldsymbol{\pi}}$: 0.524 $\boldsymbol{Z_1}$: **2.012** $\boldsymbol{Z_2}$: 0.049	$\hat{\boldsymbol{\pi}}$: 0.528 $\boldsymbol{Z_1}$: 1.774

Performance not against the spread

Table 25 contains additional strategies based on performance not against the spread. Many of the strategies are similar to those discussed above, but the past information is now measured by the game outcome itself and not against the spread. The null of randomness is never rejected at the multiple test benchmark.

Table 25: Overview of papers implementing "bet conditional on performance not against the spread characteristics" strategies.

Authors	Data set	Conditioning	$\hat{\pi}$	Z_1	Z_2
Lacey (1990)	NFL 1984-1986	bet on teams that qualified for post season play last season when facing a team that did not	0.550	1.825	
Vergin (1998)	NFL 1984-1995	bet on teams that qualified for post season play last season when facing a team that did not	0.486	-0.901	
Vergin (2001)	NFL 1981-1995	bet on teams that qualified for post season play last season when facing a team that did not	0.503	0.256	
Fodor et al. (2013)	NFL 2004-2012	bet on teams that qualified for post season play last season when facing a team that did not	0.496	-0.276	
Woodland and Woodland (2000)	NFL 1985-1997	bet against favorite teams that won last game	0.520	1.583	
Woodland and Woodland (2000)	NFL 1985-1997	bet against favorite teams that won at least 2 consecutive games	0.532	1.874	
Woodland and Woodland (2000)	NFL 1985-1997	bet against favorite teams that won at least 3 consecutive games	0.510	0.456	
Woodland and Woodland (2000)	NFL 1985-1997	bet against favorite teams that won at least 4 consecutive games	0.515	0.492	
Woodland and Woodland (2000)	NFL 1985-1997	bet on underdog teams that lost or tied last game	0.526	**2.039**	0.170
Woodland and Woodland (2000)	NFL 1985-1997	bet on underdog teams that lost or tied at least 2 consecutive games	0.523	1.350	
Woodland and Woodland (2000)	NFL 1985-1997	bet on underdog teams that lost or tied at least 3 consecutive games	0.525	1.147	
Woodland and Woodland (2000)	NFL 1985-1997	bet on underdog teams that lost or tied at least 4 consecutive games	0.536	1.254	
Vergin (2001)	NFL 1981-1995	bet on teams that lost their previous game by 10 points or more	0.498	-0.155	
Vergin (2001)	NFL 1981-1995	bet on teams that lost their previous game by 15 points or more	0.511	0.693	
Vergin (2001)	NFL 1981-1995	bet on teams that lost their previous game by 20 points or more	0.522	1.065	
Vergin (2001)	NFL 1981-1995	each week, bet on the team has been outscored by its opponents by the largest average amount last week	0.500	0.000	
Vergin (2001)	NFL 1981-1995	each week, bet on the team has been outscored by its opponents by the largest average amount last 2 weeks	0.471	-0.840	
Vergin (2001)	NFL 1981-1995	each week, bet on the team has been outscored by its opponents by the largest average amount last 3 weeks	0.479	-0.583	

Table 25 continued: Overview of papers implementing "bet conditional on performance not against the spread characteristics" strategies.

Vergin (2001)	NFL 1981-1995	each week, bet on the team has been outscored by its opponents by the largest average amount last 4 weeks	0.503	0.076	
Vergin (2001)	NFL 1981-1995	each week, bet on the team has been outscored by its opponents by the largest average amount last 5 weeks	0.490	-0.239	
Paul et al. (2011)	NBA 2003-2009	bet on teams on a 2-game win streak	0.498	-0.262	
Paul et al. (2011)	NBA 2003-2009	bet on teams on a 4-game win streak	0.496	-0.307	
Paul et al. (2011)	NBA 2003-2009	bet against teams on a 2-game loss streak	0.504	0.503	
Paul et al. (2011)	NBA 2003-2009	bet against teams on a 4-game loss streak	0.502	0.150	
The rows below show the strategy "bet against teams that qualified for the playoffs last season when they face a team that did not qualify in game k of the next season"					
Authors	**Data set**	**Conditioning**	$\hat{\pi}$	Z_1	Z_2
Fodor et al. (2013)	NFL 2004-2012	k = 7	0.516	0.254	
		k = 8	0.404	-1.457	
		k = 9	0.473	-0.405	
		k = 10	0.475	-0.384	
		k = 11	0.469	-0.500	
		k = 12	0.492	-0.126	
		k = 13	0.564	0.944	
		k = 14	0.507	0.119	
		k = 15	0.508	0.128	
		k = 16	0.444	-0.882	
		k = 17	0.620	**2.018**	1.618

Table 26: Overview of papers implementing "bet on teams in top of AP poll in first game of next season when playing against a team not in the top 25" strategy. The strategy is further conditioned on the team being the favorite (F) (which is of course often the case for last season top 25 teams) and playing against a power 5 team (P) or not (N).

Authors	**Data set**	**Top 25 P**	**Top 25 F P**	**Top 25 N**	**Top 25 F N**	**Top 10 P**	**Top 10 F P**
Bennett (2019)	College football 2008-2016	$\hat{\pi}$: 0.458 Z_1: -0.577	$\hat{\pi}$: 0.500 Z_1: 0.000	$\hat{\pi}$: 0.413 Z_1: -1.960 Z_2: -1.427	$\hat{\pi}$: 0.416 Z_1: -1.878	$\hat{\pi}$: 0.500 Z_1: 0.000	$\hat{\pi}$: 0.520 Z_1: 0.200
		Top 10 N	**Top 10 N F**	**Top 11-25 P**	**Top 11-25 F P**	**Top 11-25 N**	**Top 11-25 N F**
		$\hat{\pi}$: 0.309 Z_1: **-2.832** Z_2: **-2.481**	$\hat{\pi}$: 0.321 Z_1: **-2.610** Z_2: **-2.266**	$\hat{\pi}$: 0.448 Z_1: -0.557	$\hat{\pi}$: 0.500 Z_1: 0.000	$\hat{\pi}$: 0.480 Z_1: -0.346	$\hat{\pi}$: 0.480 Z_1: -0.346

Spread movements

Table 27 supplements the strategy discussed in the spread movements section of the main text. The null of randomness is never rejected. Baryla Jr, Borghesi, Dare, and Dennis (2007) zoom in on the efficiency of the betting market during the first four games of a season. They compare early season price formation with that of the IPO process banks face when pricing a new security. At the start of a season, the betting market has some indications about the strength of a team, but true values are only revealed gradually as the season progresses. More concretely, they test whether movements in the point spread between the opening line and closing line contain useful information in the first four games of a season. As shown in Table 27, the null of randomness is never rejected.

Table 27: Overview of papers implementing "bet conditional on spread movements" strategies

	Authors	**Data set**	**Conditioning**	$\hat{\pi}$	Z_1	Z_2
Rows 1-8 show the strategy "bet on home team when the spread for the home team moves by k points".						
1	Gandar et al. (1998)	NBA 1985-1994	k = -3.5	0.433	-0.730	
2			k = -2.5	0.536	0.805	
3			k = -1.5	0.490	-0.469	
4			k = -0.5	0.500	0.000	
5			k = 0.5	0.490	-0.819	
6			k = 1.5	0.520	0.986	
7			k = 2.5	0.563	1.463	
8			k = 3.5	0.467	-0.365	
Rows 9-25 show the strategy "bet on home team when the spread moved by k points from the opening line to the closing line in the first four home games of a season".						
9	Baryla Jr et al. (2007)	NBA 1985-2005	k ≤ -4	0.488	-0.152	
10			k ≤ -3.5	0.423	-0.784	
11			k ≤ -3	0.375	-1.414	
12			k ≤ -2.5	0.520	0.283	
13			k ≤ -2	0.440	-1.153	
14			k ≤ -1.5	0.422	-1.584	
15			k ≤ -1	0.510	0.280	
16			k ≤ -0.5	0.504	0.124	
17			k ≤ 0	0.544	1.677	
18			k ≤ 0.5	0.506	0.197	
19			k ≤ 1	0.448	-1.405	
20			k ≤ 1.5	0.571	1.604	
21			k ≤ 2	0.456	-0.887	
22			k ≤ 2.5	0.424	-1.172	
23			k ≤ 3	0.589	1.336	
24			k ≤ 3.5	0.458	-0.577	
25			k ≤ 4	0.500	0.000	
26	Gandar et al. (1988)	NFL 1980-1985	bet on the team that becomes less favored (more of an underdog) over the course of the week's betting for games in weeks following "winning" weeks for the public. "Winning" weeks were those for which at least 50% of line changes from the opening to the closing line moved the betting line closer to the eventual game outcome	0.570	**2.669**	**1.762**

Chapter 3: How the Internet Can Shape Markets: the Case of Sports Betting

Abstract

The rise of the internet impacts many vested businesses through e.g. decreases in marginal costs and increases in informed customers. However, measuring the impact of digitalization remains difficult. Betting markets provide a clean setting to estimate the price impact of digitalization since products are homogeneous and pricing behavior is easily measurable. Consistent with mixed strategy models under imperfect information, we show that between 2000 and 2018, transaction costs for gamblers have dropped by almost 70%, although transaction cost dispersion persists. Furthermore, we also find that betting markets became more informationally efficient by documenting a significant decrease in the favorite-longshot bias.

JEL Classification

D40, D83, L83, O3

Keywords

Fixed-odds sports betting, online markets, favorite-longshot bias, price dispersion, digitalization

3.1 Introduction

3.1.1 Context

In this chapter we study Premier League football[28] fixed-odds betting markets. The Premier League is the top tier of the English football league system and the biggest football league in the world in terms of revenue (Statista, 2019). In the UK, gross sports bet wins[29] amounted to three billion euro in 2018 and football is the most popular sport to bet on (Gambling Commission, 2018; H2 Gambling Capital, 2019).

The betting industry has been disrupted by the meteoric rise of online gambling. In less than two decades, the share of online gambling grew from a marginal phenomenon with only 0.1% of betting turnover in 1999 (Paton, Siegel, & Williams, 2002) to the dominant channel with over 70% market share in 2018. Between 2003 and 2018, the UK online sports betting market grew at a fabulous compound annual growth rate of 18% as shown in Figure 2 (H2 Gambling Capital, 2019). As online gambling strongly reduces search and switching costs for consumers and drastically lowers barriers to entry and marginal costs to producers, its introduction is expected to have an impact on bookmaker behavior.

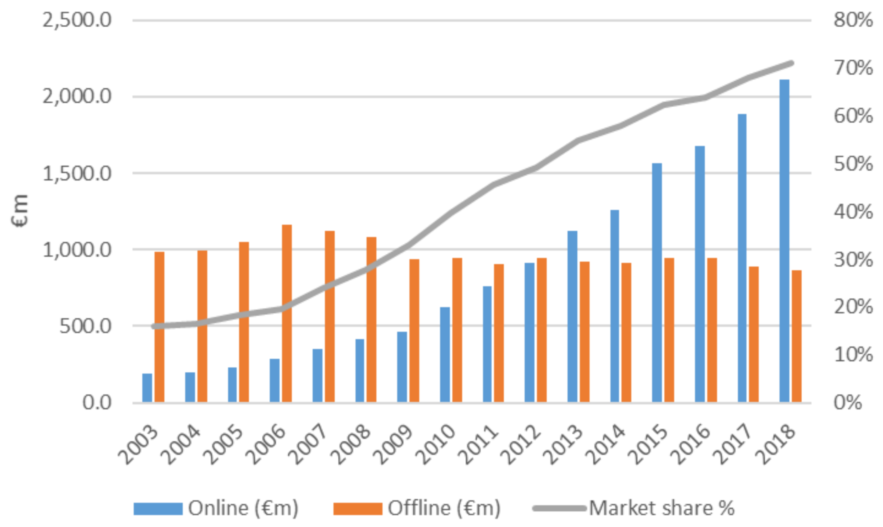

Figure 2: Online (blue bars) and offline (orange bars) sports betting gross wins in the UK in million euro (left-hand side axis). Market share is indicated by the grey line (the right-hand side axis) (H2 Gambling Capital, 2019).

3.1.2 Results

Using a dataset of odds set by 13 major bookmakers on all Premier League games between 2000 and 2018, we find that bookmakers use a dynamic pricing strategy, i.e. the transaction costs[30] they charge are a function of season and game characteristics. In particular, we show that the average transaction costs were squeezed by a factor of three, from 13.3% in the first

[28] Soccer, not American football.

[29] The total amount wagered minus all payouts.

[30] Synonyms include the margin, the juice, the cut, the overround or the vigorish.

season of our sample to 4.2% in the last season. This decrease is both statistically and economically very significant and strongly associated with the rising importance of online gambling. Within seasons, we show that bookmakers offer statistically significantly lower transaction costs on derbies and games between top teams. This is in line with Franck et al. (2013) and Grant, Oikonomidis, Bruce, and Johnson (2018) who hypothesize bookmakers offer deliberately attractive odds from time to time for marketing purposes. Factors that capture game attendance or the competitive balance do not seem to affect bookmaker pricing.

Although average transaction costs decreased dramatically over the seasons in our sample, large differences prevail between the transaction costs different bookmakers charge. This is intuitively puzzling as bets placed on the same outcome with different (online) bookmakers are virtually perfect substitutes. However, the observation is in line with the related industrial organization literature (Cavallo, 2017). The rise of the internet spawned a substantial body of theoretical and empirical work on its effects on markets. Internet evangelists anticipated “frictionless commerce” closely approximating textbook Bertrand competition (Ellison & Ellison, 2005; Litan & Rivlin, 2001). The empirical work finds that online prices are indeed lower than their offline counterparts, however, significant price dispersion is more persistent than intuitively expected “the law of one price is still no law at all” (Baye, Morgan, & Scholten, 2006, p. 46).

Theoretically, our observations that average transaction costs move down while transaction cost dispersion remains can be explained by the canonical model of Varian (1980). In this model, both decreases in marginal costs and increases in the fraction of informed gamblers lower average prices. These evolutions are both associated with increases in online gambling. However, as long as not all customers are perfectly informed, it is optimal for firms to randomize their prices which leads to an equilibrium with price dispersion, which is what we observe.

Lastly, we study the informational efficiency of the odds in our sample. The main stylized facts emerging from related empirical work show that market implied probabilities of event outcomes are generally very close to their empirical counterparts, but that consistent mispricing occurs at the extremes (Griffith, 1949; Sauer, 2005). The returns of low probability bets are systematically too low relative to the returns of high probability bets. This observation is called the “favorite-longshot bias” (FLB) and is considered “one of the most robust anomalous empirical regularities in economics” (Walls & Busche, 2003). The favorite-longshot bias puzzles economists as it implies a negative risk premium in betting markets. The standard neoclassical explanation of the favorite-longshot bias points to utility functions in which agents have a preference for variance or skewness (Golec & Tamarkin, 1998; Weitzman, 1965), but alternative behavioral theories have been developed (see Ottaviani and Sørensen (2008) for an overview). We establish a statistically and economically significantly positive link between transaction costs and the favorite-longshot bias. More specifically, we find that our most extreme favorite-longshot bias measure drops by 4.6 percentage points for every percentage point decrease in transaction costs. This means that the market became more informationally efficient over the course of our sample as the mispricing of extreme events decreased significantly. Furthermore, this interaction between transaction costs and the favorite-longshot bias is consistent with the cost-based explanation of the bias discussed by i.a. Hurley and McDonough (1995), Paton and Williams (1998) and M. A. Smith, Paton, and Williams (2006).

3.1.3 Advantages of our setting

We argue that the sports betting setting is an interesting empirical lab for a couple of reasons. First, bets on the same game are as homogeneous a product as it gets. Apart from differences in price, the products are completely identical and most betting sites are indistinguishable from each other if you would hide the logos. This removes the need to control for differences in product characteristics, which is often necessary in related studies.

Second, pricing behavior is easily measurable in a betting context. In the literature, the transaction costs bookmakers charge are measured by the sum of the inverse odds (more on this in section 2). This is convenient as we only need data on odds, which is easily accessible, to measure the transaction cost level over all outcomes jointly. Furthermore, next to the general level of the transaction costs, we can also easily get insight into how bookmakers price extreme events as we can systematically compare the returns on high odds events to those of low odds events. The observation that betting markets allow us to directly compare market implied probabilities with their empirical counterparts is also the reason why a substantial strand of the informational efficiency literature has focused on betting markets. "Economists have given great attention to stock markets in their efforts to test the concept of market efficiency, yet wagering markets are, in one key respect, better suited for testing efficiency and rationality. The advantage of wagering markets is that each asset (bet) has a well-defined termination point at which its value becomes certain" (Thaler & Ziemba, 1988, pp. 161-162).

Third, our dataset spans almost two decades, which is much longer than what is common in related work (Gorodnichenko & Talavera, 2017; Lünnemann & Wintr, 2011). As a result, we can study the long-term evolutions in transaction costs and transaction cost dispersion. Furthermore, the transaction costs in our sample continuously decrease, which is very different from for example Gorodnichenko and Talavera (2017), who find that increases and decreases are equally likely in online markets in general.

Fourth, the prices that many goods or service providers set are to some extent sticky, even online. Gorodnichenko, Sheremirov, and Talavera (2018) for example report price spells of 7 to 20 weeks for internet retailers in the US and the UK. This price rigidity is often explained by menu costs or costs associated with information processing and coordination with customers (see for example Zbaracki, Ritson, Levy, Dutta, and Bergen (2004)). The prices of bookmakers are not sticky as they are forced to price every single game, i.e. set the correct odds for the different outcomes. Next to a high variation in odds, we will show that the transaction costs bookmakers charge their customers are also highly variable between different games. As an example, bookmaker Ladbrokes charged over 560 different transaction costs in our sample. This high variability in transaction costs charged is an indication that bookmakers are very good at extracting consumer surplus[31].

Relatedly, notice that while informational efficiency is heavily researched in sports betting, relatively little work is done on transaction costs. The lack of research in this area can be attributed to the limited transaction cost variation in other much studied betting microstructures like spread betting[32] (mostly associated with American football and basketball) and pari-mutuel betting[33] (mostly associated with horse racing). In fixed-odds betting, the setting we operate in, transaction costs are implicitly embedded into the offered odds and can vary greatly between bookmakers and games as discussed above.

The rest of this chapter is structured as follows. In section 3.2 we introduce the fixed-odds betting microstructure. The theoretical model and hypotheses are examined in section 3.3. In section 3.4 the data set is introduced and descriptive statistics are examined. In section 3.5 the analyses involving the evolution of the transaction costs and transaction cost dispersion are carried out. Section 3.6 connects the transaction costs with the favorite-longshot bias and section 3.7 concludes.

3.2 Fixed-Odds Microstructure

European football betting is mainly organized via fixed-odds betting. This means bookmakers offer odds for specific game-related events on which a gambler can bet. In this chapter, we use the decimal/European odds convention. These decimal odds represent the payout on a winning unit bet. More formally, if o_{ijm} are the decimal odds for outcome i (where $i = 1$ refers to a home win, $i = 2$ refers to a draw and $i = 3$ refers to an away win) set by bookmaker j, for match m, the return on a home win bet r_{1jm} is given as follows.

$$r_{1jm} = \begin{cases} o_{1jm} - 1 & \text{if home goals} > \text{away goals ("home win")} \\ -1 & \text{if home goals} \leq \text{away goals ("draw" or "away win")} \end{cases}$$

[31] In this respect, the bookmaker industry is similar to the airline or hotel industry where dynamic pricing practices are often employed (Bilotkach, Gorodnichenko, & Talavera, 2010).

[32] In spread betting transaction costs are very rigidly set by the 11-for-10 rule (which means you earn 10 for every 11 bet if you win the bet). Levitt (2004, p. 237) remarks in a footnote that the absence of price competition is a “major puzzle in this industry”. The constancy of the transaction costs in this market microstructure is further investigated by Sandford and Shea (2013) who attribute it to the first mover disadvantage bookmakers face when setting their lines. Only recently, bookmaker competition starts to bring down transaction costs (Berkowitz et al., 2018).

[33] In pari-mutuel betting mostly associated with horse racing, commission rates tend to differ between different racetracks, but are generally equal for different races on the same track which again limits variability. A tangentially related literature in this microstructure examines the factors influencing racetrack gambling demand which includes discussions of the revenue maximizing commission rate (Gruen, 1976; W. D. Morgan & Vasche, 1979; Suits, 1979).

The returns on draw and away bets are analogous. Sports bets can thus be regarded as European binary options where the underlying is the outcome of a game. A home bet is in the money whenever the number of home goals is strictly larger than the number of away goals.

The relation between sports bets and Arrow-Debreu securities is apparent (Shin, 1992). The state space consists of the three game outcomes "home win", "draw" and "away win" ($i = 1,2,3$) and each bet on outcome i has a similar payoff vector namely $(o_{1jm}, 0, 0)$, $(0, o_{2jm}, 0)$, $(0,0, o_{3jm})$.

Without transaction costs, the odds (o_{ijm}) are the inverse state prices, i.e. the price for a security that pays off 1 when state i substantiates and zero else. As a simple example, a bet with odds of 2 provides the gambler a potential payoff of 2 for each unit bet. Equivalently, the inverse of the odds (1/2, the state price) is the stake the gambler must put up for a potential unit payoff. Under the law of one price and absence of arbitrage, the primal asset pricing equation tells us that

$$p = E(mx), \tag{1}$$

where p is the price of an asset, m is the stochastic discount factor and x is the risky payoff of the asset. In this context, we can write

$$p_i = \frac{1}{o_{ijm}} = E(mx_i) = E(m)E(x_i) + cov(m, x_i), \tag{2}$$

where p_i is the state price associated with state i and x_i the payoff in state i. We can very reasonably assume that the stochastic discount factor is not correlated with the payoffs. If we furthermore assume that the risk-free return is negligible over the maturity of the contract (which is a realistic assumption given the short-term nature of sports bets) such that $E(m) = 1$, we can write

$$p_i = \frac{1}{o_{ijm}} = E(m)E(x_i) = 1 \times E(x_i) = \pi_i\,, \tag{3}$$

where π_i is the probability associated with state i (the probability of a unit payoff for this bet). Consequently, there is a direct link between odds and probabilities:

$$o_{ijm} = \frac{1}{\pi_i}. \tag{4}$$

Note that odds are inversely related to outcome probabilities: high (low) odds reflect a low (high) probability event. In this case the expected return can be written as:

$$E(r_i) = o_{ijm}\pi_i - 1 = 1 - 1 = 0. \tag{5}$$

3.2.1 Transaction costs

The zero transaction cost assumption we introduced earlier to derive equation (5) is clearly not realistic. Bookmakers charge a fee for their services by systematically skewing the odds in their favor, i.e. setting the odds lower than they should be. As for market makers in traditional financial markets, this fee allows bookmakers to make a profit and acts as an insurance against adverse selection (i.e. customers trading on superior information) and adverse market movements (i.e. information that arrives after a bet is locked in which makes the offer more favorable for the gambler). Introducing a transaction cost distorts the link between odds and probabilities. As now

$$o_{ijm} = o^*_{ijm} - \tau_{ijm} < \frac{1}{\pi_i} \tag{6}$$

where o_{ijm} are again the odds advertised by the bookmaker, o^*_{ijm} are the "true odds", which are now different from the advertised odds because a spread of τ_{ijm} is subtracted. To derive probabilities from odds under this more realistic assumption, we have to know the pricing system employed by the bookmaker (τ_{ijm}). Equality (5) now becomes an inequality,

$$E(r_i) = (o^*_{ijm} - \tau_{ijm}) \times \pi_i - 1 < 0 \tag{7}$$

resulting in an expected negative return for the gambler. A natural metric of the transaction cost charged by the bookmaker would be τ_{ijm}, but this is unobservable. In the literature, the transaction cost is traditionally computed via the booksum, which is conveniently computable from the odds. This transaction cost metric is an aggregate measure of unfairness over all possible outcomes of an event.

The booksum Π_{jm} is defined as

$$\Pi_{jm} = \sum_{i=1}^{3} \frac{1}{o_{ijm}} \tag{8}$$

and the transaction cost as

$$\mu_{jm} = \Pi_{jm} - 1. \tag{9}$$

We have shown that the bookmaking business evolves around setting odds such that

$$\frac{1}{o_{ijm}} > \pi_i. \tag{10}$$

As

$$\sum_{i=1}^{3} \pi_i = 1, \tag{11}$$

it follows that

$$\Pi_{jm} = \sum_{i=1}^{3} \frac{1}{o_{ijm}} > 1 \tag{12}$$

and thus

$$\mu_{jm} > 0. \tag{13}$$

In any case, the smaller the transaction cost, the less the link between the odds and probabilities is distorted. The more the inverse of the advertised odds approximates the true state prices, the closer we approach an ideal efficient market.

3.2.2 The link between transaction costs and the favorite longshot bias

Notice that although the computation of the transaction cost via the booksum is common practice and convenient, it makes the implicit assumption that the proportional spreads are equal:

$$\frac{\tau_{1jm}}{o_{1jm}} = \frac{\tau_{2jm}}{o_{2jm}} = \frac{\tau_{3jm}}{o_{3jm}},$$

which is not necessarily the case. Therefore, it is interesting to look at the transaction cost and the favorite-longshot bias simultaneously. As mentioned previously, the favorite-longshot bias is the stylized fact that low odds bets have higher expected returns that high odds bets. Such differences in expected returns are induced by non-proportional spreads, i.e. bookmakers subtract a proportionally higher spread from high odds compared to low odds.

We measure the favorite-longshot bias by the difference in expected returns between odds deciles. The favorite-longshot bias metric for bookmaker j in season s is given by:

$$FLB_{js}^{X\%} = E\left(r_{js}^{X\%}\right) - E\left(r_{js}^{100-X\%}\right),$$

where $E\left(r_{js}^{X\%}\right)$ and $E\left(r_{js}^{100-X\%}\right)$ indicate the expected returns of all bets with the $X\%$ lowest odds (favorites) and the $X\%$ highest odds (longshots) respectively. By looking at transaction costs and the favorite-longshot bias jointly, we can both say something about the aggregate unfairness of bets and the transaction cost differential between high odds and low odds bets.

3.3 Theoretical model and hypotheses

The rise of the internet rejuvenated interest in theoretical models that can explain the observed levels of price dispersion even though search costs, spatial differentiation and information asymmetries have been reduced drastically (see Baye et al. (2006) for an overview). Of particular interest for us are the models with an information clearinghouse, i.e. a third party that provides an overview of all the prices in the market. Bettors have access to many real-time price comparison websites[34] such that they can be easily informed if they want to and can pick the bookmaker with the most attractive offering. However, these tools are not as established as economists would predict such that a considerable fraction of customers is not perfectly informed. A study by the European Commission (2014) finds that 48% of consumers know that price comparison websites exist, but they are not really familiar with them. Furthermore, 47% of consumers who do not use them indicate they only buy from websites they already know.

[34] Examples include oddsportal.com, betmonitor.com and oddspedia.com.

To stich the different elements of this chapter (transaction costs, transaction cost dispersion and the favorite-longshot bias) together, we use a two-stage model where gamblers first choose a bookmaker and then a team to bet on. We model the bookmaker choice via Varian (1980), which can be interpreted as a situation where customers have access to an "information clearinghouse" where all supplier prices are visible. In a second stage, the gamblers choose which outcome they want to bet on. Here we will draw on the cost-based favorite-longshot bias of which versions are discussed in Hurley and McDonough (1995), Paton and Williams (1998) and M. A. Smith et al. (2006). In these models, the choice crucially depends on the level of transaction and information costs as we will discuss below.

3.3.1 Bookmaker choice and pricing

In the model, n price-setting, identical bookmakers can take bets on games at a constant marginal cost c and charge a transaction cost of μ. Aside from differences in transaction costs charged, bets with different bookmakers are completely homogeneous. The gamblers have unit demand and have a reservation transaction cost of v such that they will not make a bet if $\mu > v$. There are two types of gamblers. Informed gamblers I first visit an odds comparison website where all bookmakers are listed and make a bet with the bookmaker with the most favorable offer (given that this offer is not above their reservation transaction cost). Uninformed gamblers M pick a bookmaker at random and make a bet whenever the offered transaction cost is less than their maximum wiliness to pay. These uninformed clients for example do not know that odds comparison sites exist, they can be influenced by bookmaker marketing such that they do not only look at the price or choose the same bookmaker as their friends. As there are n bookmakers, each bookmaker has $U = \frac{M}{n}$ uninformed clients.

For every game, the bookmakers choose a transaction cost μ from a density function $f(\mu)$. The bookmaker with the lowest transaction cost attracts $I + U$ clients, i.e. all informed gamblers and their share of uninformed gamblers. The other bookmakers still do business with their share of uninformed clients U. The bookmakers face a tradeoff between charging bargain and ripoff prices in the Salop and Stiglitz (1977) terminology, i.e. setting a competitive price increases the probability to attract the price sensitive gamblers but decreases the markup they charge on their loyal clients.

Varian (1980) shows that there is an equilibrium in mixed strategies, i.e. firms randomize their prices. If firms would predicably charge a certain transaction cost instead of randomizing their prices, a bookmaker could easily undercut the competition and attract all informed gamblers. The continuous cumulative distribution function of transaction costs can be found via the situation where a bookmaker is indifferent between simply charging the reservation price and earning $(v - c)$ from its fraction of uninformed customers $\frac{U}{M+I}$, and charging a price μ between $[\mu_0, v]$, where it captures the informed gamblers $\frac{I}{I+M}$ with a probability[35] of $[1 - F(\mu)]^{n-1}$ while also charging this price to its fraction of uninformed gamblers $\frac{U}{M+I}$. More formally,

$$(\mu - c)\left(\frac{I}{I+M}[1 - F(\mu)]^{n-1} + \frac{U}{M+I}\right) = (v - c)\frac{U}{M+I},$$

from which we can isolate the equilibrium cumulative distribution function:

$$F(\mu) = 1 - \left[\frac{(v - \mu)U}{(\mu - c)I}\right]^{\frac{1}{n-1}} \; on \; [\mu_0, v].$$

The lower bound μ_0 can be found by solving for $F(\mu_0) = 0$:

$$\mu_0 = c + (v - c)\frac{U}{U+I}.$$

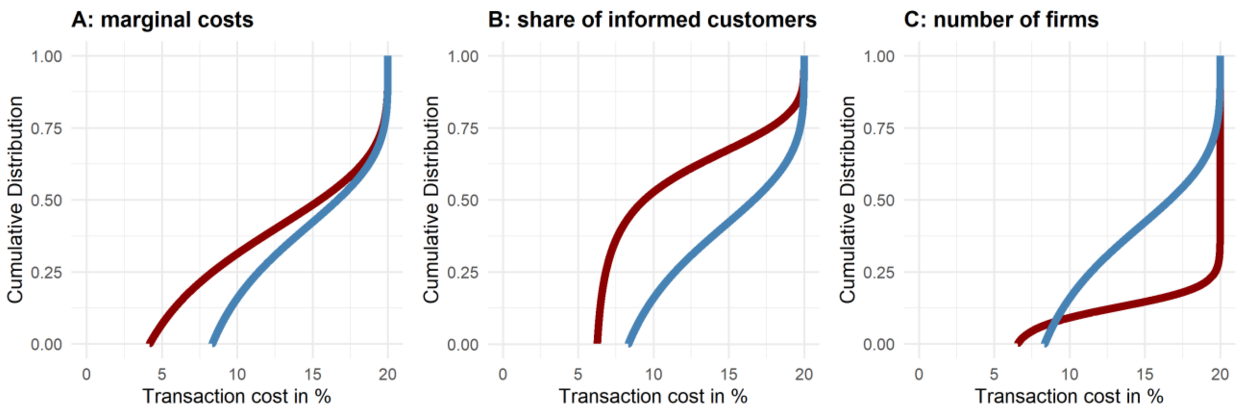

Figure 3: Comparative statics of the Varian (1980) model. Panel A shows the CDF for low (red line) versus high (blue line) marginal costs. Panel B shows the CDF for a large share of informed customers (red line) versus a low share of informed customers (blue line). Panel C shows the CDF for a large number of firms (red line) versus a low number of firms (blue line)[36].

Visual comparative statics of the model are shown in Figure 3 (for a more formal discussion of the effects of changes in the parameters, see for example Varian (1980) or (Pennerstorfer, Schmidt-Dengler, Schutz, Weiss, & Yontcheva, 2020)). The model is especially useful for our purposes for a couple of reasons.

[35] To capture the informed gamblers, all other bookmakers have to charge a higher price, which happens with a probability of $[1 - F(\mu)]^{n-1}$.

[36] For replicability purposes, in the graphs $v = 20$, $M = 5$, $c = 1$ & 6, $I = 5$ & 50, $n = 5$ & 25.

First, decreases in marginal costs lead to lower average transaction costs as shown in panel A of Figure 3. As the rise of online betting drastically lowered marginal costs, we would expect average transaction costs to decrease over time.

Second, the model also predicts that increases in the fraction of informed gamblers lead to lower average transaction costs[37] as shown in panel B of Figure 3. As online gambling drastically reduces search costs (Paton et al., 2002) and information asymmetries, we would expect the proportion of informed gamblers to rise. Furthermore, switching costs largely disappear as a gambler can browse from one bookmaker to another by the click of a mouse. Decreasing switching costs are another incentive for customers to become informed, which can further fuel an increase in informed gamblers.

Third, only when every customer is perfectly informed, we have a degenerate distribution function where the transaction cost collapses to the marginal cost. This means that we will have transaction cost dispersion in equilibrium as long as we have some uninformed customers, "price dispersion is […] the measure of ignorance in the market" (Stigler, 1961, p. 214). This prediction is in line with the empirical work that finds price dispersion is more persistent than intuitively expected, even in online markets (Baye et al., 2006; Goldfarb & Tucker, 2019; Gorodnichenko et al., 2018).

One reason why a significant part of customers stays uninformed is that non-trivial switching costs could remain, even online (Franck et al., 2013). Another possibility is that firms engage in obfuscation or bait-and-switch strategies, i.e. practices that frustrate consumer search and learning. Examples include the sales of add-ons at high unadvertised prices, offering some products at low prices to attract customers while charging high prices for others (loss-leader strategy) or complicated product descriptions (Ellison & Ellison, 2009). In our context, there are many examples of such activities including bookmakers offering "odds boosts", i.e. extremely favorable odds on some events (while charging much higher transaction costs on other events) or bookmakers multiplying the initial balance gamblers post in their betting account. However, unbeknown to some customers, the profits made via such promotions can often only be withdrawn from their accounts under very specific and stringent circumstances. Furthermore, research shows that advertised odds during live games are focused on complex products (Newall, Thobhani, Walasek, & Meyer, 2019).

Lastly, the rise of online gambling increases the number of bookmakers. From a producer perspective, there is no more need for bookmakers to invest in brick-and-mortar gambling facilities, which significantly reduces the barrier to entry. While only a handful of bookmakers offered odds online in the late nineties, more than 1200 bookmakers are active on the

[37] See also Pennerstorfer et al. (2020) who empirically confirm this observation in the retail gasoline market.

internet today[38]. From a consumer perspective, bettors can now access a wide range of gambling services online, from the comfort of their homes (Vlastakis, Dotsis, & Markellos, 2009). This globalization of gambling markets effectively increases the number of bookmakers as well, as it allows gamblers to place bets with bookmakers that where inhibitingly expensive to visit physically. As noted by J. Morgan, Orzen, and Sefton (2006), the effect of an increase in the number of firms in the Varian (1980) model is more subtle. An increase in the number of firms reduces the average transaction costs of the informed gamblers but increases the average transaction costs of the uninformed. With more firms, there is a smaller probability that a bookmaker will be able to attract the informed, reducing the incentive to charge a competitive transaction cost, but at the same time, the informed customers will have more firms to choose from.

3.3.2 Bet choice and pricing

In the second stage of the model, gamblers choose which bet they want to make. For this stage, we rely on the cost-based FLB model, which establishes a positive relationship between the level of information and transaction costs and the size of the favorite-longshot bias. Versions of this model are discussed in Hurley and McDonough (1995), Paton and Williams (1998) and M. A. Smith et al. (2006). In its most condensed form, the model stipulates that without information or transaction costs, risk-neutral agents know the true probabilities of both outcomes and bet such that the market implied probabilities are equal to the empirical probabilities. When transaction and information costs increase, gamblers are less capable of distinguishing between favorites and longshots. These gamblers will bet more randomly such that the longshot is overbet and the favorite is underbet, inducing the favorite-longshot bias. As a result, there exist a positive relationship between the size of the transaction and information costs and the magnitude of the favorite-longshot bias. This explanation echoes Grossman and Stiglitz (1980) who argue informational efficiency is impossible when information is costly.

As argued earlier, we expect the rise in online gambling to be associated with lower average transaction costs, which would decrease the FLB as well. Furthermore, also note that the internet makes vast oceans of data on games, teams and players easily available to both consumers and producers. This strongly decreases information asymmetries and costs. This allows bookmakers to improve their predictions of game outcomes while bettors have more tools at their disposal to identify bookmaker mispricing and enforce market discipline.

---

[38] This is the number of online sportsbooks and racebooks listed at online.casinocity.com. "Online.casinocity.com is probably the world's most comprehensive and widely used online gambling portal." (R. J. Williams & Wood, 2007, p. 8)

3.4. Data & Descriptive Statistics

The data set comprises 18 Premier League seasons from 2000-2001 to 2017-2018 totaling 6840 games[39]. Odds are provided by the following bookmakers: bet365, Blue Square, Bet&Win (bwin), Gamebookers, Interwetten, Ladbrokes, Pinnacle, Sportingbet, Stan James, Sporting Odds, Stanleybet, VC Bet (BetVictor) and William Hill. The bookmakers in the sample are a good representation of the football gambling market as the most important firms are included. GVC Holdings, which as of 2018 owns Ladbrokes, bwin, Sportingbet and Gamebookers is the largest online sports betting operator in the world with an annual revenue of over £3.5 billion in 2018 (GVC Holdings, 2019). bet365 had an operating revenue of over £2.8 billion in 2018 and is the biggest bookmaker worldwide in terms of monthly online traffic (azBookmakers, 2019; bet365 Group Limited, 2018). William Hill had an operating revenue of over £1.6 billion in 2018 (William Hill PLC, 2019). Also note that both incumbent bookmakers like Ladbrokes, which dates back to 1886 (azBookmakers, 2019), William Hill, founded in 1934 (William Hill PLC, 2019) or BetVictor, founded in 1946 (Harris, 2016) and online disrupters like Interwetten, founded in 1990 (Interwetten, n.d.), Pinnacle, founded in 1998 (Cronin, 2018) and bet365, founded in 2000 (bet365, n.d.) are present in the sample. Our sample only starts in 2000 which is a few years after the birth of online betting. However, in the early years online betting was a marginal phenomenon and responsible for only 0.1% of betting turnover in 1999 (Paton et al., 2002). For every game, the bookmakers provide three odds: home win odds, draw odds, and away win odds. For weekend fixtures, the odds are recorded on Friday afternoon, for midweek fixtures, the odds are collected on Tuesday morning.

Early literature usually assumed that bookmakers are merely balancing their books, such that they lock in a profit independent of the outcome of the game (see for example Woodland and Woodland (1991)). If the books are perfectly balanced, the winners can be paid by the losers while the bookmaker pockets a small transaction cost. In more recent work, Grant et al. (2018) classify bookmakers with this business model as "book-balancing bookmakers" (BBBs). In seminal work however, Levitt (2004) argued that bookmakers can increase their profits if they exploit the preferences of their clients, i.e. making sure there are more losers than winners. Grant et al. (2018) classify bookmakers with this business model as "position-taking bookmakers" (PTBs). In contrast with book-balancing bookmakers who live off volume, position-taking bookmakers can achieve higher margins if they actively manage their client portfolio. As they take positions against their customers, they try to reduce their exposure to sophisticated gamblers by monitoring the behavior of their clients and limiting their action radius if they win too much (Franck et al., 2013; Grant et al., 2018). The majority of bookmakers are position-taking bookmakers, which is also reflected in our sample (Pinnacle is the only BBB).

[39] The data are collected by Joseph Buchdahl and can be accessed via his website: http://www.football-data.co.uk/data.php.

Table 28: Descriptive statistics of the data set.

	N games	N odds observations	Home win %	Draw %	Away win %	Mean goals per match
2000-2001	380	5,535	48%	27%	25%	2.61
2001-2002	380	6,363	43%	27%	30%	2.63
2002-2003	380	7,941	49%	24%	27%	2.63
2003-2004	380	7,977	44%	28%	28%	2.66
2004-2005	380	7,974	46%	29%	26%	2.57
2005-2006	380	10,260	51%	20%	29%	2.48
2006-2007	380	10,257	48%	26%	26%	2.45
2007-2008	380	11,079	46%	26%	27%	2.64
2008-2009	380	11,397	46%	26%	29%	2.48
2009-2010	380	11,400	51%	25%	24%	2.77
2010-2011	380	11,400	47%	29%	24%	2.80
2011-2012	380	11,400	45%	25%	31%	2.81
2012-2013	380	11,394	44%	28%	28%	2.80
2013-2014	380	9,120	47%	21%	32%	2.77
2014-2015	380	8,097	45%	25%	30%	2.57
2015-2016	380	7,977	41%	28%	31%	2.70
2016-2017	380	7,980	49%	22%	29%	2.80
2017-2018	380	7,980	46%	26%	28%	2.68
Sum	6,840	165,531				
Average			46%	26%	28%	2.66

Descriptive statistics are reported in Table 28. Each Premier League season consists of 380 games. The number of odds observations is much larger than the number of games as we observe home win, draw and away win odds of many bookmakers for each game. The number of odds observations varies because of attrition of bookmakers from the data set (mainly because of mergers or acquisitions) and new bookmakers entering the sample. The last four columns of Table 28 show game related statistics: the home win percentage, the draw percentage, the away win percentage and the number of goals scored on average. We notice the well-known home field advantage: in 46% of the games the home team wins, in 28% of the games the away team wins and 26% of the games end in a draw. There is some idiosyncratic variation in these game outcome statistics between seasons, but no clear trends. Simple t-tests that compare the first nine seasons with the last nine seasons teach us that the differences in these game outcome statistics are not statistically significant (results not reported). On average, 2.66 goals are scored each game. The average number of goals is statistically significantly higher in the second half of the data set (result not reported), but the difference is economically meaningless. We conclude that the nature of the game did not change drastically and thus we can rule out major game related changes impacting our results.

3.5 On transaction costs and transaction cost dispersion

Table 29 shows some descriptive statistics of the bookmaker transaction costs. As our unit of observation is now the transaction cost as opposed to individual odds, the number of observations drops by a factor three as we need odds on the full outcome space to compute the transaction costs. As noted before, the number of observations per bookmaker varies because not all bookmakers were active or included in the sample during the entire sample period. We have some missing values because for some games we do not have all three odds so we cannot compute the transaction cost. The statistics in Table 29 point to a large heterogeneity in pricing between bookmakers and seasons. The average transaction cost over all games and bookmakers is 7.9%, but note that the bookmakers that were active in the early seasons have much higher average transaction costs than the bookmakers who entered the sample at a later point.

Table 29: Descriptive statistics of the transaction costs of the bookmakers in the sample.

	N	Coverage	Average	Median	Std dev	Min	Max
Stanleybet	373	2001-2002	11.62%	11.63%	0.27%	10.85%	12.98%
Sporting Odds	760	2002-2004	10.70%	11.54%	1.56%	7.62%	12.27%
Interwetten	6,822	2000-2018	10.57%	10.40%	3.14%	4.70%	19.21%
Sportingbet	4,504	2000-2012	10.02%	10.15%	1.35%	4.34%	15.87%
Ladbrokes	6,778	2000-2018	8.97%	7.26%	3.36%	2.49%	14.42%
William Hill	6,726	2000-2018	8.95%	6.86%	3.30%	2.30%	19.30%
Gamebookers	4,856	2000-2013	8.49%	7.95%	1.81%	0.00%	18.36%
bwin	5,319	2004-2018	7.73%	7.44%	2.12%	2.35%	12.23%
Blue Square	2,280	2007-2013	7.34%	7.03%	1.49%	3.48%	11.14%
Stan James	3,459	2005-2015	7.18%	6.31%	2.17%	2.05%	11.86%
Bet365	6,080	2002-2018	5.62%	5.52%	2.67%	1.69%	16.70%
BetVictor	4,940	2005-2018	4.69%	3.44%	2.57%	0.21%	14.29%
Pinnacle	2,280	2012-2018	2.03%	2.03%	0.16%	1.25%	4.54%
Sum	55,177						

Figure 4 plots the distribution of some summary statistics of the transaction costs per match and per season. The main takeaway is the dramatic reduction in bookmaker transaction costs over the last two decades as shown in panel A of Figure 4. The average transaction cost was 13.3% in the season 2000-2001 and dropped almost monotonically to 4.2% in the season 2017-2018, i.e. a factor three squeeze in just 18 seasons[40]. The median, minimum and maximum transaction costs shown in panels B, E and F of Figure 4 dropped substantially as well.

[40] It is worth mentioning that multiple legislative changes took place during the sample period: the Gambling Act 2005 (Light, 2007) and the 2014 Gambling (Licensing and Advertising) Act. However, attempts to attribute transaction cost changes to these Acts specifically were fruitless.

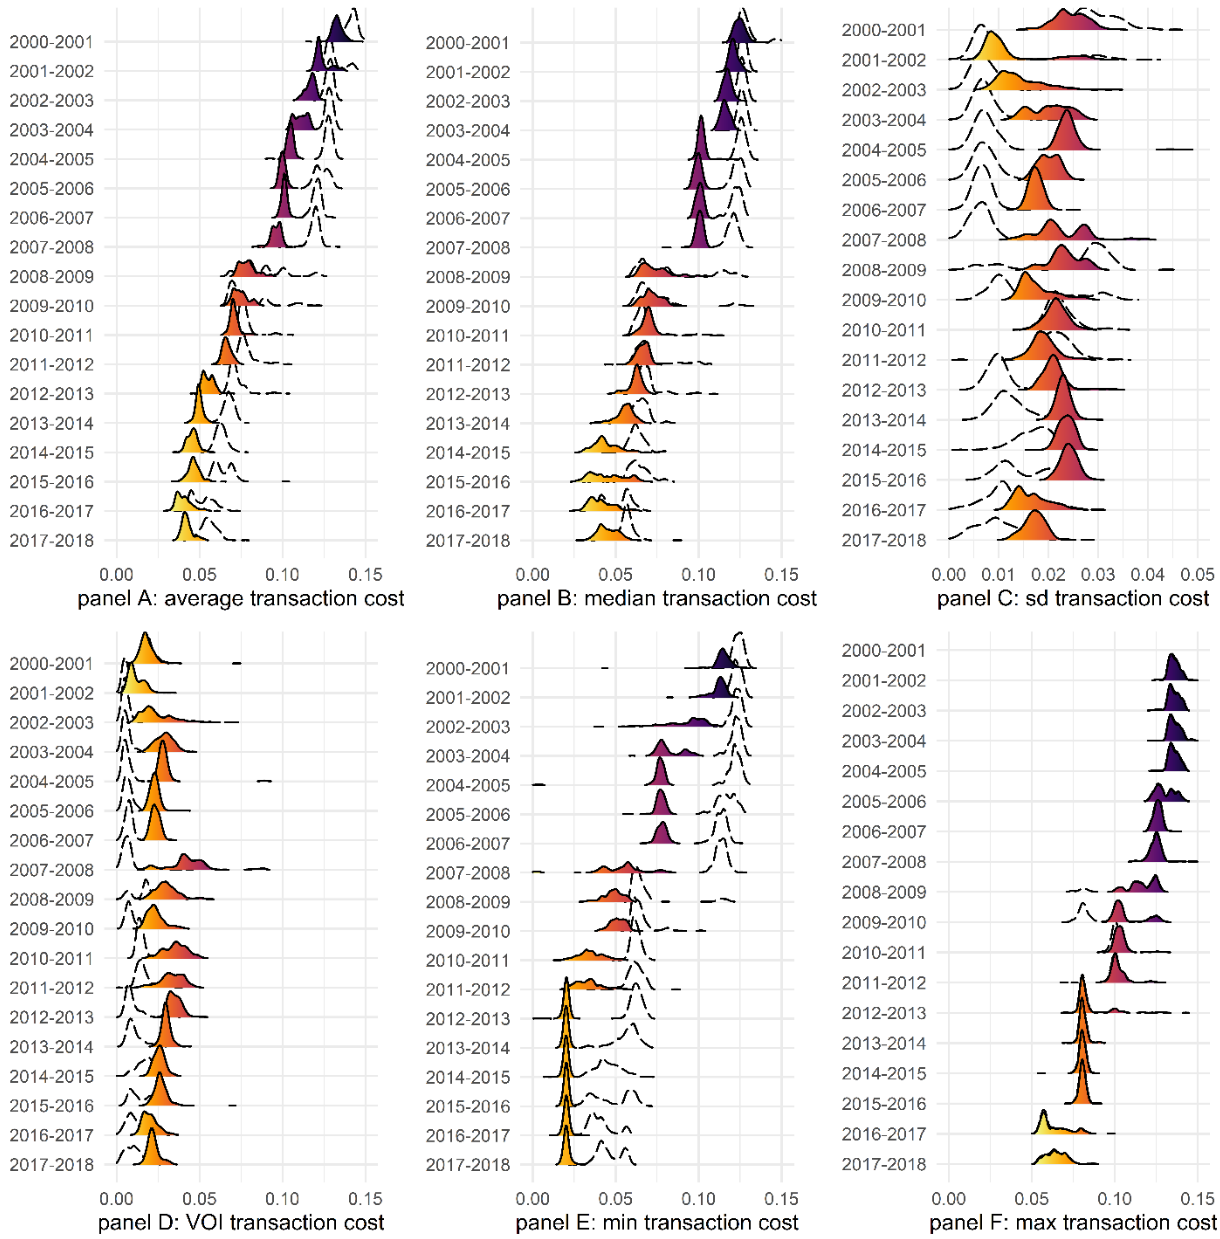

Figure 4: Panel A shows the average transaction cost per match for seasons 2000-2001 to 2017-2018 for the Premier League. For every match, the transaction costs of all bookmakers were computed and averaged. The distribution of these averages for all matches in a season is plotted in panel A. Analogously, panel B shows the median, panel C the standard deviation, panel D the value of information, panel E the minimum and panel F the maximum. The colored distributions are computed by using all available transaction costs. For the dashed line distributions, only the transaction costs of the three bookmakers that are in the sample over the entire period (2000-2018) are taken into account.

To measure the transaction cost dispersion, we use the standard deviation and the value of information, two popular dispersion metrics. The value of information is the transaction cost cut a customer can achieve by systematically betting with the cheapest bookmaker instead of picking a random bookmaker, i.e. $VOI_m = E(\mu_{mj}) - \min(\mu_{mj})$. The advantage of this metric is that it can be easily interpreted as the benefit of becoming fully informed. The transaction cost dispersion metrics are plotted in panels C and D of Figure 4 and no clear trends substantiate. This in itself is surprising. The transaction costs have come down drastically, but apparently, the dispersion between transaction costs of different bookmakers has not substantially decreased[41]. The VOI metrics teach us that a gambler could have achieved a transaction cost decrease of around 2.5 percentage points in any season by choosing the bookmaker with the lowest transaction cost.

As discussed earlier, some of the bookmakers disappeared from the sample while others entered the sample at a later point. To rule out that the transaction cost drops are attributable merely to a changing sample constituency, the above analysis is repeated using data only from the three bookmakers that were active during the entire sample period: Interwetten, Ladbrokes and William Hill. The statistics for these bookmakers are plotted via the dashed line densities shown in Figure 4. The results are identical to the full sample analysis: a very large drop in transaction costs but no clear trends with respect to the dispersion.

An important question to answer next is whether the observed price dispersion arises because of permanent differences in pricing between bookmakers or because bookmakers employ a mixed strategy as predicted by the Varian (1980) model. To shed light on this issue, we follow Chandra and Tappata (2011) and compute the temporal price dispersion via rank reversals, i.e.

$$r_{kl} = \frac{1}{T_{kl}} \sum_{m=1}^{T_{kl}} \boldsymbol{I}_{\{\mu_{km} > \mu_{lm}\}},$$

where we compute the fraction of games where bookmaker k has a higher transaction cost than bookmaker l where T_{kl} stands for the total number of games where both bookmakers are active in. In a mixed strategy, we would observe a significant proportion of reversals as firms randomize their prices to a large extent. Overall, if we express all rank reversals as a proportion between 0% and 50% (such that bookmaker k is on average cheaper than bookmaker l), we find an average of 18% across all bookmakers, which can be interpreted as evidence for temporal price dispersion. To put this number in perspective, Chandra and Tappata (2011) and Pennerstorfer et al. (2020) find a rank reversal proportion in gasoline markets of 15% and 10.5% respectively. The average however does not reveal the full picture as not all firms appear to engage in mixed strategies.

Table 30 shows the proportion of games where the bookmaker in the columns charges a higher transaction cost than the bookmaker in the rows. In general, the evidence presented again points in the direction of mixed strategies. For 15% of the bookmaker pairs however,

[41] As a robustness exercise, we computed the standard deviation of the residuals of the regression model introduced below. These residuals can be interpreted as the price of a homogenous goods as we control for other factors that influence pricing. The analysis is virtually analogous to the one presented in figure 3; although there is some idiosyncratic variation in the standard deviation of the residuals between seasons, there are no clear trends.

there is not a single reversal, i.e. the bookmaker that is generally cheaper never has a higher transaction cost than the bookmaker that is generally more expensive. For the large majority of cases without any rank reversals, bookmaker Pinnacle is involved. Pinnacle advertises itself with the slogan "always best odds with the lowest margin" so its ads seem to align with reality. It appears that bookmakers are indeed deploying mixed strategies except for Pinnacle which consistently sets lower prices. Furthermore, Pinnacle also explicitly welcomes sophisticated customers such as arbitrageurs whereas other bookmakers ban them from their platforms. Grant et al. (2018) classify Pinnacle as a "book-balancing bookmaker", i.e. a bookmaker that behaves as an infinitely risk averse market maker with profits that are independent of the outcome of the game. As they merely match buyers with sellers, they have no adverse selection issues which can allow Pinnacle to set lower transaction costs. All other bookmakers in our sample are "PTBs", "position-taking bookmakers" in the Grant et al. (2018) taxonomy.

Table 30: The rank reversals between all bookmakers in the sample over the entire sample period. Rank reversals measure the proportion of observations where the transaction cost of the bookmaker in the columns is larger than that of the bookmaker in the rows.

	IW	**LB**	**SB**	**WH**	**SY**	**B365**	**SO**	**BW**	**SJ**	**VC**	**BS**	**PS**
GB	0.98	0.68	0.74	0.64	0.56	0.22	0.80	0.75	0.41	0.37	0.38	0.00
IW		0.24	0.16	0.18	0.00	0.01	0.00	0.08	0.01	0.03	0.02	0.00
LB			0.28	0.52	0.04	0.02	0.02	0.43	0.14	0.06	0.43	0.00
SB				0.67	0.67	0.10	0.37	0.43	0.24	0.13	0.12	/
WH					0.01	0.02	0.00	0.51	0.13	0.08	0.53	0.00
SY						/	/	/	/	/	/	/
B365							0.55	0.94	0.83	0.49	0.93	0.04
SO								/	/	/	/	/
BW									0.23	0.02	0.30	0.00
SJ										0.13	0.72	0.00
VC											0.83	0.06
BS												0.00

In Figure 5, we show the fraction of games where a bookmaker holds a certain position in the transaction cost ranking for every fourth season of our sample (the other seasons are shown in appendix and the conclusions are similar). Again, we see that most bookmakers engage in mixed strategies. Particularly striking examples are VC bet (VC) and William Hill (WH). In season 2009-2010, they are represented in every position of the transaction cost ranking. This means that, depending on the game, they can charge the highest or lowest transaction cost of all bookmakers in the sample. Furthermore, also notice that the relative position of bookmakers can change quite dramatically between seasons. Ladbrokes (LB) is a mid-priced bookmaker in all seasons except for seasons 2006-2007 & 2007-2008 (shown in appendix) where it is consistently expensive. Another example is Sportingbet (SB), which starts as the cheapest bookmaker in the sample in 2000-2001, but consistently becomes more expensive until it leaves the sample. These results confirm the validity of the taxonomy of Grant et al. (2018) in the sense that PTBs and BBBs coexist but have very different business models. PTBs have a much more active pricing strategy, i.e. a wider transaction cost variance both within and between seasons, while BBBs simply try to maximize their volume by charging predictably low fees.

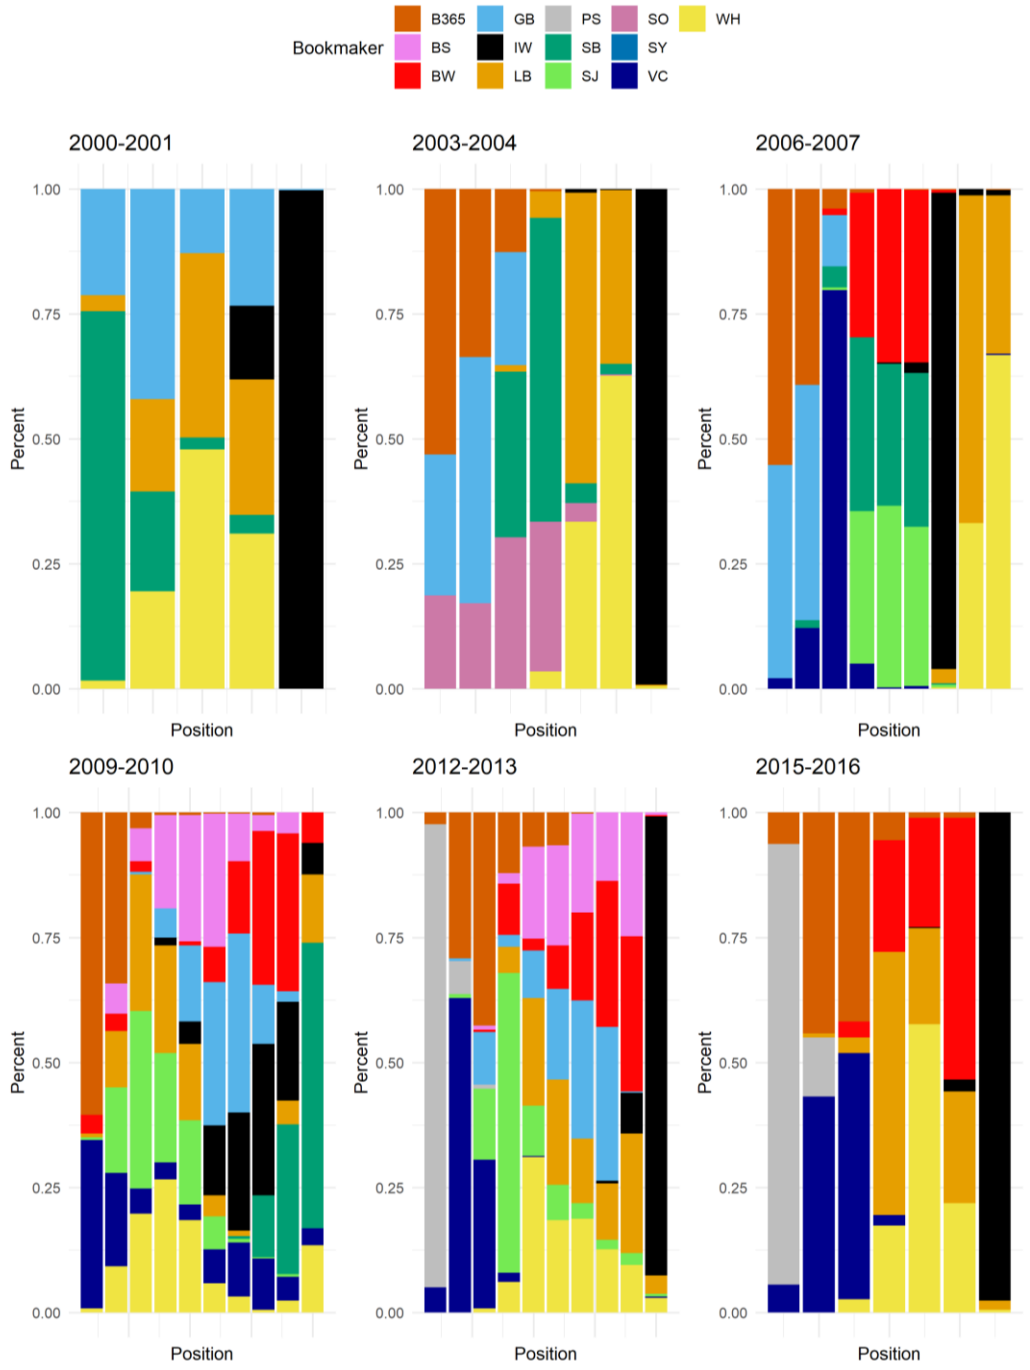

Figure 5: The horizontal axes of the stacked bar charts show the transaction cost positions (lowest transaction cost on the left and highest transaction cost on the right). The vertical axes indicate the percentage of games where a bookmaker occupies a certain position.

To further strengthen our understanding of the pricing strategies of individual bookmakers, we visualize the transaction costs of five bookmakers in Figure 6 (graphs of the other bookmakers are similar but are not shown). Two observations stand out. First, bookmakers appear to use a dynamic pricing model within seasons, leading to different transaction costs for different games in the same season. Furthermore, these pricing models appear to be quite stable over time as the transaction cost distributions are very similar over different seasons, take for example William Hill in the first 8 seasons. We will zoom in on factors that drive transaction cost variation between games in the section below. From time to time, however, bookmakers review the overall level of their transaction costs and to cut them drastically, leading to large discrete movements of their transaction cost distributions. Take William Hill again as an example, the bookmaker cut its transaction cost from around 12.5% to around 6% over a single season. Also note that the bookmakers do not all change their transaction cost level at the same time, the adjustments are staggered over different seasons. A possible explanation for this behavior is that bookmakers avoid within season menu costs via the use of an automated, dynamic pricing system. However, between different seasons, bookmakers face considerable menu costs with respect to the overall transaction cost level they will charge. These menu costs are not related to the act of changing the price level itself, but rather to managerial menu costs which include information gathering and decision-making costs which can be substantial (Stamatopoulos, Bassamboo, & Moreno, 2021). Furthermore, Zbaracki et al. (2004) argue that firms might be reluctant to cut prices to avoid retaliation of competitors, which can incite a negative price spiral.

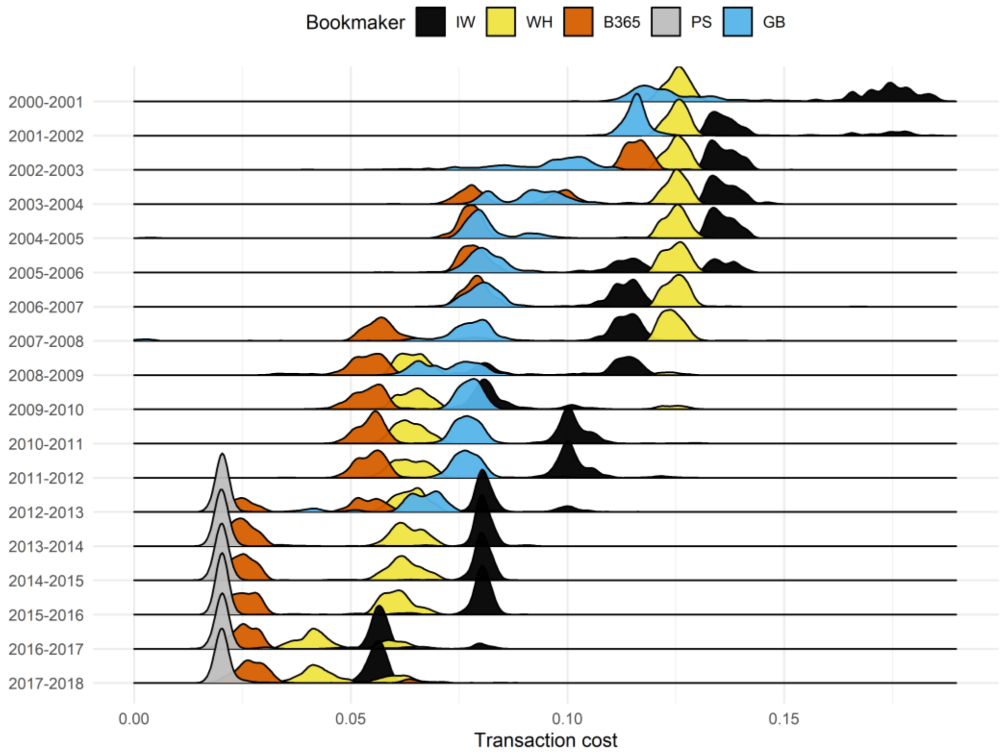

Figure 6: Distribution of transaction costs per game per season for five selected bookmakers: Interwetten (IW), William Hill (WH), b365, Pinnacle (PS) and Gamebookers (GB).

3.5.1 Regression models

In this section, we introduce regression models that exploit the cross-sectional variation in the transaction costs per game between different bookmakers to statistically underpin our visual analyses and to further explore which variables explain the transaction costs. We estimate the following model via pooled OLS:

$$\begin{aligned}\mu_{jm} = \alpha + \beta_1 online\_market\_share_m + \beta_2 seasons2010\_2018_m \\ + \beta_3 weekend\_dummy_m + \beta_4 attendance\_home\_and\_away_m \\ + \beta_5 team\_quality\_difference_m + \beta_6 derby\_dummy_m \\ + \beta_7 cracker\_dummy_m + \beta_k month_m + \beta_l bookmaker_j + \varepsilon_{jm}\end{aligned}$$

Our dependent variable is μ_{jm}, the transaction cost charged by bookmaker j in match m in percentage points. Our independent variables include industry, time, team and bookmaker related variables that could be expected to drive the transaction costs. The industry variable is $online\_market\_share_m$, the market share of online sports betting in the UK[42]. As discussed previously, we expect a larger online gambling market share to be correlated with lower bookmaker transaction costs. To capture other potential changes in the bookmaking industry, we include $seasons2010\_2018_m$, a dummy variable equaling zero for the first half of the data set, and one for the second half of the data set[43].

Next, we include a number of team-related variables. In seminal work, Levitt (2004) shows that bookmakers in the NFL take advantage of bettor preferences to increase their returns. We could expect bookmakers to increase transaction costs on games where gamblers have a strong preference to bet on and/or that attract relatively large numbers of uninformed gamblers. On the other hand, increased bookmaker competition for such bets might reduce transaction costs and bookmakers could set deliberately low prices for marketing purposes (Franck et al., 2013). The variable $attendance\_home\_and\_away_m$ proxies the attention the game received. It is the sum of the average home attendance of the home team and the average home attendance of the away team in the respective season[44]. Forrest and Simmons (2008) show that heavily supported teams receive better odds. Furthermore Franck, Verbeek, and Nüesch (2011) find that this effect is amplified in weekend games, so we also include a weekend dummy. In a similar vein, we include $derby\_dummy_m$, a dummy variable equaling one if the game is considered a derby. Eleven different matchups were considered derbies, a full list is provided in appendix. $cracker\_dummy_m$ is a dummy equaling one if the game is considered a “cracker”: a game between top teams. For each season, the crackers are defined as the three matchups where the sum of total points obtained in the last two seasons of the

---

[42] Data were obtained from H2 Gambling Capital (2019) (and visualized in Figure 2).

[43] This measure is crude by design. More detailed measures or season fixed effects run into identification issues with the market share of online gambling.

[44] This data was collected from https://www.worldfootball.net/attendance/eng-premier-league-2000-2001/1/.

home and away team is largest. As Deutscher, Ötting, Schneemann, and Scholten (2019) show that betting volumes are larger for games where the outcomes are more certain, we include the variable $team\_quality\_difference_m$, which proxies the competitive balance of the game. It is the difference between total points obtained by the teams in the two previous seasons[45].

Furthermore, we include month fixed effects as games towards the end of the season are generally more important than games in the beginning of the season (teams competing for the championship, for Champions League tickets or fighting relegation). For NCAA basketball, Humphreys et al. (2013) show that betting action increases as the season progresses. Furthermore, game outcome uncertainty could be higher in the months following a transfer window. Lastly, $bookmaker_j$ are dummy variables indicating the bookmaker that set transaction cost μ_{jm}. ε_{jm} is an error term. We estimate standard errors clustered at both the bookmaker and season level as we expect the regression model errors could be correlated within these clusters[46].

In a similar vein, we regress our measure of transaction cost dispersion, the standard deviation of the transaction costs per match, sd_margin_m on the abovementioned variables. This results in the following model. Parameters are again estimated via pooled OLS.

$$\begin{aligned} sd\_margin_m = {} & \alpha + \beta_1 online\_market\_share_m + \beta_2 seasons2010\_2018_m \\ & + \beta_3 weekend\_dummy_m + \beta_4 attendance\_home\_and\_away_m \\ & + \beta_5 team\_quality\_difference_m + \beta_6 derby\_dummy_m \\ & + \beta_7 cracker\_dummy_m + \beta_8 number\_of\_bookies_m + \beta_k month_m + \varepsilon_{jm} \end{aligned}$$

Note that we cannot include bookmaker fixed effects as our dependent variable is already an aggregation of the transaction costs of multiple bookmakers. We do include the $number\_of\_bookies_m$ variable that measures the number of bookmakers for which data was available for game m. We estimate standard errors clustered at the season level[47].

The model estimates are shown in Table 31. In columns 1 and 2, the dependent variable is the transaction cost per game per bookmaker. The only difference between these two models is that the market share of UK online sports betting is not included in the first model. In model (1) the seasons dummy is highly significant, both statistically and economically. In the last half of the sample, the transaction cost is 3.7 percentage points lower than in the first half. This is a lot given the order of magnitude of the transaction costs (the average transaction cost over

[45] If a team did not compete in the Premier League in at least one of the last two previous seasons, it receives no points for these seasons.
[46] The results are virtually similar under different clustering choices, these results are available upon request.
[47] Again, the results are virtually similar under different clustering choices, these results are available upon request.

all seasons is 7.9%). This result supplements the conclusion of the previous graphical analyses: the transaction costs in Premier League betting dropped greatly over the last decades. Interestingly, our result is in contrast with Angelini and De Angelis (2019) who find transaction costs are time-invariant via a different methodology and a smaller sample.

In model (2) the online market share variable (expressed in percentage points) is introduced. First, note that the coefficient is negative and statistically significant. A larger online gambling market share is associated with lower transaction costs. Second, the size of the coefficient is economically meaningful. A 10 percentage points increase in the online sports betting market share is associated with a 1.3 percentage point drop in transaction costs. This is a lot given that the online sports betting market share in the UK increased by more than 50 percentage points over the sample period, which would induce a 6.5 percentage point drop[48]. Third, note that the size of the coefficient of the seasons dummy is much smaller in model (2) and the statistical significance vanishes. The market share variable seems to capture the transaction cost decrease quite well. This link between online gambling and falling transaction costs was already hinted at by Elaad et al. (2020).

Of the team related features, $attendance\_home\_and\_away$ and $team\_quality\_difference$ are not significant. These characteristics apparently do not influence bookmaker pricing. However, the $derby\_dummy$ and $cracker\_dummy$ variables are significant at the 1% significance level in the full model. The transaction costs in derby games and top games are on average respectively 12 and 19 basis points lower. Increased competition between bookmakers for these important games is a plausible explanation of the decreased transaction costs. This is consistent with Franck et al. (2013) who argue bookmakers offer deliberately attractive odds from time to time for marketing purposes. Lastly, model (2) appears to capture the variation in transaction costs rather well, reflected by the high adjusted R^2. The variables explain more than 80% of the transaction cost variance.

[48] Another evolution that could impact transaction costs in the same direction is the introduction of the betting exchanges in 2000. As opposed to the traditional quote driven betting markets where bookmakers unilaterally set odds, betting exchanges are order driven markets where the exchanges merely match buyers with sellers in a continuous double auction (Flepp, Nüesch, & Franck, 2017). However, between 2014 and 2019, the market share of betting exchanges hovered between just 10.3% and 7.6% (Gambling Commission, 2019) so it would be a stretch to these exchanges were responsible for the large transaction cost changes we document in this chapter.

Table 31: Output of regression models. In model (1) and (2), the dependent variable is the transaction cost in percentage points charged by bookmaker j in match m. Online_market_share is the market share of online sports betting in the UK expressed in percentage points. seasons2010_2018 is a dummy variable capturing the second half of the data set. weekend_dummy is a dummy equaling one if the game was played on Friday, Saturday or Sunday. attendance_home_and_away proxies the attention the game received. It is the sum of the average home attendance of the home team and the average home attendance of the away team in the respective season. team_quality_difference proxies the quality difference between the teams. It is the difference between total points obtained by both teams in the two previous seasons. derby_dummy is a dummy equaling one if the game is a derby. Cracker_dummy is a dummy equaling one if the game is considered a "cracker". For each season, the crackers are defined as the three matchups where the sum of total points obtained in the last two seasons of the home and away team is largest. The model includes bookmaker and month fixed effects and standard errors are clustered at the bookmaker and season level. In model (3) the dependent variable is the standard deviation of the transaction costs in game m. number_of_bookies measures the number of bookmakers for which data was available on a given match. Model (3) includes month fixed effects and the standard errors are clustered at the season level.

	Dependent variable:		
	transaction_cost		sd_transaction_costs
	(1)	(2)	(3)
Online_market_share		-13.091	-0.193
		t = -8.997***	t = -0.213
seasons2010_2018	-3.717	0.327	0.158
	t = -5.070***	t = 0.663	t = 0.564
weekend_dummy	-0.080	-0.056	0.031
	t = -1.348	t = -1.301	t = 1.187
attendance_home_and_away	-0.002	0.0003	0.0001
	t = -1.941*	t = 0.455	t = 0.261
team_quality_difference	-0.0002	-0.0001	0.00002
	t = -0.331	t = -0.250	t = 0.154
derby_dummy	0.004	-0.120	0.069
	t = 0.060	t = -3.334***	t = 2.484**
cracker_dummy	-0.090	-0.189	0.057
	t = -1.034	t = -4.930***	t = 0.810
number_of_bookies			0.027
			t = 0.451
Constant	6.026	9.981	1.756
	t = 8.346***	t = 17.443***	t = 2.459**
Bookmaker fixed effects	Yes	Yes	No
Month fixed effects	Yes	Yes	Yes
SE method	CL B+S	CL B+S	CL S
Observations	51,211	51,211	6,080
Adjusted R^2	0.714	0.822	0.025
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 32: Bookmaker fixed effects. All variables of model (2) shown in table 3 were included in the model that was used to estimate these parameters. The transaction cost of bookmaker Pinnacle is the base rate. Bookmakers are coded as follows: bet365 (B365), Blue Square (BS), Bet&Win (BW), Gamebookers (GB), Interwetten (IW), Ladbrokes (LB), Sportingbet (SB), Stan James (SJ), Sporting Odds (SO), Stanleybet (SY), VC Bet (VC) and William Hill (WH).

	Dependent variable:
	transaction_cost
bookmakerB365_transaction_cost	0.963
	t = 5.198***
bookmakerBS_transaction_cost	2.687
	t = 14.577***
bookmakerBW_transaction_cost	3.533
	t = 22.909***
bookmakerGB_transaction_cost	1.972
	t = 7.167***
bookmakerIW_transaction_cost	5.261
	t = 28.477***
bookmakerLB_transaction_cost	3.916
	t = 21.149***
bookmakerSB_transaction_cost	3.533
	t = 11.924***
bookmakerSJ_transaction_cost	2.234
	t = 10.888***
bookmakerSO_transaction_cost	2.846
	t = 6.997***
bookmakerVC_transaction_cost	0.749
	t = 5.493***
bookmakerWH_transaction_cost	3.875
	t = 21.170***
Constant	9.981
	t = 17.141***
Includes all model variables	Yes
SE method	CL B
Observations	51,211
Adjusted R^2	0.822
Note:	*p<0.1; **p<0.05; ***p<0.01

With respect to the transaction cost dispersion, the results are very different. Only the derby dummy is statistically significant at the 5% confidence level as shown in column 3 of Table 31. Furthermore, the adjusted R^2 is very low. The persistence of the price dispersion is remarkable, but in line with the literature on the effect of e-commerce on prices. Brynjolfsson and Smith (2000) for example find that online book and CD prices are 9-16% lower than their offline equivalents but also find relatively high price dispersion. In the life insurance industry, J. R. Brown and Goolsbee (2002) show that increases in internet usage decrease the price of term life insurance by 8-15%. Furthermore, the authors find that price dispersion initially increased, but decreased as internet usage became more widespread. For the airline industry, Orlov (2011) finds that a higher internet adaptation rate is associated with lower average ticket prices. Internet access does not seem to affect interfirm price dispersion. However, intrafirm price dispersion increases, suggesting that the internet allows firms to better price discriminate between their own customers.

To further analyze the transaction cost dispersion, it is interesting to zoom in on the bookmaker fixed effects estimates which are shown in Table 32. While only the bookmaker effects are shown, all variables discussed earlier are included in the model that was used to estimate these coefficients. The transaction cost of Pinnacle is the base rate. Apparently, all bookmakers have both economically and statistically significantly higher transaction costs than Pinnacle. The differences in bookmaker pricing are striking. The transaction costs of Interwetten (IW) for example, are 5.3 percentage points higher than those of Pinnacle. This is large if you consider that the average transaction cost over all seasons and bookmakers is 7.9%. This pricing power of bookmakers is remarkable given that bets placed on the same outcome with different bookmakers are virtually perfect substitutes. These results are in line with earlier research that shows gamblers can reduce costs considerably by comparing the offerings of different bookmakers (Angelini & De Angelis, 2019; Forrest, Goddard, & Simmons, 2005). Brand loyalty induced by intense marketing could be one of the reasons explaining this relative price inelasticity. In 2017, 50% of the Premier League teams were sponsored by gambling operators (Deutscher et al., 2019). Alternatively, gamblers could be held back by the (minor) inconvenience of setting up an account with another bookmaker. Furthermore, as the transaction cost are not directly observable, some gamblers might just not be aware of the large differences in pricing.

3.6 On the link between transaction costs and the favorite-longshot bias

Economists have long been interested in sports betting markets for testing market efficiency as they can be regarded as “simple financial markets” (Sauer, 1998, p. 2021). Fundamentally, participants in both financial and betting markets require an appropriate return for the risk they take. Both types of markets bring together large pools of investors with heterogeneous beliefs and information and like derivatives trading and active asset management, sports betting is zero sum in nature (before commissions) (Levitt, 2004). However, sports bets are uniquely qualified for tests of informational efficiency as their true values are quickly and exogenously revealed, which circumvents the joint-hypothesis problem (Fama, 1991; Thaler & Ziemba, 1988). As a result, a sizable part of the betting literature examines the efficient market hypothesis in i.a. American football (Gray & Gray, 1997), baseball (Woodland & Woodland, 1994), basketball (Berkowitz et al., 2015), horse racing (Snowberg & Wolfers, 2010; M. Sung & Johnson, 2010) and football (Croxson & Reade, 2014; Vlastakis et al., 2009).

In this section, we look at the link between transaction costs and the favorite-longshot bias (FLB), i.e. the stylized fact that bets on favorites outperform bets on longshots. As discussed before, there are theoretical arguments that would predict the existence of a positive relationship between the level of the transaction costs and the FLB.

We link our analyses on transaction costs and the FLB via the cost-based FLB model, which predicts a positive relationship between the size of the FLB and the size of the transaction and information costs (Hurley & McDonough, 1995). Interestingly, Hurley and McDonough (1995) cannot confirm the model in an experimental setting, but later work appears to find an empirical link between the size of transaction costs and the FLB. Paton and Williams (1998) compare fixed odds betting with spread betting for Premier league football in the 1996-1997 season and find that the FLB is higher in fixed odds betting, where transaction costs are higher. Similarly, M. A. Smith et al. (2006) and Bruce, Johnson, Peirson, and Yu (2009) find a larger favorite-longshot bias in betting contexts where transaction costs are higher for horse racing in 2002 and 1996 respectively. Sobel and Raines (2003) show that the favorite-longshot bias is more pronounced when the complexity of a bet, and thus the information cost, is higher. Our setup allows us to build upon this literature by investigating whether cross-sectional variation in transaction costs between bookmakers in fixed-odds betting predicts the size of the FLB. Again, the size of our sample is an advantage as earlier work usually relied on a single season to test their hypotheses.

First, we test for a FLB over the entire sample by computing the average return for odds deciles, which are shown in Table 33. The average return on bets in the first decile (smallest odds, the favorites) is -1.40% while the average return on bets in the tenth decile (the larges odds, the longshots) is -21.49%. Similarly, bets in the second decile outperform bets in the ninth decile by more than 13 percentage points. Although we find evidence of a favorite-longshot bias, the relationship is between odds and the expected return is noisy, which is consistent with earlier literature (Direr, 2011).

Table 33: Descriptive statistics of the odds deciles over the entire sample.

Decile	N	Min odds	Max odds	Mean odds	Standard deviation odds	Empirical Event Probability	Average Return
1	16554	1.05	1.62	1.39	0.14	71.57%	-1.40%
2	16553	1.62	2.00	1.81	0.11	53.45%	-3.80%
3	16553	2.00	2.40	2.22	0.12	41.94%	-7.24%
4	16553	2.40	3.00	2.69	0.16	32.41%	-13.09%
5	16553	3.00	3.20	3.14	0.07	29.33%	-8.05%
6	16553	3.20	3.40	3.27	0.04	28.96%	-5.26%
7	16553	3.40	3.60	3.48	0.08	26.36%	-8.29%
8	16553	3.60	4.33	3.93	0.20	22.88%	-10.28%
9	16553	4.33	6.00	4.97	0.46	16.74%	-17.25%
10	16553	6.00	34.00	9.12	3.38	9.70%	-21.49%

Next, we compute different FLB metrics per bookmaker and season. We measure the FLB by the difference in expected returns between odds deciles:

$$FLB_{js}^{X\%} = E\left(r_{js}^{X\%}\right) - E\left(r_{js}^{100-X\%}\right),$$

where $E\left(r_{js}^{X\%}\right)$ and $E\left(r_{js}^{100-X\%}\right)$ indicate the mean returns of all bets with the $X\%$ lowest odds (favorites) and the $X\%$ highest odds (longshots) respectively for bookmaker j in season s. These metrics are visualized in Figure 7. Although there is significant variation in the metrics between seasons, what stands out is that they all decline significantly over the sample period. Furthermore, the decline is largest for the most extreme FLB metrics, which intuitively makes sense as there is a more outspoken FLB for these metrics.

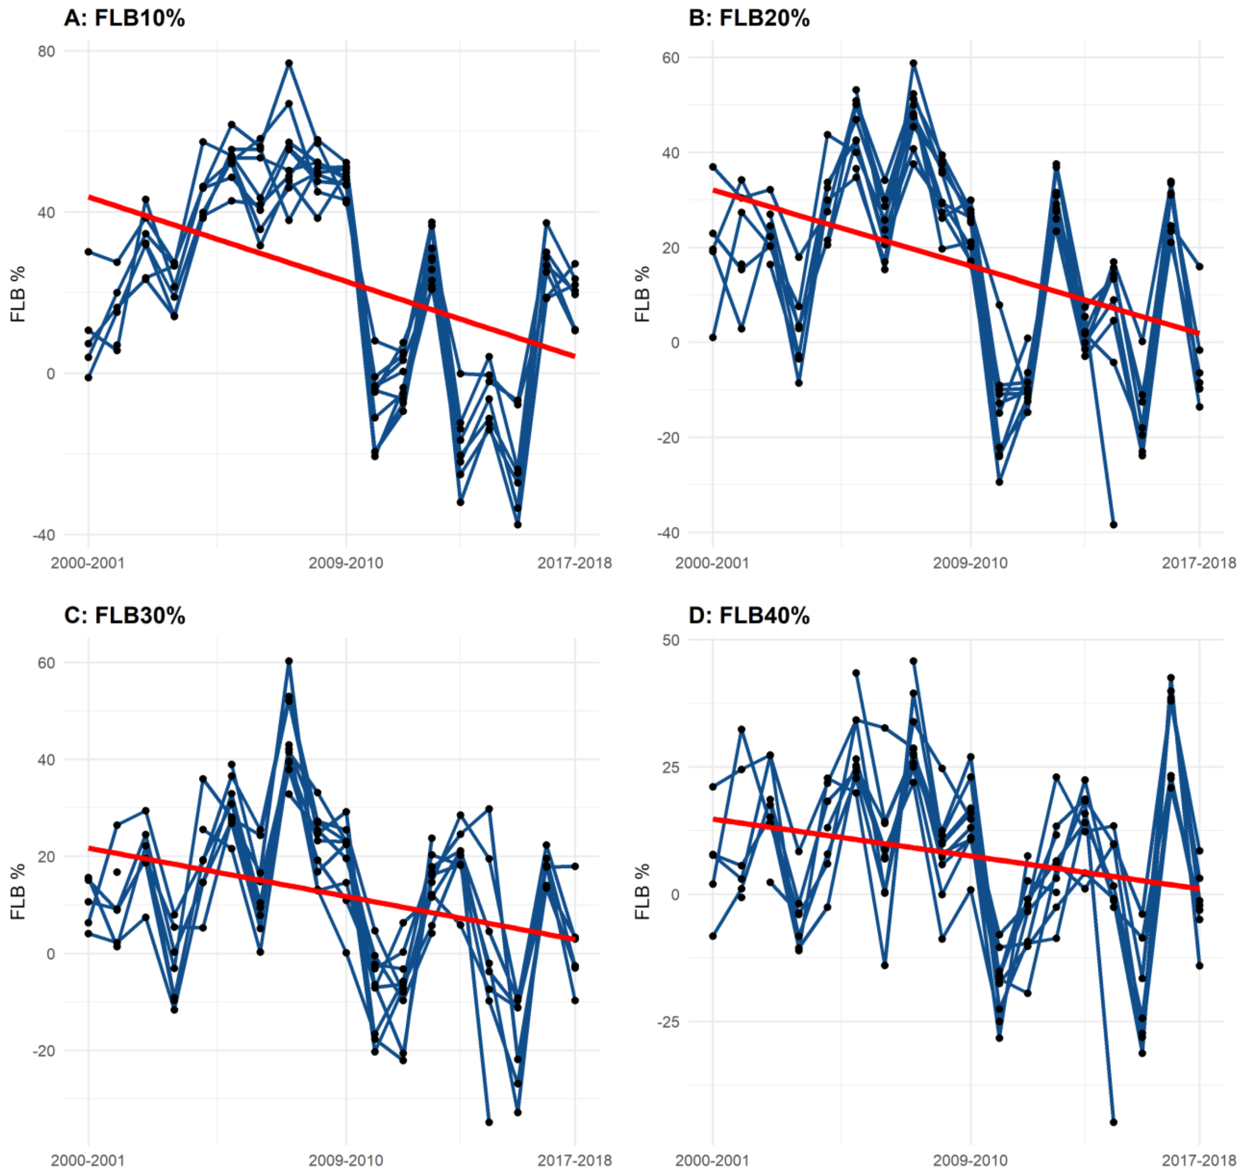

Figure 7: Different favorite-longshot bias metrics per season and bookmaker. FLB10% measures the difference in mean returns between the 10% lowest odds and the 10% highest odds. The other FLB metrics are analogous.

To investigate whether the FLB is smaller when transaction costs are lower as predicted by the cost-based explanation of the FLB, we run the following regression where we exploit the cross-sectional variation in bookmaker pricing:

$$FLB_{js}^{X\%} = \alpha + \beta_1 \mu_{js} + \beta_l bookmaker_j + \varepsilon_{js}$$

where $FLB_{js}^{X\%}$ is the expected return difference between bets with the lowest and highest $X\%$ of odds with bookmaker j in season s, μ_{js} is the average transaction cost, $bookmaker_j$ are bookmaker dummies and ε_{js} is an error term.

The regression outputs are shown in Table 34. The columns indicate different FLB metrics in decreasing level of extremeness (FLB10% is the most extreme metric as it only takes the bottom and top 10% of the sample into account while FLB50% measures the difference in return between the bottom and top half of the data set). As shown in Table 34, we find a positive and statistically significant relationship between the transaction costs charged and the FLB, which is consistent with the cost-based explanation of the FLB. For every percentage point decrease in transaction costs, the FLB drops by 4.6 percentage points for the most extreme FLB metric. This is economically meaningful if we consider the average transaction costs dropped from 13.3% to 4.2% over the sample period. Note that the size of the coefficients drops monotonically when we take less extreme FLB metrics into account, which is again what we would expect ex-ante.

The coefficients of the bookmaker dummies give insight into the pricing strategies of the respective bookmakers, i.e. whether bookmakers charge a relatively higher transaction cost for longshots compared to favorites. The base rate is again bookmaker Pinnacle and as shown in Table 34, 4 bookmakers have a significantly lower FLB on at least two metrics.

Table 34: Output of the FLB regression models. The dependent variable measures the FLB by the expected return differential between odds deciles. (1) measures the expected return difference between the bottom 10% odds (favorites) and the top 10% (longshots), (2) measures the expected return difference between the bottom 20% odds (favorites) and the top 20% (longshots) and so on. mean_transaction_cost measures the mean transaction cost charged by the respective bookmaker in the respective season. For the bookmaker dummies, the FLB of Pinnacle is the base rate. Bookmakers are coded as follows: bet365 (B365), Blue Square (BS), Bet&Win (BW), Gamebookers (GB), Interwetten (IW), Ladbrokes (LB), Sportingbet (SB), Stan James (SJ), Sporting Odds (SO), Stanleybet (SY), VC Bet (VC) and William Hill (WH).

	Dependent variable:				
	FLB				
	FLB10%	FLB20%	FLB30%	FLB40%	FLB50%
mean_transaction_cost	4.611	3.290	1.843	1.420	1.270
	t = 4.738***	t = 6.748***	t = 3.604***	t = 2.444**	t = 2.067**
bookmakerB365	1.894	-1.005	1.181	-6.739	-1.075
	t = 0.543	t = -0.576	t = 0.645	t = -3.236***	t = -0.488
bookmakerBS	2.725	-0.509	-0.563	-2.227	-9.822
	t = 0.528	t = -0.197	t = -0.208	t = -0.722	t = -3.012***
bookmakerBW	-4.036	-4.534	-1.405	-7.988	-3.130
	t = -0.728	t = -1.631	t = -0.482	t = -2.411**	t = -0.893
bookmakerGB	1.574	-1.532	0.328	-2.922	-4.278
	t = 0.249	t = -0.483	t = 0.099	t = -0.773	t = -1.069
bookmakerIW	-9.689	-5.907	0.539	-3.198	-7.015
	t = -1.166	t = -1.418	t = 0.123	t = -0.644	t = -1.336
bookmakerLB	-7.063	-9.819	-4.462	-8.527	-5.695
	t = -1.042	t = -2.891***	t = -1.252	t = -2.106**	t = -1.330
bookmakerSB	-9.185	-6.246	-1.746	-3.936	3.083
	t = -1.180	t = -1.601	t = -0.427	t = -0.846	t = 0.627
bookmakerSJ	-11.230	-3.359	-1.000	-5.130	-14.304
	t = -2.293**	t = -1.369	t = -0.389	t = -1.754*	t = -4.624***
bookmakerSO	-15.304	-23.758	-14.232	-17.642	-22.384
	t = -1.821*	t = -5.642***	t = -3.223***	t = -3.514***	t = -4.215***
bookmakerSY	-37.838	-16.892	-3.677	-18.925	-19.941
	t = -4.056***	t = -3.614***	t = -0.750	t = -3.396***	t = -3.384***
bookmakerVC	1.558	1.013	-1.129	-4.433	-0.224
	t = 0.603	t = 0.783	t = -0.832	t = -2.872***	t = -0.137
bookmakerWH	-12.945	-8.388	-5.201	-5.542	-7.855
	t = -1.908*	t = -2.468**	t = -1.459	t = -1.368	t = -1.833*
Constant	-8.732	-4.879	-0.979	1.832	1.900
	t = -4.412***	t = -4.921***	t = -0.941	t = 1.550	t = 1.519
SE method	CL B	CL B	CL B	CL B	CL B
Observations	147	147	147	147	147
Adjusted R^2	0.145	0.124	0.034	-0.0002	0.017
Note:			*p<0.1; **p<0.05; ***p<0.01		

3.7 Conclusion

The rise of online gambling has disrupted UK fixed-odds betting markets. The theoretical model of Varian (1980) predicts that decreases in marginal costs and increases in the fraction of informed gamblers, which are both associated with the increasing importance of online gambling, lowers average prices but does not eliminate price dispersion. Our observations are consistent with these predictions.

We find that bookmakers engage in mixed strategies. Transaction costs vary greatly between different seasons, bookmakers and games. We find that average transaction costs decreased from 13.3% in the first season of the data set (2000-2001) to 4.2% in the last season of the data set (2017-2018). This decrease is both statistically and economically exceptionally significant and driven by the rise of online gambling. Our model furthermore teaches us that game attendance and the quality difference between teams are not systematically priced in the transaction costs. Derby games or games between top teams however enjoy statistically significant lower transaction costs. Although average transaction costs have come down, the transaction cost dispersion is very persistent over time. This is remarkable in a context where customers can compare bookmaker prices with a few mouse clicks, but it is in line with the empirical work on online price dispersion.

Furthermore, we study informational efficiency by looking at the interaction between the favorite-longshot bias (FLB) and transaction costs. We find a strong and statistically significant relationship between the FLB and transaction costs, which is consistent with the cost-based explanation of the FLB. It appears that lower transaction costs, induced by the rise of online gambling, have also driven down the main empirical anomaly in betting markets and thus made odds more informationally efficient.

3.8 Appendix to chapter 3

Derby list

Table 35: List of all Derbies

Name	Teams Involved
North West Derby	Manchester United - Liverpool FC
North London Derby	Arsenal FC - Tottenham Hotspur
Merseyside Derby	Liverpool FC - Everton FC
Manchester Derby	Manchester United - Manchester City
Tyne-Wear Derby	Sunderland AFC - Newcastle United
Second City Derby	Birmingham City - Aston Villa
Black Country Derby	West Bromwich Albion - Wolverhampton Wanderers
South Coast Derby	Southampton FC - Portsmouth FC
A23 Derby	Crystal Palace - Brighton & Hove Albion
East Lancashire Derby	Blackburn Rovers - Burnley FC
South Wales Derby	Swansea City - Cardiff City

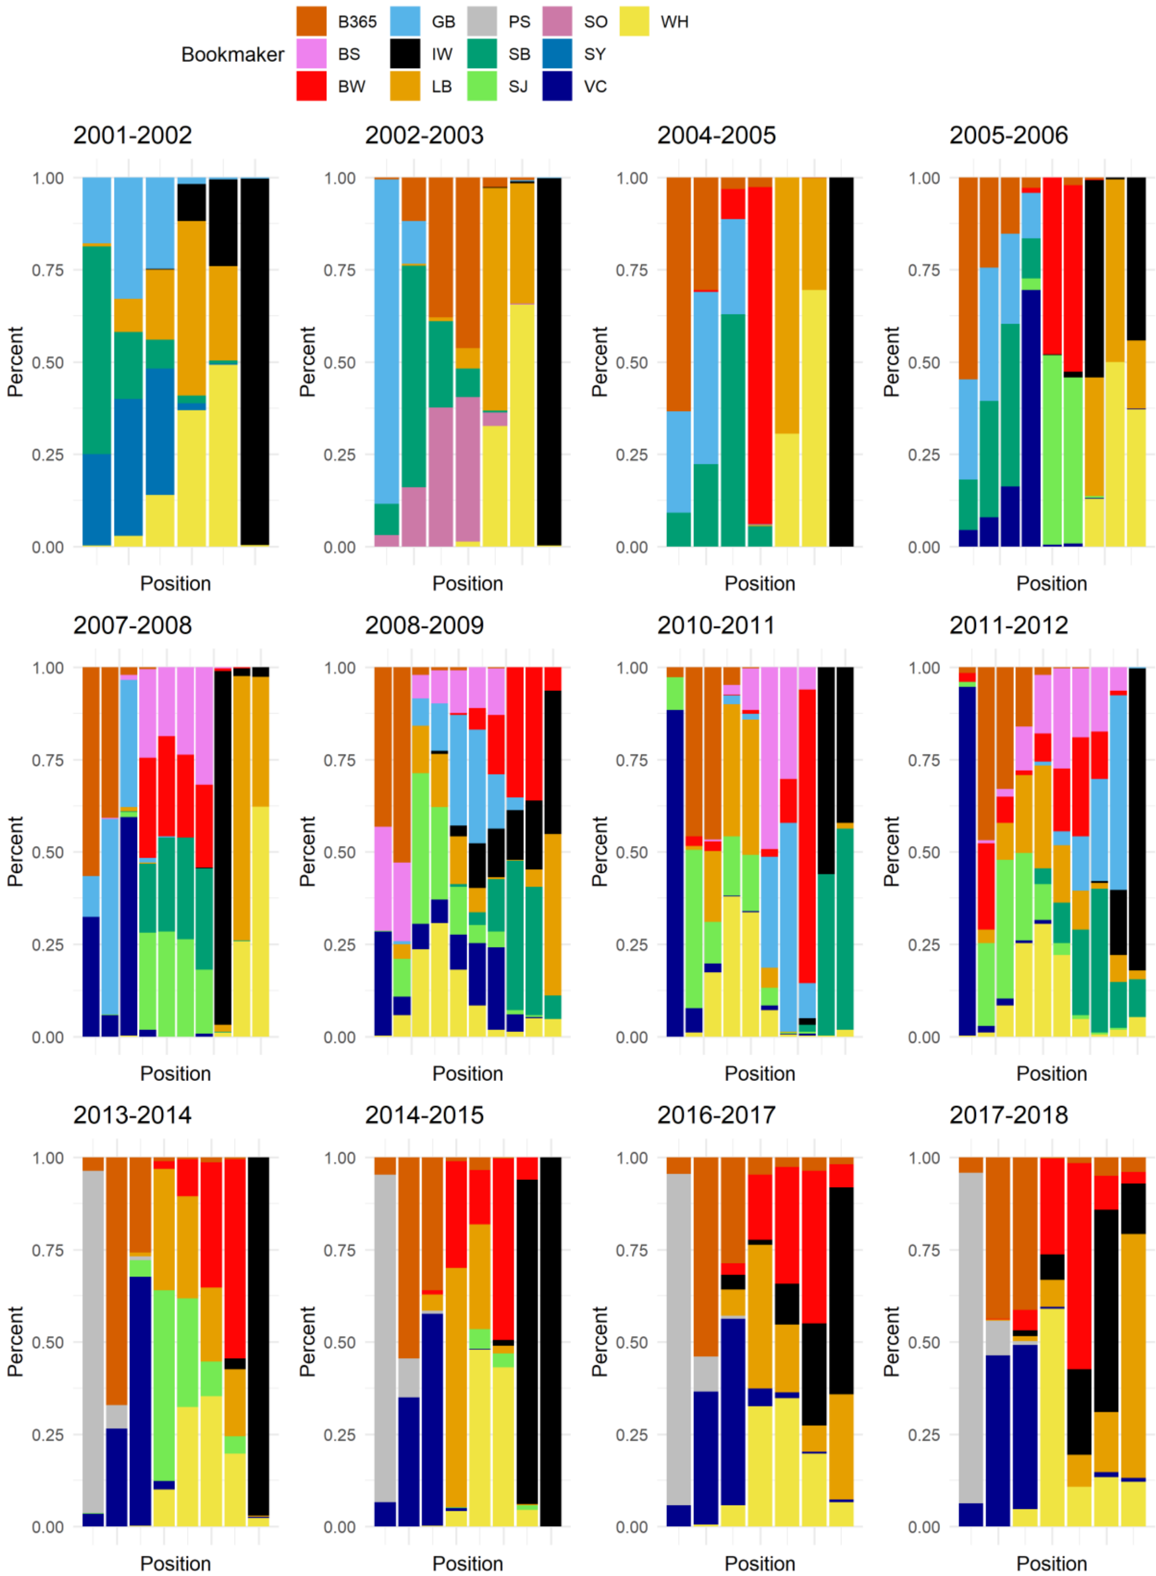

Figure 8: The x-axes of the stacked bar charts show the transaction cost positions (lowest transaction cost on the left and highest transaction cost on the right). The y-axes indicate the percentage of games where a bookmaker occupies a certain position.

Chapter 4: Does time series momentum also exist outside traditional financial markets? Near-laboratory evidence from sports betting.

Abstract

The presence of time series momentum in the returns of financial assets puzzles economists. We show that this anomaly is also present in sports betting, a seemingly unrelated market and a near-laboratory setting. We find both a statistically significant and economically meaningful difference between the returns of bets on recent winners compared to recent losers. These differences are not due to rational compensations for variance or skewness, but are consistent with underreaction. The bookmakers, the market makers in this context, do not appear to react efficiently to new information.

JEL Classification

G14, G40, Z2

Keywords

Time series momentum, sports betting, underreaction, asset pricing, behavioral finance

In this chapter we document the existence of time series momentum in sports betting. We use a data set of hourly pre-game odds of 32 bookmakers on the outcomes of 17380 football (soccer) games played in 50 major football leagues between 2015 and 2016. First, we find evidence of return predictability based on past odds movements via a regression framework. More specifically, we find that outcomes of which the odds have decreased (and thus the probability of the underlying event has increased) generally have higher expected returns compared to outcomes of which the odds have increased (and thus the probability of the underlying event has decreased). Second, we deploy a time series momentum trading strategy which confirms the results from the regression analysis. The average returns of betting on odds which have been decreasing are statistically significantly higher compared to returns on bets with increasing odds. Furthermore, the differences are economically meaningful. For example, betting on home outcomes of which the odds have decreased by at least 10% in the runup to the game outperforms bets of which the odds have increased by at least 10% by 13.83 percentage points (t-stat of 3.57). These effects are robust for a battery of different design choices. It appears that bookmakers, i.e. the market makers, do not incorporate new information efficiently into their odds.

We argue that these effects cannot be explained by rational risk premia as assets that have become more risky (higher odds, lower probability of the underlying event happening) have a lower return, while assets that have become less risky (lower odds, higher probability of the underlying event happening) enjoy a higher return. Furthermore, the results do not seem driven by a rational skewness compensation either. The explanation most consistent with the empirical observations is a behavioral underreaction model where information is only slowly absorbed into market prices. These results are consistent with earlier work in both experimental asset markets (Gillette, Stevens, Watts, & Williams, 1999; Kirchler, 2009; Page & Siemroth, 2021; Stevens & Williams, 2004; Weber & Welfens, 2007) and in empirical work on the cross-section of expected stock returns (Abarbanell & Bernard, 1992; Bernard & Thomas, 1989; L. K. Chan, Jegadeesh, & Lakonishok, 1996; Hui & Yeung, 2013).

The contribution of this chapter is twofold. First, it adds to the sports betting literature by providing evidence of time series momentum in pre-game bookmaker odds. Second, it contributes to the broader asset pricing literature by providing an indication that time series momentum in financial markets could be driven at least partly by behavioral forces and by underreaction more specifically. The sports betting setting is an interesting empirical lab as terminal values of the assets, i.e. the outcomes of the bets, are observable and independent of bettor behavior (Moskowitz, 2021; Thaler & Ziemba, 1988). Furthermore, as we know the terminal values of the assets, we can distinguish between underreaction (which is characterized by a continued drift) and overreaction (which is characterized by a reversal). Generalizing results obtained in sports betting to capital markets should be done cautiously. However, similar patterns in seemingly unconnected markets could expose fundamental symmetries. A common criticism to behavioral theories is that "allowing for irrationality opens a Pandora's box of ad hoc stories that will have little out-of-sample predictive power"

(Daniel et al., 1998, p. 1841). A general behavioral theory should be able to explain cognitive glitches irrespective of whether sports bets or capital market securities are involved.

The rest of this chapter is structured as follows. Section 4.1 introduces the context. Section 4.2 discusses the institutional setting and the dataset. In section 4.3, we carry out return predictability tests. In section 4.4 time series momentum strategies are tested while section 4.5 discusses the results and section 4.6 concludes.

4.1 Context

Whether past returns predict future returns is a question that is probably older than financial economics itself. Early empirical work showed that the dependence of returns on their own past is "either extremely slight or else non-existent" (Fama, 1965, p. 90). Historical returns were believed to contain no information that could be profitably leveraged in a trading strategy as expected in an informationally efficient market. Since seminal work by Jegadeesh and Titman (1993) and more recently by Moskowitz, Ooi, and Pedersen (2012), interest in both cross-sectional and time-series momentum respectively has been rejuvenated. Fundamentally, both types of momentum refer to the observation that historical returns predict future returns. Momentum is typically cross-sectional, meaning that assets which have been outperforming relative to other assets keep on outperforming in the future. Time series momentum or trend momentum directly tries to predict future returns of an individual asset based on its past returns.

Significant cross-sectional and time series momentum premia have been found in major asset classes including stocks, government bonds, corporate bonds, currencies and commodities (Asness et al., 2013; Georgopoulou & Wang, 2016; Jostova, Nikolova, Philipov, & Stahel, 2013; Moskowitz et al., 2012). Momentum is also remarkably persistent over time as it shows up in long run historical data sets (Annaert & Mensah, 2014; Geczy & Samonov, 2015; Goetzmann & Huang, 2018). This meaningful return predictability based on past returns is awkward as it seems to imply that markets are not even weakly efficient.

Despite truckloads of empirical evidence on the existence of momentum, its origin puzzles financial economists and is still heavily debated. Multiple, sometimes mutually exclusive, explanations for this "premier anomaly" (Fama & French, 2008, p. 1653) have been proposed[49]. A first possibility is that momentum simply does not exist in the real world. Momentum could for example be the result of data mining (Jegadeesh & Titman, 2001), but this is very unlikely given the magnitude of the empirical evidence. Another possibility is that momentum profits exist on paper, but not in practice because of transaction costs (Lesmond, Schill, & Zhou, 2004; Patton & Weller, 2020).

[49] Many models are developed for momentum in the cross-sectional sense. However they also pertain to individual assets and thus are relevant for time series momentum as well (Moskowitz et al., 2012).

The most common explanations however are either risk related or behavioral. A risk premium for assets that have been performing well would fit the standard asset pricing paradigm if good past performance increases an asset's risk but this is counterintuitive (Lewellen, 2002) although some rational stories have been proposed (Galariotis, 2013; Johnson, 2002; Li, 2018). Another explanation fitting the rational paradigm is that the momentum premium is a compensation for systematic skewness (Harvey & Siddique, 2000). Behavioralists argue that agents do not properly react to new information and that under- and/or overreaction to news could explain momentum (e.g. Barberis, Shleifer, and Vishny (1998), Daniel et al. (1998), Frazzini (2006)).

Empirically testing these theories is generally hindered by the joint hypothesis problem (Fama, 1991). If prices deviate from theoretical models, it is not clear whether the prices or the models are wrong, or both. To make matters worse, both risk and behavioral forces could be at work at the same time making it harder to distinguish between different alternatives. Betting markets could provide a way out of this gridlock (Moskowitz, 2021; Thaler & Ziemba, 1988). Betting markets have traditionally received substantial attention by economists because they can be regarded as "simple financial markets" (Sauer, 1998, p. 2021). These markets are interesting asset pricing labs for several reasons (Vandenbruaene, De Ceuster, & Annaert, 2022).

First, uncertainty is quickly resolved exogenously from bettor behavior and beliefs. Sports bets are typically short-term contracts with maturities of days to just minutes after which the terminal values of the assets are revealed. This allows researchers to detect systematic mispricing and makes betting markets "better suited for testing market efficiency and rationality" (Thaler & Ziemba, 1988, p. 162) compared to traditional financial markets. Furthermore, the short maturity of the contracts can enhance efficiency as agents can quickly get feedback on their decisions which experimental research shows is important to eliminate mispricing (Forsythe et al., 1982; Haruvy, Lahav, & Noussair, 2007). Second, bets take place in a vacuum, i.e. the payoffs are arguably uncorrelated with aggregate risks that drive the returns of traditional financial assets. This lack of correlation between the payoffs of bets and marginal utility means that the classic risk factors are not applicable to the pricing of sports bets (Moskowitz, 2021). Third, as betting markets are real markets where real money is at risk, external validity concerns that tend to plague experimental results are alleviated. Moreover, the size of the available data, both cross-sectionally and in the time series, is multiple orders of magnitude larger than would be economically feasible in an experimental setting. As a result of these useful research characteristics, a sizable literature on the efficiency of sports betting markets has been developed for many sports including American football (Gray & Gray, 1997), baseball (Woodland & Woodland, 1994), basketball, (Berkowitz et al., 2015), horse racing (L. V. Williams & Paton, 1997) and football (Croxson & Reade, 2014)[50].

Momentum strategies in sports betting have been tested previously (for an overview see Vandenbruaene et al. (2022)). However, almost all strategies that have been proposed take past game performance as the trading signal. Such strategies bet for example on a team that won its last 5 games which is somewhat analogous to earnings momentum strategies in the traditional empirical work in finance. Much less attention has been given to time series momentum strategies that take the evolution of the prices (the odds in a betting context) in the run-up to the game as the trading signal. There are a few notable exceptions[51] that use the evolution between the first odds (opening odds) and the last odds before the game starts (closing odds), but none of them evaluates the full time series of pre-game odds which is the focus of this chapter and which much more closely resembles the time series momentum tests carried out with traditional financial assets.

4.2 Setting, Data & Descriptive Statistics

Contemporary European football betting is mainly organized via fixed-odds betting. This means that bookmakers, who are market makers, offer odds for game outcomes on which gamblers can bet. An important characteristic of this microstructure is that the potential payouts are fixed and known when a bet is made, which is not the case in parimutuel betting popular in horse racing. This does not mean the odds cannot change over time, it just means that when a gambler chooses to place a bet, the odds at that time are locked in and known by the gambler. In this chapter, we use the decimal/European odds convention. These decimal odds represent the payout on a winning unit bet. For example, a gambler betting 1 on the home team when the home team odds are 1.66 gets a payoff of 1.66 when the home team wins (i.e. a return of 66%) and a payoff of 0 when the home team does not win. Note that odds are inversely related to outcome probabilities: high (low) odds reflect a low (high) probability event.

More formally, let o_{ijmt} be the decimal odds for outcome *i*, set by bookmaker *j*, for game *m* at time *t* where $i = 1$ refers to a home win, $i = 2$ refers to a draw and $i = 3$ refers to an away win. The return on a home win bet r_{1jmt+1} is given as follows:

$$r_{1jmt+1} = \begin{cases} (o_{1jmt} - 1) & \text{if home goals} > \text{away goals ("home win")} \\ -100\% & \text{if home goals} = \text{away goals ("draw")} \\ -100\% & \text{if home goals} < \text{away goals ("away win").} \end{cases}$$

The returns on draw and away win bets are analogous.

[50] It is worth noting that parallel to the establishment of the empirical literature on sports bets, an experimental literature that uses the betting microstructures (like parimutuel betting) has been developed. See Noussair and Tucker (2013) for an overview.

[51] See for example Baryla Jr et al. (2007), Crafts (1985), Gandar et al. (1988) or Shank (2018).

Sports bets can thus be regarded as European binary options where the underlying is the outcome of a game. A home bet is in the money whenever the number of home goals is strictly larger than the number of away goals. Also note that the inverse odds are essentially state prices in the Arrow and Debreu (1954) framework.

To make their business profitable, bookmakers charge a transaction cost for their services. This transaction cost is implicitly embedded into the odds, i.e. bookmakers skew the odds into their favor by setting them lower than they would be if the bets were fair. The property that transaction costs are directly embedded into the prices makes life easy for us as the returns on bets are already the post-transaction cost returns. These transaction costs should result in a positive expected return for the bookmaker and a negative expected return for the gambler.

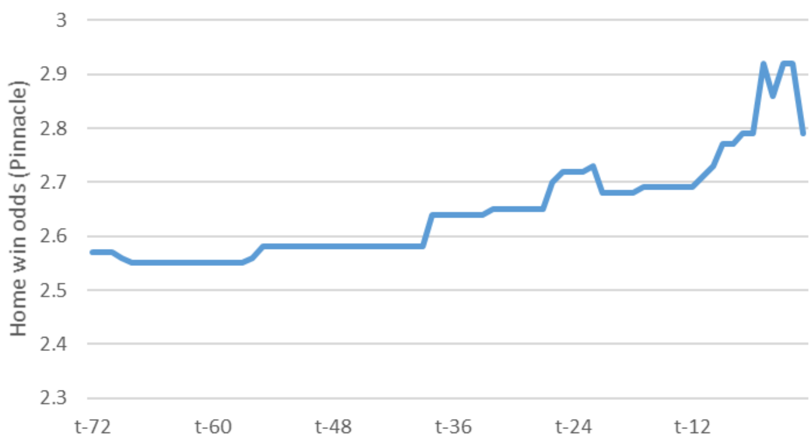

Figure 9: Liverpool – Manchester city (2/3/2016), hourly pre-game home win odds as quoted by Pinnacle.

As mentioned earlier, the data set consists of time series of pre-game odds on 17380 football games[52]. For each game, the home win, draw and away win odds of 32 bookmakers[53] are recorded every hour. In essence, we thus observe evaluations of the outcome probabilities of the games of 32 agents (bookmakers). Note that these evaluations are fully transparent as both bookmakers and gamblers can monitor the odds online in real-time, resulting in a competitive market. Each time series starts with the quoted odds 72 hours before the game and ends with the quoted odds 1 hour before the game. The bookmakers can update their odds as new information is released (like the starting lineup or the injury of a key player or because of significant volumes placed on one of the outcomes which can be an indication of the existence of insider information). As an example, Figure 9 shows the time series of pre-

[52] The data are collected by Kaunitz, Zhong, and Kreiner (2017) and are available on Kaggle: https://www.kaggle.com/austro/beat-the-bookie-worldwide-football-dataset#odds_series_b.csv.gz.

[53] These bookmakers are: Interwetten, bwin, bet-at-home, Unibet, Stan James, Expekt, 10Bet, William Hill, bet365, Pinnacle Sports, DOXXbet, Betsafe, Betway, 888sport, Ladbrokes, Betclic, Sportingbet, myBet, Betsson, 188BET, Jetbull, Paddy Power, Tipico, Coral, SBOBET, BetVictor, 12BET, Titanbet, youwin, ComeOn, Betadonis, Betfair Sports.

game home win odds quoted by bookmaker Pinnacle for a game between Liverpool and Manchester City on 2 March 2016. The odds were revised from 2.57 to 2.79 in the 72 hours before the game.

The fundamental value of the assets (the bets) thus fluctuates over time and is driven by the relative strengths and weaknesses of the teams involved and not by the experimental design. This is different from related work on under- and overreaction in experimental asset markets where the fundamental values of the assets typically decline deterministically over time as is the case in the seminal model of V. L. Smith et al. (1988). As an exception to this design, Stöckl, Huber, and Kirchler (2015) run an experiment in which the fundamental values fluctuate randomly over time. They find that the assets are overvalued (undervalued) when the fundamental values predominantly decline (increase), which is also consistent with our results reported in sections 4.3 and 4.4.

The games in the data set were played in 50 different football leagues between September 2015 and November 2016. Table 36 reports the leagues and game-related summary statistics of the games in the sample. The number of games per league varies because of some missing values, but mainly because of differences in league organization like the number of teams in the league and the timing of the games. On average, we have just under 350 games per league. It is interesting to point to the well-documented home field advantage. In all but one of the leagues, the probability of a home win is larger than the probability of an away win. On average home wins occur 50% more often than away wins. Lastly, the mean number of goals per match is 2.62.

To make the time series in our sample more manageable, we summarize the odds of the 32 individual bookmakers into market odds $\bar{o}_{imt}$ which could be interpreted as the market price of the respective outcome[54]. These market odds are the trimmed cross-sectional average odds for outcome i over all bookmakers j for match m at time t. We remove the three bookmakers with the highest and lowest odds to prevent that our results are driven by extreme values that are potentially pricing errors. (Analyses with different trim choices are similar.)

Table 37 reports summary statistics of the market odds $\bar{o}_{imt}$. We have approximately three million market odds observations distributed uniformly across the three outcome categories (home win, draw, away win). Note that the mean and median home odds (2.46 and 2.15) are significantly lower than the corresponding away odds (3.98 and 3.29), which reflects the strong home field advantage. Furthermore, it is worth pointing to the standard deviations and the extrema of the different odds categories. Draw market odds live in the smallest interval while away win market odds vary most wildly.

[54] These market odds are not directly tradable. In our trading strategies, we will use these market odds as an information signal but make bets on odds of individual bookmakers to make sure the strategies are implementable.

Table 36: Descriptive statistics of the data set.

	League	Number of games	Home win %	Draw %	Away win %	Mean goals per match
1	Argentina: Primera Division	510	46.86%	28.43%	24.71%	2.32
2	Austria: Tipico Bundesliga	206	47.57%	24.27%	28.16%	2.61
3	Belarus: Vysshaya Liga	274	38.69%	28.10%	33.21%	2.45
4	Belgium: Jupiler League	341	49.85%	24.63%	25.51%	2.88
5	Belgium: Proximus League	294	41.84%	25.85%	32.31%	2.94
6	Brazil: Série A	499	52.51%	24.05%	23.45%	2.45
7	Bulgaria: A PFG	122	46.72%	25.41%	27.87%	2.37
8	Chile: Primera Division	288	44.79%	27.78%	27.43%	2.88
9	China: Super League	273	49.08%	26.01%	24.91%	2.65
10	Colombia: Liga Aguila	523	48.57%	27.53%	23.90%	2.37
11	Croatia: 1. HNL	209	43.06%	27.27%	29.67%	2.23
12	Cyprus: First Division	286	41.96%	26.22%	31.82%	2.69
13	Czech Republic: Synot liga	186	51.61%	22.58%	25.81%	2.76
14	Denmark: Superliga	260	45.38%	25.00%	29.62%	2.76
15	Ecuador: Serie A	311	47.59%	26.05%	26.37%	2.63
16	England: Championship	654	43.43%	28.59%	27.98%	2.46
17	England: League One	661	42.81%	25.26%	31.92%	2.66
18	England: League Two	652	38.50%	26.53%	34.97%	2.67
19	England: Premier League	434	43.32%	27.42%	29.26%	2.75
20	France: Ligue 1	443	45.15%	27.31%	27.54%	2.58
21	France: Ligue 2	451	40.80%	32.59%	26.61%	2.29
22	Germany: 2. Bundesliga	359	43.73%	27.02%	29.25%	2.69
23	Germany: Bundesliga	355	46.48%	23.38%	30.14%	2.81
24	Greece: Super League	287	47.04%	27.53%	25.44%	2.23
25	Israel: Ligat ha'Al	278	38.85%	28.06%	33.09%	2.31
26	Italy: Serie A	457	45.73%	25.60%	28.67%	2.62
27	Italy: Serie B	587	46.51%	29.98%	23.51%	2.40
28	Japan: J-League	373	37.80%	23.86%	38.34%	2.61
29	Kazakhstan: Premier League	229	50.22%	19.65%	30.13%	2.49
30	Mexico: Primera Division	407	44.23%	29.24%	26.54%	2.81
31	Netherlands: Eredivisie	386	43.26%	25.65%	31.09%	2.90
32	Norway: Tippeligaen	299	46.82%	24.75%	28.43%	2.82
33	Paraguay: Primera Division	303	37.29%	23.43%	39.27%	2.97
34	Poland: Ekstraklasa	346	42.20%	28.61%	29.19%	2.64
35	Portugal: Primeira Liga	357	42.58%	24.09%	33.33%	2.66
36	Romania: Liga 1	304	41.78%	28.29%	29.93%	2.52
37	Russia: Premier League	283	45.94%	28.62%	25.44%	2.33
38	Scotland: Premiership	261	40.23%	25.29%	34.48%	2.83
39	Serbia: Super Liga	346	45.95%	26.30%	27.75%	2.38
40	Slovakia: Fortuna liga	236	51.27%	22.03%	26.69%	2.56
41	Slovenia: Prva liga	213	35.21%	27.70%	37.09%	2.43
42	South Korea: K-League Classic	290	38.62%	28.28%	33.10%	2.65
43	Spain: Primera Division	344	48.55%	23.55%	27.91%	2.79
44	Spain: Segunda Division	423	43.74%	31.68%	24.59%	2.21
45	Sweden: Allsvenskan	288	46.53%	22.57%	30.90%	3.10
46	Switzerland: Super League	208	47.12%	25.00%	27.88%	3.24
47	Turkey: Super Lig	351	47.86%	25.64%	26.50%	2.64
48	Ukraine: Pari-Match League	212	40.57%	23.11%	36.32%	2.65
49	Uruguay: Primera Division	291	39.18%	25.77%	35.05%	2.57
50	USA: MLS	430	51.40%	27.91%	20.70%	2.81
	Sum	17380				
	Average	347.60	44.53%	26.19%	29.28%	2.62

Table 37: Summary statistics for all market odds, home win market odds, draw market odds and away win market odds. The market odds are the trimmed hourly cross-sectional averages of the odds (3 highest and lowest odds are removed).

	All market odds	Home win market odds	Draw market odds	Away win market odds
N	3559957	1186728	1186471	1186758
Mean	3.34	2.46	3.58	3.98
Median	3.15	2.15	3.34	3.29
Standard deviation	1.92	1.39	0.82	2.7
Skewness	4.97	4.65	4.91	4.19
Kurtosis	49.04	37.4	40.47	30.61
Minimum	1.02	1.02	1.60	1.02
Maximum	55.94	33.76	18.74	55.94

Table 38 provides summary statistics of both absolute and percentage hourly market odds changes, the latter will be our variable of interest in the next section. First, notice that on average market odds do not change. The mean and median values are all very close to zero, both for absolute and percentage changes. However, from time to time very dramatic odds changes occur between two consecutive odds, as indicated by the extreme values on the last two rows of the two panels in Table 38. Double digit percentage increases and decreases in the market odds occur for all outcome categories.

Table 38: Summary statistics for changes in all market odds, home win market odds, draw market odds and away win market odds. The changes are the percentage and absolute change in market odds, which are the trimmed hourly cross-sectional averages of the odds (3 highest and lowest odds are removed).

	All market odds	Home win market odds	Draw market odds	Away win market odds
N percentage changes	3,507,701	1,169,309	1,169,053	1,169,339
Mean percentage changes	0.02%	0.01%	0.01%	0.03%
Median percentage changes	0.00%	0.00%	0.00%	0.00%
Standard deviation percentage changes	0.96%	0.94%	0.53%	1.27%
Skewness percentage changes	9.54	7.76	5.65	8.85
Kurtosis percentage changes	998.42	772.83	399.83	739.52
Minimum percentage change	-53.57%	-53.57%	-25.85%	-52.00%
Maximum percentage change	185.93%	133.08%	55.95%	185.93%
N absolute changes	3,507,701	1,169,309	1,169,053	1,169,339
Mean absolute changes	0.00	0.00	0.00	0.00
Median absolute changes	0.00	0.00	0.00	0.00
Standard deviation absolute changes	0.06	0.04	0.02	0.08
Skewness absolute changes	6.05	7.02	3.85	4.40
Kurtosis absolute changes	1915.86	3657.17	621.90	895.92
Minimum absolute change	-8.52	-8.52	-1.86	-8.00
Maximum absolute change	8.69	6.27	2.50	8.69

The dynamism in pregame odds is more pronounced when we consider the total change between the opening odds and the closing odds as depicted in Figure 10. From these figures, we learn that bookmakers do not systematically revise their odds in either direction in the runup to a game. This could be an indication that without the arrival of new information, the odds 72 hours before the start are equally indicative of the outcome of the game as the odds just an hour before the game as would be expected in an efficient market. However, notice that remarkable changes in odds, both in relative and absolute terms, occur. For percentage changes, we witness extrema as low as -71.78% and as high as 248.74%, in absolute terms, changes in market odds of -14.33 to 22.56 are observed which are clearly economically meaningful (remember that these market odds are already trimmed to minimize pricing errors in the data). As shown in Figure 11, in absolute value terms, the odds[55] tend to change the most in the hours just before the game. This is intuitively reasonable as important information like starting lineups is released close to the start of the game.

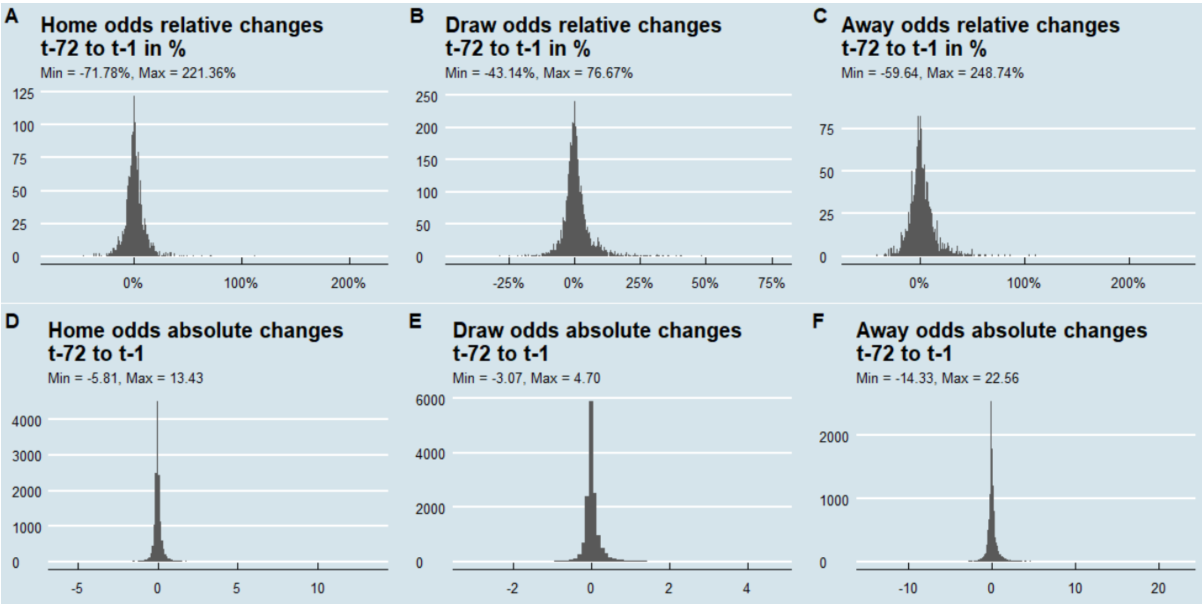

Figure 10: Histograms of the percentage changes in market odds from t-72 to t-1 per outcome category (home win, draw, away win) are shown in panels A to C while histograms of absolute changes are shown in panels D to F.

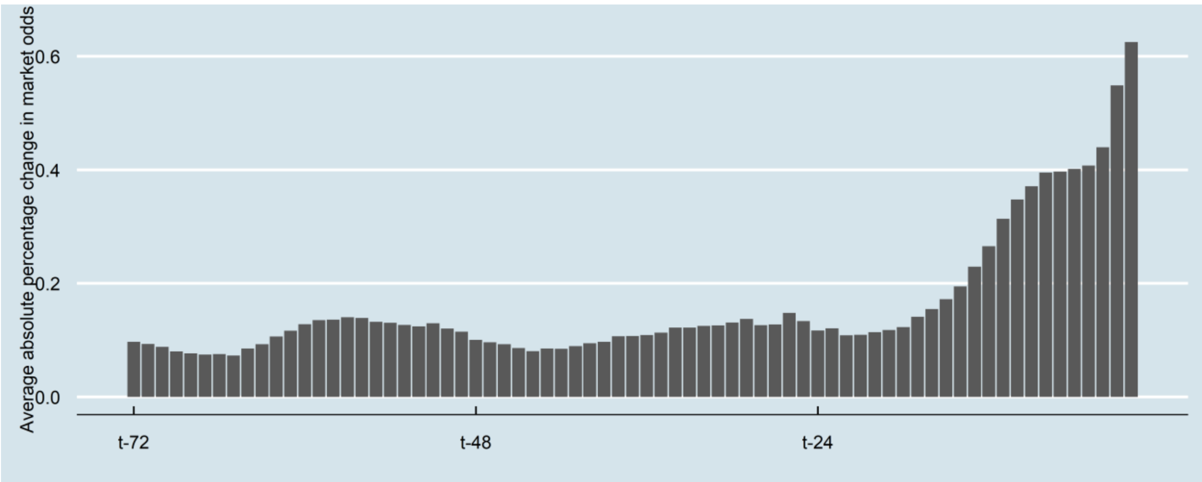

Figure 11: Average of the absolute value of the percentage changes in away market odds per hour.

[55] Figure 11 is based on away market odds. The dynamics in home and draw market odds are almost identical.

Figure 12 shows the average returns of home, draw and away bets made every hour in the runup to the game. Notice that there are significant differences between the returns of the different outcome states. The time series average of the home returns is -6.77%, the averages for draw and away odds are -8.09% and -12.78% respectively. This appears to be a reflection of the favorite-longshot bias, i.e. the empirical regularity in betting markets that returns on low odds bets significantly outperform returns on high odds bets (Ottaviani & Sørensen, 2008). As already indicated in Table 37, home odds are generally lower than draw odds, which are again lower than the away odds. The differences within the time series are much smaller than between the time series. T-tests that compare the mean returns at t-72 and t-1 of each of the time series reveal that none of the differences are statistically significant (results not reported).

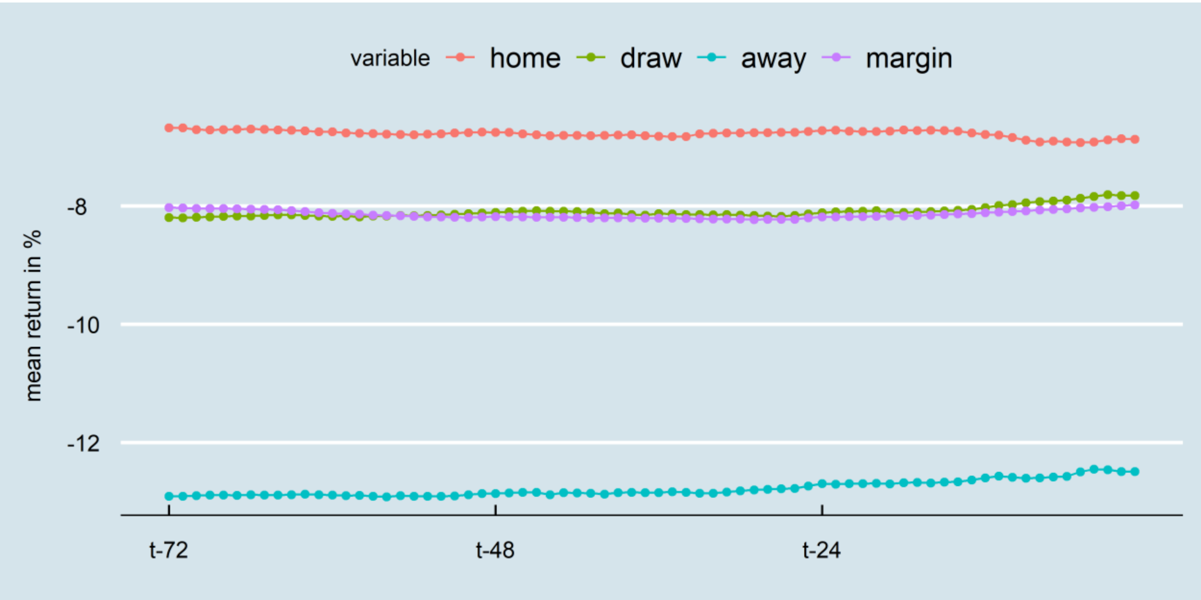

Figure 12: Returns in % of bets on home, draw and away outcomes over time. For example, the first blue dot indicates the mean return of betting on the away team at t-72. The margin is computed as the inverse of the booksum as discussed in chapter 3.

4.3 Return predictability

In an efficient betting market, information is instantaneously and adequately incorporated into the pre-game odds. In such a market, past price movements do not predict future returns. To test for return predictability, we estimate the following equation:

$$r_{imt} = \alpha_k + \beta_k \Delta\bar{o}_{imt-k}^{t-1} + \beta_m League_m + \varepsilon_{imt} \quad (1)$$

where r_{imt} is the return achieved by betting on the market odds for outcome i in match m at time t. We will refer to the period between t and the end of the game as the investment window. $\Delta\bar{o}_{imt-k}^{t-1}$ represents the relative change in market odds between $t-k$ and $t-1$, referred to as the information window, for outcome i in match m with $k > 1$. We end the information window at $t-1$ instead of t to avoid information spillovers[56]. We control for the league the game is played in as bookmakers could charge different transaction costs in leagues that receive a lot of attention like the English Premier League versus more humble competitions like the Vysshaya Liga in Belarus.

In an efficient market, we expect β_k to be indistinguishable from zero.

Hypothesis 1a: $\beta_k = 0$: Efficient sports betting market.

Alternatively, bookmakers do not incorporate information efficiently into their odds. In this case, odds change too little or too much depending on whether bookmakers under- or overreact. Because terminal values of sports bets are observable (we know the outcome of the game), we can easily distinguish between under- and overreaction. Suppose important information 24 hours before the game induces the home odds to decline, as depicted in panel A of Figure 13. This means the probability of a home win has gone up. In an efficient market, the odds move from o_a to o_b. If bookmakers underreact to the news, the odds will not move down by the full extent but end up at o_{b1}. These odds are a good deal for gamblers as they are higher than they should be, these odds are too high given the outcome probability. If bookmakers overreact to news, the odds will move to o_{b2}. These odds are not a good deal for gamblers as they are lower than they should be given the outcome probability. Analogously, in panel B of Figure 13 information 24 hours before the game induces the home odds to increase. This means that the probability of a home win has decreased. In an efficient market the home odds move from o_a to o_b. If bookmakers underreact, the odds will not move up by the full extent and end up at o_{b1}. These odds are a bad deal as they are lower than they should be. If bookmakers overreact to news, the odds will move to o_{b2}. These odds are a good deal as they are higher than they should be.

We can conclude that if $\beta_k > 0$ in equation (1), a past increase (decrease) of the odds increases (decreases) subsequent returns. This happens when bookmakers overreact. If $\beta_k < 0$ in equation (1), a past increase (decrease) in the odds decreases (increases) subsequent returns. This happens when bookmakers underreact. We summarize this in the following hypotheses.

Hypothesis 1b: $\beta_k > 0$: Inefficient sports betting market where agents overreact

Hypothesis 1c: $\beta_k < 0$: Inefficient sports betting market where agents underreact

[56] For robustness, we repeat the analysis with an information window ending at $t-2$. The results are similar and shown in appendix.

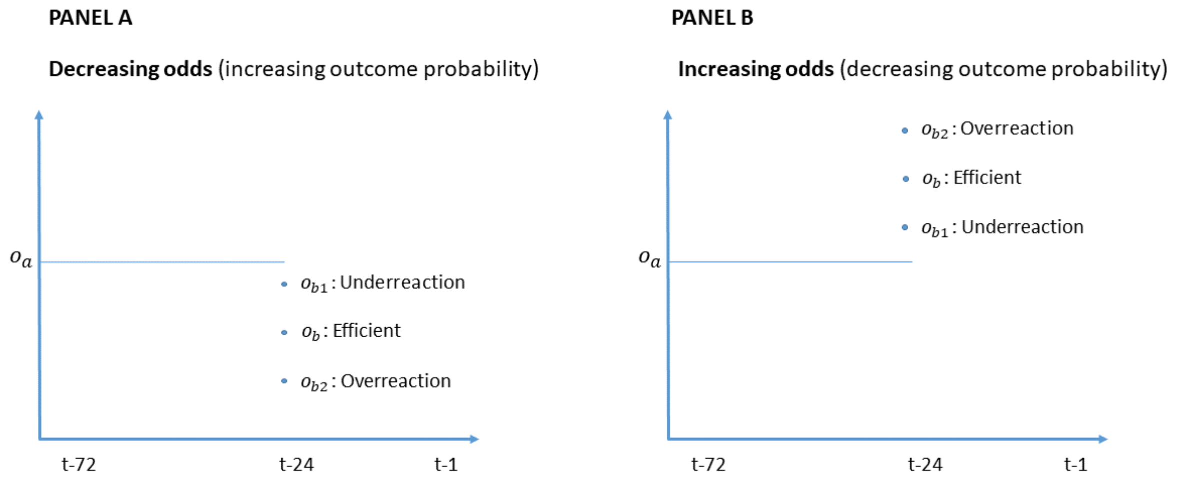

Figure 13: Efficient (o_b), under- (o_{b1}) and overreaction (o_{b2}) to news at t-24

To avoid data mining issues, we run regression (1) for all possible information and investment windows, i.e. all $t-k$ to $t-1$ intervals that are feasible. By doing this we get a full overview of the possible combinations of information and investment windows and avoid cherry-picking values for k and t that work by accident. Note that we have 72 hourly odds observations per game, resulting in $\frac{72\times71}{2} = 2556$ unique information windows. As we require at least 1 hour between the end of the information window and the start of the investment window, 71 information windows become uninvestable, resulting in 2485 regressions. We estimate heteroskedasticity and autocorrelation consistent standard errors via the Newey-West procedure[57].

Table 39 summarizes the results of the regressions. We estimate the regressions for home, draw and away odds individually to bypass the correlation between odds on the same game (i.e. if home odds move down, away odds have to move up to keep the transaction cost constant). Overall, the evidence is heavily tilted towards market inefficiencies induced by underreaction. $\widehat{\beta_1}$ is negative in over 93% of cases for home odds, draw odds and away odds. Moreover, it is not uncommon to observe coefficients smaller than -1, meaning that a change in odds results in a more than proportional change in returns. This is especially the case for draw odds where almost 50% of coefficients are smaller than -1. Next to being economically significant, many of the coefficients are also statistically significant. The number of t-stats smaller than -1.96 hovers between 15% and 42% and extreme t-stats under -4 are observed. Furthermore, none of the positive t-stats is larger than 1.96.

[57] For the lag parameter, we use $T^{1/4} = 72^{1/4} \approx 3$. The results for other standard error estimation choices (e.g. clustering by league) are similar as shown in appendix.

Table 39: Overview of the $\widehat{\beta_1}$ coefficients and t statistics of regression (1). The table shows the number of regressions with coefficients and t-stats in the respective range.

Coefficient	**Home**	**Draw**	**Away**	t-stat	**Home**	**Draw**	**Away**
b > 0.1	123	72	95	t > 1.96	0	0	0
0 < b < 0.1	58	18	41	-1.96 < t < 1.96	2026	1451	2111
-0.1 < b < 0	259	13	182	-2.5 < t < -1.96	267	594	205
-0.5 < b < -0.1	1551	181	1800	-3 < t < -2.5	120	375	103
-1 < b < -0.5	286	1062	293	-4 < t < -3	67	65	65
b < -1	208	1139	74	t < -4	5	0	1

An objection that can be made is that we run many regressions and that our regressions are clearly not independent. In terms of statistical significance, Chen (2020) shows[58] that for an arbitrary correlation structure the expected number of rejections of the null is given by:

$$E[Number\ of\ observed(|t_i| > \bar{t})] = Np(|Z| > \bar{t}),$$

where t_i are the t-statistics of N tests, $\bar{t}$ is a critical value and Z a standard normal random variable. For the common critical value of $\bar{t} = -1.96$, we would only expect $2485 \times 0.05 = 124$ rejections of the null per odds category. In contrast, we respectively find 459, 1034 and 374 such results. The results are even more extreme when we would consider the number of t-stats below -3 or -4 compared to what we would expect. Another way to place the results in context is by applying the Bonferroni correction for multiple hypothesis tests. In this case, we scale the significance level α by the number of tests: $\alpha^{Bonferroni} = \frac{\alpha}{N}$. As we run 2485 regressions, we evaluate the results at $\alpha^{Bonferroni} = 2 \times 10^{-5}$, which corresponds to a t-stat of -4.26. Given that we observe t-stats smaller than this extreme value for both home and away odds, these results are statistically significant even when we take the large number of tests into account. Similar results are obtained via other popular multiple hypotheses tests like the Holm correction or the Benjamini, Hochberg and Yekutieli (BHY) correction. Notice furthermore that the Bonferroni correction is especially harsh in our context as many of our t-stats are highly correlated. For example, suppose we run two regressions of which the underlying data is almost perfectly correlated. This would result in $\alpha^{Bonferroni} = \frac{0.05}{2}$ which is clearly too extreme as we essentially run the same regression twice. As a result, the Bonferroni hurdle rate we apply is an extreme upper bound to test statistical significance in our context.

In summary, the results make a case for significant return predictability driven by underreaction, or in other words:

$$E(r_{imt}|\Delta\bar{o}_{imt-k}^{t-1} > 0) < E(r_{imt}|\Delta\bar{o}_{imt-k}^{t-1} < 0).$$

[58] Under the assumption that the marginal distribution of t-statistics is standard normal under the null.

Note that these results are hard to reconcile with a rational risk story. The returns increase on bets for which the odds have decreased, i.e. which became less risky. For bets that have become riskier, the returns go down. This is the opposite of what one would expect in a framework where predictable patterns in returns must be induced by risk factors.

As a robustness exercise we follow Moskowitz et al. (2012) and estimate a variation to regression (1) where we check whether just the sign of the past returns contains predictive information.

$$r_{imt} = \alpha_k + \beta_k sign(\Delta\bar{o}_{imt-k}^{t-1}) + \beta_m League_m + \varepsilon_{imt} \qquad (2)$$

The results of regression (2) are displayed in Table 40 and are in line with the earlier discussion.

Table 40: Overview of the $\widehat{\beta_1}$ coefficients and t statistics of regression (2). The table shows the number of regressions with coefficients and t-stats in the respective range.

Coefficient	**Home**	**Draw**	**Away**	t-stat	**Home**	**Draw**	**Away**
b > 0.1	20	144	61	t > 1.96	0	1	0
0 < b < 0.1	4	18	13	-1.96 < t < 1.96	1598	1419	1725
-0.1 < b < 0	6	18	7	-2.5 < t < -1.96	575	594	445
-0.5 < b < -0.1	55	103	94	-3 < t < -2.5	249	349	185
-1 < b < -0.5	223	148	255	-4 < t < -3	63	113	111
b < -1	2177	2054	2055	t < -4	0	9	19

4.4 Momentum portfolios

In the previous section we put forward evidence that bookmakers underreact to new information which is in line with earlier research in experimental asset markets (e.g. Kirchler (2009); Weber and Welfens (2007)). In this section, we move beyond and analyze whether the documented return predictability can be monetized in a real betting market via a time series momentum strategy. To do so, we make two portfolios, one for "winners" and one for "losers". We define winners as the bets of which the odds have decreased i.e. the probabilities of the underlying events have increased. Losers are defined as the bets of which the odds have increased i.e. the probabilities of the underlying events have decreased. In an efficient market betting on winners will yield the same return as betting on losers as all bets are priced correctly. However, if bookmakers underreact, the returns on winners are expected to be higher than the returns on losers as prices do not respond swiftly but continue to drift towards their efficient levels.

In Figure 14, we show in blue the mean return across all M games for every hour t and every odds type i for all observations. That is, every blue dot is:

$$\frac{1}{M}\sum_{m=1}^{M} r_{imt},$$

for its respective value of i and t. Similarly, the mean returns for the observations of which the odds decreased between $t-2$ and $t-1$ are shown in red. That is, every red dot is:

$$\frac{1}{M}\sum_{m=1}^{M} r_{imt} \mid \Delta\bar{o}_{imt-2}^{t-1} < 0$$

for its respective value of i and t. Lastly, the mean returns for the observations of which the odds increased between $t-2$ and $t-1$ are shown in green. That is, every green dot is:

$$\frac{1}{M}\sum_{m=1}^{M} r_{imt} \mid \Delta\bar{o}_{imt-2}^{t-1} > 0$$

for its respective value of i and t. We compute the returns at t instead of $t-1$ to prevent information spillovers. The group means are shown in black in Figure 14 together with their 99% confidence intervals. Notice that the mean returns of the observations with decreasing odds are systematically higher than the unconditional means, which are again higher than the means of the observations with increasing odds. Furthermore, the differences between the means of the increasing and decreasing odds are statistically significant for home odds, draw odds and away odds (the t-statistics are respectively -8.72, -3.29 and-7.34). These results are consistent with the regression analysis from the previous section and again point in the direction of underreaction.

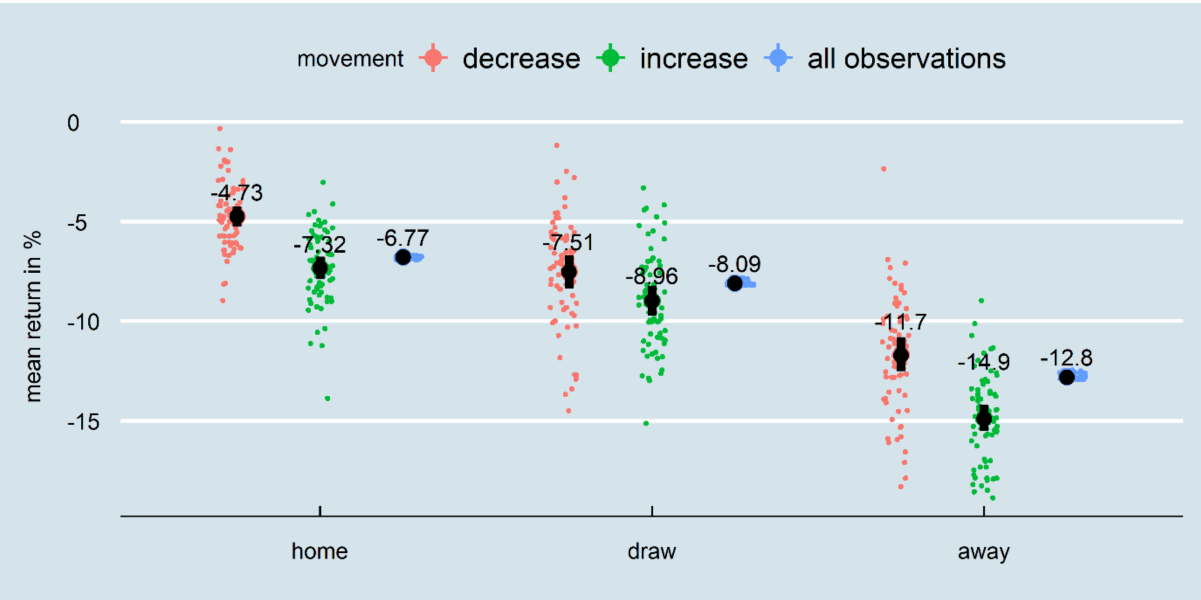

Figure 14: Mean returns for every t leading up to the game. Mean returns per hour for all observations in blue, mean returns for the observations of which the odds decreased between $t-2$ and $t-1$ in red and the mean returns for the observations of which the odds increased between $t-2$ and $t-1$ in green. The black dots are the group means and the black bars the 99% confidence intervals.

To further investigate whether the differences are exploitable, we set up the following trading strategy. We compute odds changes

$$\Delta\bar{o}_{imt-k}^{t-1}$$

for all feasible values for t and k and generate a trading signal whenever $|\Delta\bar{o}_{imt-k}^{t-1}| > |c|$ for the first time in the respective time series i.e. at the smallest possible value for $t-1$ where the critical value condition is met. If the condition is met, we bet at the most favorable odds offered by the 32 bookmakers in the sample at time t. This setup eliminates information spillovers in both the information window and investment window and is implementable in real time. Furthermore, the market odds are not directly tradable, so we only use these odds to generate a trading signal and trade at odds offered directly in the market.

Table 41: Performance of the momentum strategies when bets are made at the best available odds. c is a critical value for the relative change in odds during the time series. For example, c = -5% means betting on all odds that have decreased by at least 5% in the runup to the game. The difference in means is tested via t-tests, the difference in variances is tested via the non-parametric Fligner-Killeen test, the difference in the skewness is tested via a bootstrap setup. The stars indicate whether the results are significant at the 10%, 5% and 1% level respectively.

Panel A: HOME ODDS

c	Number of bets	Mean return	Standard deviation	Skewness	c	Number of bets	Mean return	Standard deviation	Skewness	Difference means	p-value means	p-value variances	p-value skewness
-5%	6901	4.54%	139.95%	1.93	5%	7573	-3.15%	144.83%	2.00	7.69%	0.12%***	2.60%**	79.40%
-10%	2877	8.80%	154.98%	2.39	10%	3756	-5.03%	158.36%	2.41	13.83%	0.04%***	0.32%***	94.60%
-15%	1282	10.03%	174.24%	2.83	15%	2070	-6.54%	173.01%	2.54	16.57%	0.74%***	0.58%***	68.00%
-20%	645	11.09%	178.23%	2.36	20%	1202	-3.37%	195.44%	3.09	14.46%	10.86%	2.84%**	4.40%**

Panel B: DRAW ODDS

c	Number of bets	Mean return	Standard deviation	Skewness	c	Number of bets	Mean return	Standard deviation	Skewness	Difference means	p-value means	p-value variances	p-value skewness
-5%	2842	-1.97%	181.65%	1.75	5%	4224	-6.42%	182.15%	1.82	4.45%	31.38%	96.53%	73.00%
-10%	657	-4.88%	200.63%	2.21	10%	1660	-9.60%	190.10%	1.90	4.72%	60.44%	72.43%	80.00%
-15%	226	-0.38%	216.95%	2.08	15%	748	-18.98%	186.74%	2.08	18.60%	24.47%	15.13%	50.60%
-20%	82	17.93%	239.01%	1.85	20%	370	-21.51%	188.03%	2.17	39.44%	16.42%	5.75%*	41.80%

Panel C: AWAY ODDS

c	Number of bets	Mean return	Standard deviation	Skewness	c	Number of bets	Mean return	Standard deviation	Skewness	Difference means	p-value means	p-value variances	p-value skewness
-5%	8520	-2.08%	188.79%	2.82	5%	9746	-8.69%	191.83%	3.33	6.61%	1.93%**	0.32%***	35.80%
-10%	4200	-2.08%	192.60%	2.83	10%	5668	-11.19%	208.80%	4.76	9.11%	2.50%**	0.04%***	28.80%
-15%	2137	-0.45%	201.62%	2.74	15%	3548	-10.11%	233.31%	5.31	9.66%	9.93%*	0.06%***	7.80%*
-20%	1093	2.70%	204.44%	2.37	20%	2413	-11.38%	222.65%	3.16	14.08%	6.64%*	0.15%***	0.00%***

Summary statistics of the momentum strategy are shown in Table 41. We consider 4 different values for $|c|$ (5, 10, 15 & 20) for home odds, draw odds and away odds, resulting in 12 different "momentum portfolios". First, note again that in every case, the mean returns on bets of which odds have decreased are larger than the mean returns on bets of which the odds which have increased. This is consistent with our previous analyses and with underreaction. This difference is often substantial and statistically significant. For example, bets of which the odds have decreased by at least 10% in the runup to a game outperform bets of which the odds have increased by at least 10% by 13.83 percentage points (t-stat of 3.57). Moreover, this difference increases almost monotonically when the critical values go up. Second, it is striking that the returns of many implementations are positive after transaction costs which indicates that the inefficiency could be profitably exploited. For example, the simple strategy of making a unit bet every time the home odds move down by

20% has an expected return of 11.09%. For draw odds the results are even more extreme: betting when draw odds move down by more than 20% has an expected return of 17.93%. However, keep in mind that this strategy only bets on 82 games out of 17380 games in the sample.

Table 42: Performance of the momentum strategies when bets are made at median odds. c is a critical value for the relative change in odds during the time series. For example, c = -5% means betting on all odds that have decreased by at least 5% in the runup to the game. The difference in means is tested via t-tests, the difference in variances is tested via the non-parametric Fligner-Killeen test, the difference in the skewness is tested via a bootstrap setup. The stars indicate whether the results are significant at the 10%, 5% and 1% level respectively.

Panel A: HOME ODDS

c	Number of bets	Mean return	Standard deviation	Skewness	c	Number of bets	Mean return	Standard deviation	Skewness	Difference means	p-value means	p-value variances	p-value skewness
-5%	6901	-4.52%	124.32%	1.59	5%	7573	-10.60%	129.97%	1.65	6.08%	0.41%***	17.99%	77.00%
-10%	2877	-3.80%	131.57%	1.96	10%	3756	-13.77%	138.72%	1.96	9.97%	0.28%***	6.03%*	91.80%
-15%	1282	-5.92%	140.63%	2.31	15%	2070	-16.97%	147.13%	2.01	11.05%	3.00%**	4.99%**	69.30%
-20%	645	-8.13%	137.47%	1.75	20%	1202	-16.38%	158.14%	2.33	8.25%	24.43%	13.28%	1.40%**

Panel B: DRAW ODDS

c	Number of bets	Mean return	Standard deviation	Skewness	c	Number of bets	Mean return	Standard deviation	Skewness	Difference means	p-value means	p-value variances	p-value skewness
-5%	2842	-10.34%	164.07%	1.59	5%	4224	-14.20%	164.98%	1.64	3.86%	33.43%	82.73%	64.80%
-10%	657	-16.41%	173.02%	2.00	10%	1660	-18.22%	170.40%	1.81	1.81%	82.03%	95.72%	74.30%
-15%	226	-17.22%	174.84%	1.86	15%	748	-28.25%	163.91%	1.99	11.03%	39.97%	31.31%	60.40%
-20%	82	-6.09%	184.12%	1.65	20%	370	-30.16%	165.71%	2.07	24.07%	27.79%	14.00%	29.60%

Panel C: AWAY ODDS

c	Number of bets	Mean return	Standard deviation	Skewness	c	Number of bets	Mean return	Standard deviation	Skewness	Difference means	p-value means	p-value variances	p-value skewness
-5%	8520	-12.72%	163.75%	2.49	5%	9746	-18.94%	164.44%	2.71	6.22%	1.06%**	0.36%***	50.60%
-10%	4200	-15.25%	162.32%	2.49	10%	5668	-22.84%	170.65%	2.88	7.59%	2.46%**	0.20%***	39.80%
-15%	2137	-16.40%	164.24%	2.38	15%	3548	-23.92%	183.07%	3.28	7.52%	10.94%	0.37%***	7.00%*
-20%	1093	-15.28%	164.42%	2.18	20%	2413	-25.19%	181.05%	2.77	9.91%	10.96%	0.66%***	1.40%**

An objection that could be made is that the differences in returns are rational risk compensations. In this scenario, strategies that are riskier should yield higher returns. However, this explanation seems unlikely as the variances of the returns of bets on decreasing odds are lower than those of the increasing odds in 8 out of 12 strategies and the differences are often statistically significant. Another argument that could be brought forward is that the differences in returns are the result of a skewness preference of gamblers. Empirical work on stock markets is often consistent with the hypothesis that investors like stocks with positive skewness (lottery-like characteristics) and are willing to pay a premium for them (Annaert et al., 2013). However, this explanation is not supported by our data as the skewness values of the returns of the different portfolios are seldomly statistically different from each other.

For robustness purposes, we repeat the analyses but now make bets at the median odds instead of maximum odds as show in Table 42. Again, the mean returns of bets at decreasing odds systematically outperform bets at increasing odds. This time however, none of the implementations has a positive mean return. The rational explanations (compensation for variance of skewness) are once again not supported.

4.5 Discussion

We document a strong time series momentum effect in betting markets that appears to be driven by underreaction of bookmakers. Although underreaction is indeed a prevalent behavioral fallacy in the literature, many behavioral models involve both underreaction and overreaction (e.g. Barberis et al. (1998) or Daniel et al. (1998)). There are a number of reasons why we could fail to document overreaction in our setting. First, underreaction simply appears to be a more persistent anomaly than overreaction. For example, Lin and Rassenti (2012) find that underreaction drifts substantially outnumber overreaction reversals in an experimental asset market. Similarly, Stevens and Williams (2004) find that underreaction is a common behavioral glitch and document that agents underreact to both negative and positive information in a controlled experimental setting. In their study, 24.8% of subjects underreacted to information more than 50% of the time while 0% of subjects overreacted more than 50% of the time, leading them to conclude that they do not find evidence of systematic overreaction. This is conisistent with Weber and Welfens (2007) who also find evidence of stock price underreaction but not of overreaction in an experimental setting.

Second, many behavioral models consist of short-term underreaction followed by longer-term overreaction. As our time series only contain 72 hourly observations, they might be too short to pick up reversal patterns. In the seminal work by De Bondt and Thaler (1985) for example, the observed overreaction effect is strongest between 3 and 5 years after portfolio formation. Such a timeframe is clearly of a different order of magnitude than we have in our data (and that of previous experimental work). In related work, Moskowitz (2021) finds evidence consistent with overreaction in a US sample of sportsbets. However, his momentum signals are related to the past performance of teams up to the last 8 games which is again a much longer timeframe than the one we consider in this study.

4.6 Conclusion

Does time series momentum also exist outside traditional financial markets? If we look at pregame sports bet odds, the answer is yes. The expected returns of “winners” are economically and statistically significantly higher than those of “losers”. This result is robust to a battery of different design choices. Moreover, some long only portfolios are profitable after transaction costs. Furthermore, we add to the growing evidence that agents systematically underreact to new information. The momentum patterns that we document are consistent with a prolonged drift towards efficient asset values.

Can we further generalize our results? It is obvious that betting markets are in many respects very different than traditional financial markets. However, the fact that we document a similar momentum pattern in a seemingly unrelated environment can reveal a fundamental human behavioral glitch. Especially because it is consistent with earlier results in experimental asset markets.

4.7 Appendix to chapter 4

Table 43: Overview of the $\widehat{\beta_1}$ coefficients and t statistics of regression (1) when the standard errors are clustered by league. The table shows the number of regressions with coefficients and t-stats in the respective range.

Coefficient	**Home**	**Draw**	**Away**	t-stat	**Home**	**Draw**	**Away**
b > 0.1	123	72	95	t > 1.96	1	0	0
0 < b < 0.1	58	18	41	-1.96 < t < 1.96	2041	1265	2042
-0.1 < b < 0	259	13	182	-2.5 < t < -1.96	270	494	249
-0.5 < b < -0.1	1551	181	1800	-3 < t < -2.5	125	377	111
-1 < b < -0.5	286	1062	293	-4 < t < -3	47	337	75
b < -1	208	1139	74	t < -4	1	12	8

In the main text, we defined the information window as $\Delta\bar{o}_{imt-k}^{t-1}$, representing the relative change in market odds between $t-k$ and $t-1$ with $k>1$. As a robustness exercise, we change the gap between the information window and investment window from 1 to 2 hours: $\Delta\bar{o}_{imt-k}^{t-2}$ with $k>3$. The results are shown below and are similar to those discussed in the main text.

Table 44: Overview of the $\widehat{\beta_1}$ coefficients and t statistics of regression (1) with the information window ending at $t-2$ instead of $t-1$. The table shows the number of regressions with coefficients and t-stats in the respective range.

Coefficient	**Home**	**Draw**	**Away**	t-stat	**Home**	**Draw**	**Away**
b > 0.1	133	71	100	t > 1.96	1	0	0
0 < b < 0.1	51	21	44	-1.96 < t < 1.96	2022	1249	2012
-0.1 < b < 0	241	12	194	-2.5 < t < -1.96	245	463	225
-0.5 < b < -0.1	1519	179	1733	-3 < t < -2.5	113	378	105
-1 < b < -0.5	275	1026	276	-4 < t < -3	33	315	67
b < -1	198	1106	68	t < -4	1	10	6

To further drill down on the results of regression (1), we visualize the coefficients and t-statistics of all 2485 regressions for home odds, draw odds and away odds in Figure 15 to Figure 17 respectively. The start of the information period on the vertical axis represents $t-k$, the end of the information period on the horizontal axis represents $t-1$, the investment period starts at t and runs until the end of the game as discussed earlier.

The visual results again highlight that underreaction patterns are very prevalent while significant overreaction patters are virtually non-existent. Furthermore, the graphs point to notable differences in timing of the return predictability between the different outcome categories. For home and away odds, most statistically significant results are clustered either early in the time series or in information windows that end around 8 hours before the start of the game. As also shown in Table 39, the predictability patterns are especially outspoken for draw outcomes. For draws, a large majority of all information windows that end less than 24 hours before the start of the game show signs of significant underreaction. This appears consistent with earlier work showing that draw outcomes are especially hard to price correctly, which could also imply that agents have a hard time incorporating new information into draw odds (Dixon & Pope, 2004; Pope & Peel, 1989).

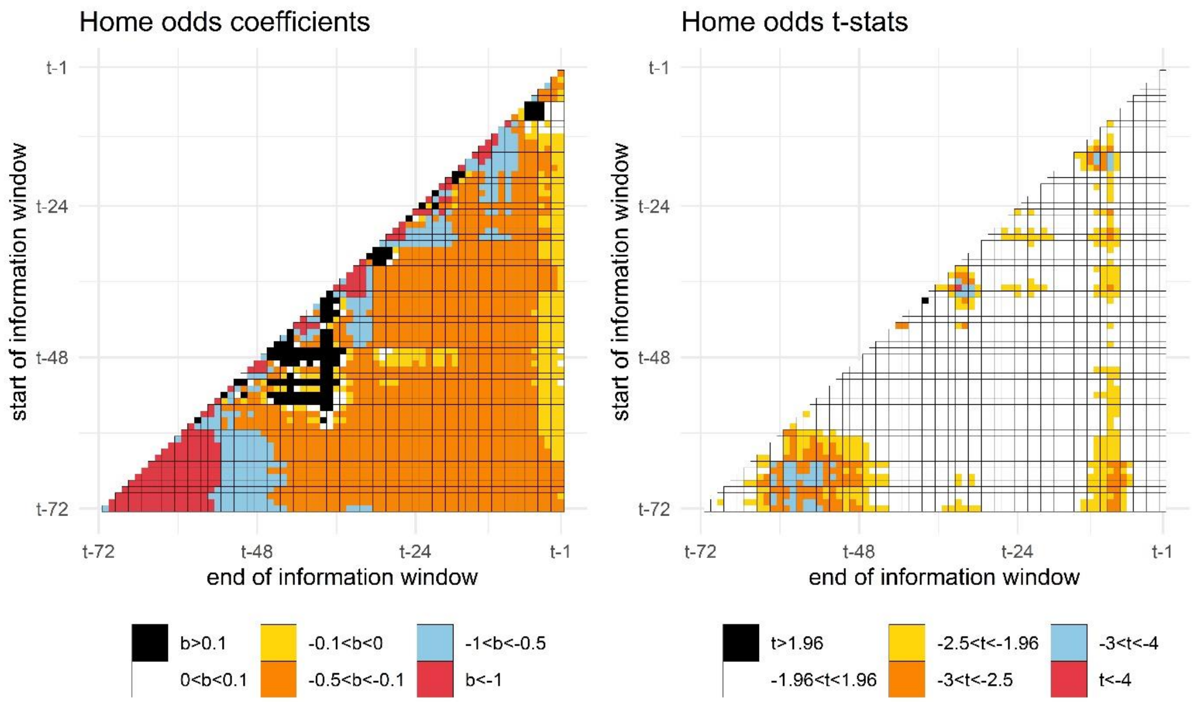

Figure 15: Coefficient estimates and t-statistics of β_1 from regression (1) for home odds. Every square of the lower triangular matrix represents a different investment window. In total 2485 regression coefficients and t-statistics are visualized.

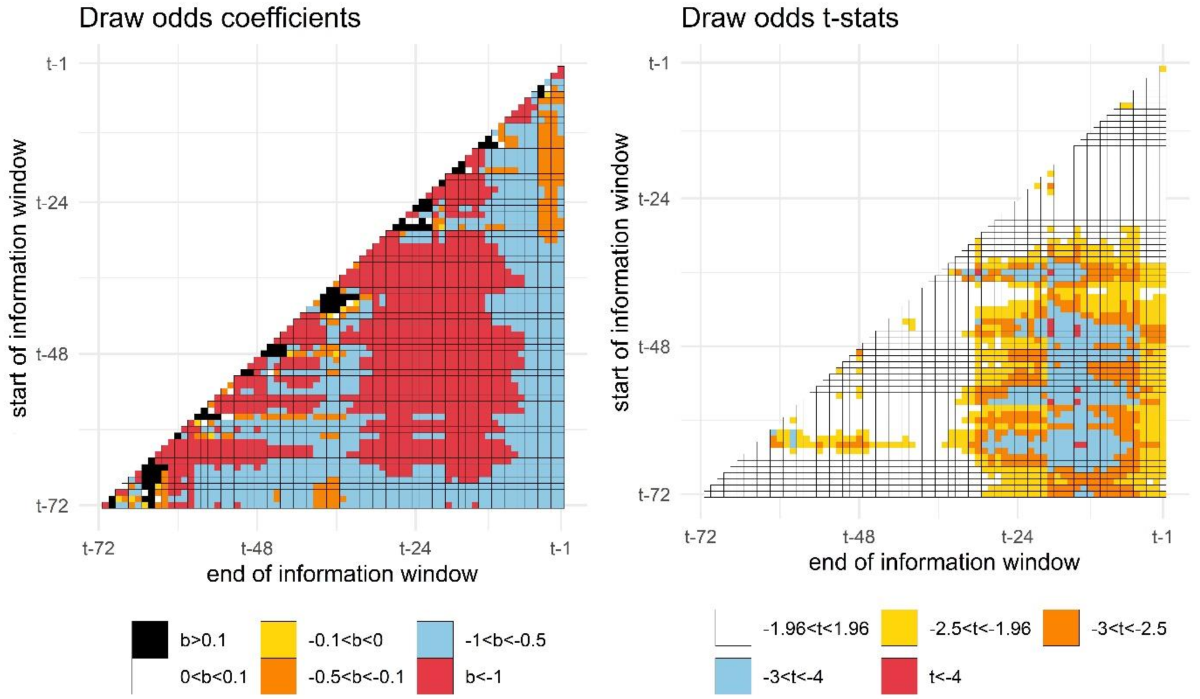

Figure 16: Coefficient estimates and t-statistics of β_1 from regression (1) for draw odds. Every square of the lower triangular matrix represents a different investment window. In total 2485 regression coefficients and t-statistics are visualized.

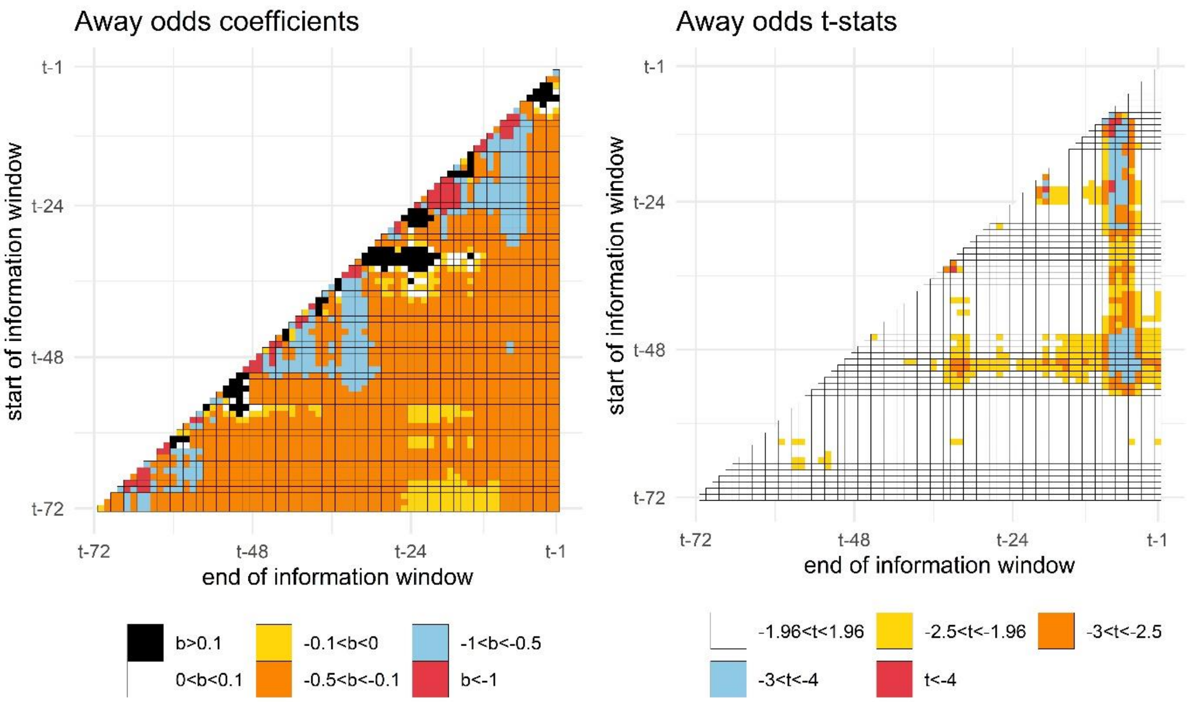

Figure 17: Coefficient estimates and t-statistics of β_1 from regression (1) for away odds. Every square of the lower triangular matrix represents a different investment window. In total 2485 regression coefficients and t-statistics are visualized

Chapter 5: Conclusions, Reflections and Avenues for Further Research

Academics who work on prediction markets often get asked what their results now actually imply for asset markets more generally. It is an important question to ask but not a straightforward one to answer. On the one hand, it is easy to get excited about the similarities. For example, the mechanics of prediction markets are often identical to those of traditional financial markets. Continuous double auctions are used in both prediction markets on US presidential elections and on the world's largest stock exchanges. Bookmakers have a role that is very similar to the market makers supplying liquidity in financial derivatives markets. If the plumbing, the hardware, is identical, could we expect that what we find for prediction markets is also true for traditional financial markets?

On the other hand, it is easy to get distracted by the differences. Millions of individuals and companies trade financial instruments for a variety of reasons including saving for their pensions, hedging financial risks, gambling on meme stocks or exploiting arbitrage restrictions. The stakes on these markets are enormous. In 2020 the total volume of stock trading alone is estimated at $138 trillion (World Federation of Exchanges, 2020). In contrast, a critic could argue that prediction markets are virtually non-existent compared to these gargantuan global financial markets and that they are mainly populated by geeks and beer-drinking sports fans who are looking for the thrill of speculating on a future event.

The ultimate guardian of informational efficiency on any market is the profit incentive. The important question is whether this incentive is strong enough on prediction markets to discipline market prices. The relevant literature firmly points in the direction that it is, which underpins the usefulness of prediction markets as an empirical research lab. Given our battery of arguments on why attaining market efficiency is easier in prediction markets, would it be a stretch to interpret prediction market results as an upper bound of informational efficiency on capital markets? If simple prediction markets would not be able to digest information efficiently, arguably all hope would be lost for our extremely complex global financial markets.

5.1 Some concluding remarks on the empirical chapters

Chapter 2 argued that the empirical work on trading strategies in prediction markets has become an elaborate fishing expedition. Most of the papers we reviewed scored bad on three metrics. First, all of them benchmarked their results at the single hypothesis test benchmark. This is not even correct for the papers individually as they tend to test many variations of the same strategy and it is even less appropriate given how many strategies have been tested throughout the years. Second, authors often do not explicitly refer to previous findings in behavioral economics or in psychology which would render their strategy test worthy in the first place. Third, the strategies are often only tested in just 1 league and only for a limited time period. In a response to these observations, a useful rubric would be to only write papers on strategies that fulfill at least two out of the following three criteria:

1) have a test statistic of at least 3
2) have explicit reference to earlier established biases
3) are tested in multiple leagues/countries/sports over a considerable period of time.

In chapter 3, we studied the evolution of transaction costs in football betting. A takeaway for future research from this chapter is that the transaction cost in fixed-odds betting is a multidimensional construct. Although transaction costs are often proxied by the sum of the inverse odds, this method will underestimate (overestimate) the true transaction cost on very large (small) odds. A solution would be to measure the transaction costs by the average return per decile as we did in table 33. However, this method requires a lot of data and is time sensitive as transaction costs change over time. Furthermore, notice that researchers often derive implied probabilities from odds by using the sum of the inverse odds which suffers from a similar bias. Lastly, chapter 3 warns us that controlling for the size of the odds might not be sufficient to precisely measure transaction costs. Factors like the importance of the game affect transaction costs as well.

In chapter 4, we brought forward evidence that bookmakers systematically underreact to new information. We looked at this in the aggregate and did not address individual news stories. A follow-up study could focus more closely what the changes in the information set actually were. Furthermore, the reaction could depend on variables such as the topic, the importance or the unexpectedness of the news.

5.2 Ideas for future research

An implicit stylized fact of the prediction market literature is that bubbles are very rare. This is surprising as bubbles and crashes are rather common in controlled asset market lab experiments. An avenue for further research could be to investigate what drives this apparent difference. One hypothesis could be that prediction market traders are more experienced and understand the situation better than the average undergraduate student who accidentally stumbles into an asset market experiment. Dufwenberg, Lindqvist, and Moore (2005) for example find that the frequency of bubble formation in experiments is greatly reduced if at least a third of participants have prior experience with such an experiment. Kirchler, Huber, and Stöckl (2012) document that mispricing and overvaluation is significantly less present when they frame their experiment such that it is very easy to understand for participants. Another hypothesis could be that the design of the contracts on prediction markets (i.e. binary options) eases efficiency. Whether the price of Bitcoin should be \$0, \$10,000 or \$100,000 is hard to say. The value is sensitive to subjective assumptions such that it can easily take on extreme values. In contrast, even the largest Real Madrid supporter will not feel comfortable betting money on their team at odds which imply Real Madrid has a 99.5% probability of winning their next league game. Is there something about the interpretability of prices as probabilities which disciplines traders to not stray too far off the true value?

Another unexplored area is prediction market volatility. For mainstream financial markets, Shiller (1981) famously argues that the observed levels of stock market volatility cannot be rationally justified by changes in dividends alone. Relatedly, Cutler, Poterba, and Summers (1989) document that the most volatile days for the S&P 500 were days on which no important news was released which puzzles efficient market theorists. As argued earlier, discount rate news does not influence prediction market prices. Furthermore, there is much less cash flow news and the news is often more transparent and accessible than that of listed companies. As a result, an intelligently designed prediction market study could revamp the old question of what drives changes in prices. A related stylized fact on financial markets is volatility clustering: small changes tend to be followed by small changes while large changes tend to be followed by large changes. This observation inspired academics to develop the GARCH toolkit to model the phenomenon. To explain volatility clustering, Engle (2004) points to clusters in news or to the price discovery process where traders are uncertain about their estimates and update their beliefs based on what others do. Again, as changes in the information set are presumably much clearer in prediction markets, we could use them to study volatility clustering from a different angle.

So far, we have been talking about avenues for future research. What about new use cases for prediction markets as a policy tool? An interesting area to think about the usefulness of prediction markets are the subsidy desks of major countries or supranational organizations. The European Union for example heavily subsidizes research and innovation through its Horizon Future program, which has a budget of €95 billion. Making sure that this money is well spent requires a huge bureaucracy and the decisions are eventually driven by judgement calls of a few individuals. Could the probability of success and impact of the proposed projects be better estimated by creating a pseudo capital market with public money? An intelligently designed prediction market where the incentives of the project applicants, the EU and the market participants are neatly aligned could increase allocative efficiency of the subsidy process, make the decision making more transparent and could cut administrative costs.

Another use case that could be further investigated is the hedging of specific (geo)political risks via prediction markets. The revenue of many companies is to some extent a function of the outcome of elections. For example, fossil fuel companies and arms producers arguably do much better under a conservative government while electric car manufacturers and renewable energy providers are much better off in a regime that subsidizes climate efforts. As the trading volumes on major political events are large, a considerable amount of risk could be hedged by using a syndicate of prediction market platforms. Yagudin et al. (2021) for example report that prices on prediction market FTX were consistent with a large actor hedging against the cancellation of the Tokyo Olympics because of the pandemic. However, the academic research on corporate hedging via prediction markets is virtually nonexistent at this point.

Ultimately, this dissertation argues and hopefully convinces the reader that prediction markets are a useful empirical laboratory for social sciences generally. So far, the usefulness of prediction markets might be somewhat underestimated as much of the literature is organized around trying to find profitable strategies and the predictive performance of these markets. The accumulated evidence quite compellingly demonstrates that consistently making money is very hard and that prediction markets predict future events very well. Let us stay away from these evident questions for a while and focus on work that tries to expose symmetries between what we talk about in traditional asset pricing and prediction markets. Some recent examples include Moskowitz and Vasudevan (2021) who try to link the favorite-longshot bias to the low risk anomaly documented in financial markets or Chiappori, Salanié, Salanié, and Gandhi (2019) who show that standard utility functions like CRRA or CARA and even the expected utility framework more generally do not model prediction market data well.

Given the noisiness of much of the data economists work with, testing whether results obtained in traditional financial markets also hold in prediction markets is a good out of sample test. From this perspective, prediction markets, just as lab experiments, are a useful addition to the toolkit of an economist.

5.3 Epilogue

We started this dissertation by discussing markets and their role in society in a very general sense. Let us close the loop and add some final thoughts in this brief epilogue.

Much of the empirical work that we carried out is consistent with market efficiency. In chapter 2, we put 40 years of literature in perspective and find a strong relationship between the sample size and the profitability of trading strategies. In contrast to the conclusions from some of the individual papers in the review, our analysis is consistent with an informationally efficient market. In chapter 3, we found that technological advances can increase efficiency by reducing anomalies. In chapter 4, we documented that market participants appear to underreact to new information, but that this behavior is not exploitable on average. In general, we find that prediction markets are good at absorbing information into their prices. In a sense, this conclusion is not surprising as it is in line with much of the theoretical and empirical work in the prediction market sphere and also with the consensus in financial economics in general.

The question that of course remains is how we should translate these academic conclusions to policy recommendations. Do these results spill over to debates on how we best organize our society? Practitioners and policy makers are often quick to jump to conclusions. Very influential think tanks like the Heritage Foundation regularly quote Hayek or Fama while making the case for free markets and low taxes. They bring forward the academic literature as a “proof” that markets work great and are best left alone. The point on which I want to end this dissertation is that although market efficiency organizes a lot of the empirical work in finance, we should not forget that it is only one of many dimensions that society cares about.

Whether markets function well is an unworkable question so we operationalize it by focusing on informational efficiency. However, finding that market prices are not predictably wrong can of course not be generalized to the idea that we should organize our society by a free market system. Again, the prediction market literature is a great case to make this point.

One of the implicit puzzles in the corporate prediction market literature is why so few companies actually use them. Prediction markets virtually always come out on top when compared to alternatives, so why do they have such a hard time conquering the world? The answer of course is that looking only at informational efficiency, which financial economists often do, is very myopic. Firms care about much more than just optimizing their forecasting accuracy. Installing a prediction market for example also leads to employees trying to monetize their own private information which makes them less collegial. Employees who repeatedly lose money can become demoralized and start to see their coworkers as competitors and so on. These dimensions are not as easily defined and measured as for example the mean absolute error (MAE), so they are often entirely disregarded.

The same is true for society at large. For some economic activities like capital allocation, informational efficiency is arguably the most important metric and competitive markets are the best vessel to achieve this. For most others, it is much less clear. Simply optimizing for efficiency because that is the dimension we talk about all the time might not be a good idea.

References

Abarbanell, J. S., & Bernard, V. L. (1992). Tests of analysts' overreaction/underreaction to earnings information as an explanation for anomalous stock price behavior. *The Journal of Finance, 47*(3), 1181-1207. doi:https://doi.org/10.1111/j.1540-6261.1992.tb04010.x

Abel, C. F., & Kunz, K. (2018). Unsettling the Memes of Neoliberal Capitalism through Administrative Pragmatism. In *From Austerity to Abundance?* (Vol. 6, pp. 35-58): Emerald Publishing Limited.

Ali, M. M. (1979). Some Evidence of the Efficiency of a Speculative Market. *Econometrica, 47*(2), 387-392. doi:https://doi.org/10.1142/9789812819192_0044

Alipourfard, N., Arendt, B., Benjamin, D. J., Benkler, N., Bishop, M., Burstein, M., . . . Clark, C. (2021). Systematizing Confidence in Open Research and Evidence (SCORE). doi:https://doi.org/10.31235/osf.io/46mnb

American Gaming Association. (2019). Americans Will Wager $8.5 Billion on March Madness [Press release]. Retrieved from https://www.americangaming.org/new/americans-will-wager-8-5-billion-on-march-madness/

American Gaming Association. (2020). A Record 26 Million Americans Will Wager on Super Bowl LIV [Press release]. Retrieved from https://www.americangaming.org/new/a-record-26-million-americans-will-wager-on-super-bowl-liv-2/

Amoako-Adu, B., Marmer, H., & Yagil, J. (1985). The efficiency of certain speculative markets and gambler behavior. *Journal of Economics and Business, 37*(4), 365-378. doi:https://doi.org/10.1016/0148-6195(85)90028-1

Andolfatto, D., Hendry, S., & Moran, K. (2008). Are inflation expectations rational? *Journal of Monetary Economics, 55*(2), 406-422. doi:https://doi.org/10.1016/j.jmoneco.2007.07.004

Andres, T. (2018). Divided Decade: How the financial crisis changed housing. Retrieved from https://www.marketplace.org/2018/12/17/what-we-learned-housing/

Angelini, G., & De Angelis, L. (2019). Efficiency of online football betting markets. *International Journal of Forecasting, 35*(2), 712-721. doi:https://doi.org/10.1016/j.ijforecast.2018.07.008

Annaert, J., De Ceuster, M., & Verstegen, K. (2013). Are extreme returns priced in the stock market? European evidence. *Journal of Banking & Finance, 37*(9), 3401-3411. doi:https://doi.org/10.1016/j.jbankfin.2013.05.015

Annaert, J., & Mensah, L. (2014). Cross-sectional predictability of stock returns, evidence from the 19th century Brussels Stock Exchange (1873–1914). *Explorations in Economic History, 52*, 22-43. doi:https://doi.org/10.1016/j.eeh.2013.10.002

Appelbaum, B. (2019). *The Economists' Hour*: Little, Brown and Company.

Arrow, K. J., & Debreu, G. (1954). Existence of an equilibrium for a competitive economy. *Econometrica: Journal of the Econometric Society*, 265-290. doi:https://doi.org/1907353

Arrow, K. J., Forsythe, R., Gorham, M., Hahn, R., Hanson, R., Ledyard, J. O., . . . Zitzewitz, E. (2008). The Promise of Prediction Markets. *science, 320*(5878), 877. doi:https://doi.org/10.1126/science.1157679

Asch, P., Malkiel, B. G., & Quandt, R. E. (1982). Racetrack betting and informed behavior. *Journal of Financial Economics, 10*(2), 187-194. doi:https://doi.org/10.1016/0304-405X(82)90012-5

Asch, P., Malkiel, B. G., & Quandt, R. E. (1984). Market efficiency in racetrack betting. *Journal of Business*, 165-175. doi:https://doi.org/10.1086/296257

Ashman, T., Bowman, R. A., & Lambrinos, J. (2010). The role of fatigue in NBA wagering markets: The surprising ‘‘Home Disadvantage Situation’’. *Journal of Sports Economics, 11*(6), 602-613.

Asness, C. S., Moskowitz, T. J., & Pedersen, L. H. (2013). Value and momentum everywhere. *The Journal of Finance, 68*(3), 929-985. doi:https://doi.org/10.1111/jofi.12021

azBookmakers. (2019). Biggest Bookmakers In The World. Retrieved from https://azbookmakers.com/biggest-bookmakers-in-the-world

Barberis, N., Shleifer, A., & Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics, 49*(3), 307-343. doi:https://doi.org/10.1016/S0304-405X(98)00027-0

Baryla Jr, E. A., Borghesi, R. A., Dare, W. H., & Dennis, S. A. (2007). Learning, price formation and the early season bias in the NBA. *Finance Research Letters, 4*(3), 155-164. doi:https://doi.org/10.1016/j.frl.2007.04.002

Bassett Jr, G. W. (1981). Point spreads versus odds. *Journal of Political Economy, 89*(4), 752-768. doi:https://doi.org/10.1086/261001

Baye, M. R., Morgan, J., & Scholten, P. (2006). Information, search, and price dispersion. *Handbook on economics and information systems, 1*, 323-375.

Beja, A. (1971). The structure of the cost of capital under uncertainty. *The Review of Economic Studies, 38*(3), 359-368. doi:https://doi.org/2296388

Benjamin, D. J., Berger, J. O., Johannesson, M., Nosek, B. A., Wagenmakers, E.-J., Berk, R., . . . Camerer, C. (2018). Redefine statistical significance. *Nature Human Behaviour, 2*(1), 6-10. doi:https://doi.org/10.1038/s41562-017-0189-z

Bennett, R. W. (2019). Holdover Bias in the College Football Betting Market. *Atlantic Economic Journal, 47*(1), 103-110. doi:https://doi.org/10.1007/s11293-019-09611-y

Bennett, R. W. (2020). Next game reaction to mispriced betting lines in college football. *Applied Economics Letters, 28*(12), 1006-1009. doi:https://doi.org/10.1080/13504851.2020.1795063

Benter, W. (2004). What Are My Odds?: ICCM.

Berg, J., Nelson, F. D., & Rietz, T. A. (2008). Prediction market accuracy in the long run. *International Journal of Forecasting, 24*(2), 285-300. doi:https://doi.org/10.1016/j.ijforecast.2008.03.007

Berg, J., & Rietz, T. (2006). The Iowa Electronic Markets: Stylized Facts and Open Issues, Information Markets: A New Way of Making Decisions in the Public and Private Sectors, Hahn, R., Tetlock, PC. In: AEI Press, Washington DC.

Berkowitz, J. P., Depken II, C. A., & Gandar, J. M. (2015). Information and accuracy in pricing: Evidence from the NCAA men's basketball betting market. *Journal of Financial Markets, 25*(C), 16-32. doi:https://doi.org/10.1016/j.finmar.2015.06.003

Berkowitz, J. P., Depken II, C. A., & Gandar, J. M. (2017). A favorite-longshot bias in fixed-odds betting markets: Evidence from college basketball and college football. *The Quarterly Review of Economics and Finance, 63*, 233-239. doi:https://doi.org/10.1016/j.qref.2015.11.011

Berkowitz, J. P., Depken II, C. A., & Gandar, J. M. (2018). Exploiting the "win but does not cover" phenomenon in college basketball. *Financial Review, 53*(1), 185-204. doi:https://doi.org/10.1111/fire.12155

Berlemann, M., & Vöpel, H. (2012). Tournament incentives and asset price bubbles: Evidence from a field experiment. *Economics Letters, 115*(2), 232-235. doi:https://doi.org/10.1016/j.econlet.2011.12.043

Bernard, V. L., & Thomas, J. K. (1989). Post-earnings-announcement drift: delayed price response or risk premium? *Journal of Accounting research, 27*, 1-36. doi:https://doi.org/2491062

Bernoulli, D. (1954). Exposition of a New Theory on the Measurement of Risk. *Econometrica, 22*(1), 23-36. doi:https://doi.org/10.2307/1909829

bet365. (n.d.). About us. Retrieved from https://www.bet365careers.com/about-us

bet365 Group Limited. (2018). *Report and Financial Statements*. Retrieved from

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy, 100*(5), 992-1026. doi:https://doi.org/10.1086/261849

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1998). Learning from the behavior of others: Conformity, fads, and informational cascades. *Journal of Economic Perspectives, 12*(3), 151-170. doi:https://doi.org/10.1257/jep.12.3.151

Bilotkach, V., Gorodnichenko, Y., & Talavera, O. (2010). Are airlines' price-setting strategies different? *Journal of Air Transport Management, 16*(1), 1-6. doi:https://doi.org/10.1016/j.jairtraman.2009.04.004

Bird, R., & McCrae, M. (1987). Tests of the Efficiency of Racetrack Betting Using Bookmaker Odds. *Management Science, 33*(12), 1552-1562. doi:https://doi.org/10.1287/mnsc.33.12.1552

Black, F., & Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy, 81*(3), 637-654. doi:https://doi.org/10.1086/260062

Blundell, J. (2008). *Margaret Thatcher: A Portrait of the Iron Lady*: Algora Pub.

Bolton, R. N., & Chapman, R. G. (1986). Searching for positive returns at the track: A multinomial logit model for handicapping horse races. *Management Science, 32*(8), 1040-1060. doi:https://doi.org/10.1287/mnsc.32.8.1040

Borghesi, R. (2007a). The home team weather advantage and biases in the NFL betting market. *Journal of Economics and Business, 59*(4), 340-354. doi:https://doi.org/10.1016/j.jeconbus.2006.09.001

Borghesi, R. (2007b). The late-season bias: explaining the NFL's home-underdog effect. *Applied Economics, 39*(15), 1889-1903. doi:https://doi.org/10.1080/00036840600690314

Borghesi, R., Paul, R., & Weinbach, A. (2009). Market frictions and overpriced favourites: evidence from arena football. *Applied Economics Letters, 16*(9), 903-906. doi:https://doi.org/10.1080/17446540802314519

Boulier, B. L., Stekler, H., & Amundson, S. (2006). Testing the efficiency of the National Football League betting market. *Applied Economics, 38*(3), 279-284. doi:https://doi.org/10.1080/00036840500368904

Boulier, B. L., & Stekler, H. O. (2003). Predicting the outcomes of National Football League games. *International Journal of Forecasting, 19*(2), 257-270. doi:https://doi.org/10.1016/S0169-2070(01)00144-3

Bowles, S., Kirman, A., & Sethi, R. (2017). Retrospectives: Friedrich Hayek and the market algorithm. *Journal of Economic Perspectives, 31*(3), 215-230.

Brown, A. (2014). Information processing constraints and asset mispricing. *The Economic Journal, 124*(575), 245-268. doi:https://doi.org/10.1111/ecoj.12057

Brown, A., & Reade, J. (2019). The wisdom of amateur crowds: Evidence from an online community of sports tipsters. *European Journal of Operational Research, 272*(3), 1073-1081. doi:https://doi.org/10.1016/j.ejor.2018.07.015

Brown, A., & Yang, F. (2016). Limited cognition and clustered asset prices: Evidence from betting markets. *Journal of Financial Markets, 29*, 27-46. doi:https://doi.org/10.1016/j.finmar.2015.10.003

Brown, J. R., & Goolsbee, A. (2002). Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry. *Journal of Political Economy, 110*(3), 481-507. doi:https://doi.org/10.1086/339714

Bruce, A. C., Johnson, J. E., Peirson, J. D., & Yu, J. (2009). An examination of the determinants of biased behaviour in a market for state contingent claims. *Economica, 76*(302), 282-303. doi:https://doi.org/10.1111/j.1468-0335.2008.00741.x

Brynjolfsson, E., & Smith, M. D. (2000). Frictionless Commerce? A Comparison of Internet and Conventional Retailers. *Management Science, 46*(4), 563-585. doi:https://doi.org/10.1287/mnsc.46.4.563.12061

Caldwell, B. (2013). *Hayek's The Road to Serfdom: A Brief Introduction*: University of Chicago Press.

Camerer, C. F. (1989). Does the Basketball Market Believe in the 'Hot Hand,'? *The American Economic Review, 79*(5), 1257-1261. doi:https://www.jstor.org/stable/1831452

Campbell, J. Y., Lo, A. W., & MacKinlay, A. C. (1997). *The Econometrics of Financial Markets*: Princeton University Press.

Carlin, B. I., Kogan, S., & Lowery, R. (2013). Trading complex assets. *The Journal of Finance, 68*(5), 1937-1960. doi:https://doi.org/10.1111/jofi.12029

Cassidy, J. (2013). *How markets fail: The logic of economic calamities*: Penguin UK.

Cavallo, A. (2017). Are online and offline prices similar? Evidence from large multi-channel retailers. *American Economic Review, 107*(1), 283-303. doi:https://doi.org/10.1257/aer.20160542

Chan, K. (1988). On the contrarian investment strategy. *Journal of Business, 61*(2), 147-163. doi:https://www.jstor.org/stable/2352897

Chan, L. K., Jegadeesh, N., & Lakonishok, J. (1996). Momentum strategies. *The Journal of Finance, 51*(5), 1681-1713. doi:https://doi.org/10.1111/j.1540-6261.1996.tb05222.x

Chandra, A., & Tappata, M. (2011). Consumer search and dynamic price dispersion: an application to gasoline markets. *The RAND Journal of Economics, 42*(4), 681-704. doi:http://dx.doi.org/10.2139/ssrn.1116054

Chen, A. Y. (2020). The Limits of p-Hacking: Some Thought Experiments. *The Journal of Finance*. doi:https://doi.org/10.1111/jofi.13036

Chiappori, P.-A., Salanié, B., Salanié, F., & Gandhi, A. (2019). From Aggregate Betting Data to Individual Risk Preferences. *Econometrica, 87*(1), 1-36. doi:https://doi.org/10.3982/ECTA11165

Cochrane, J. H. (2011a). How did Paul Krugman get it so wrong? . *Economic Affairs, 31*(2), 36-40. doi:https://doi.org/10.1111/j.1468-0270.2011.02077.x

Cochrane, J. H. (2011b). Presidential Address: Discount Rates. *The Journal of Finance, 66*(4), 1047-1108. doi:https://doi.org/10.1111/j.1540-6261.2011.01671.x

Coleman, B. J. (2017). Team travel effects and the college football betting market. *Journal of Sports Economics, 18*(4), 388-425. doi:https://doi.org/10.1177/1527002515574514

Colquitt, L. L., Godwin, N. H., & Shortridge, R. T. (2007). The effects of uncertainty on market prices: evidence from coaching changes in the NBA. *Journal of Business Finance & Accounting, 34*(5-6), 861-871. doi:https://doi.org/10.1111/j.1468-5957.2006.02000.x

Conlisk, J. (1993). The utility of gambling. *Journal of risk and uncertainty, 6*(3), 255-275. doi:https://doi.org/10.1007/BF01072614

Cowgill, B., Wolfers, J., & Zitzewitz, E. (2009). Using Prediction Markets to Track Information Flows: Evidence from Google. *AMMA, 3*. doi:https://doi.org/10.1007/978-3-642-03821-1_2

Cowgill, B., & Zitzewitz, E. (2015). Corporate prediction markets: Evidence from google, ford, and firm x. *The Review of Economic Studies, 82*(4), 1309-1341. doi:https://doi.org/10.1093/restud/rdv014

Crafts, N. F. (1985). Some evidence of insider knowledge in horse race betting in Britain. *Economica, 52*(207), 295-304. doi:https://doi.org/10.2307/2553853

Cronin, B. (2018). The history of Pinnacle. Retrieved from https://www.pinnacle.com/en/betting-articles/educational/the-history-of-pinnacle/358269QJKLJVQ2R8

Crotty, J. (2013). The Realism of Assumptions Does Matter: Why Keynes-Minsky Theory Must Replace Efficient Market Theory as the Guide to Financial Regulation Policy. In M. H. Wolfson & G. A. Epstein (Eds.), *The Handbook of the Political Economy of Financial Crises*. New York, NY: Oxford University Press.

Croxson, K., & Reade, J. (2014). Information and Efficiency: Goal Arrival in Soccer Betting. *The Economic Journal, 124*(575), 62-91. doi:https://doi.org/10.1111/ecoj.12033

Cutler, D. M., Poterba, J. M., & Summers, L. H. (1989). What moves stock prices? *The Journal of Portfolio Management, 15*(3), 4. doi:10.3905/jpm.1989.409212

Dana, J. D., & Knetter, M. M. (1994). Learning and Efficiency in a Gambling Market. *Management Science, 40*(10), 1317-1328. doi:https://doi.org/10.1287/mnsc.40.10.1317

Daniel, K., Hirshleifer, D., & Subrahmanyam, A. (1998). Investor Psychology and Security Market Under- and Overreactions. *The Journal of Finance, 53*(6), 1839-1885. doi:10.1111/0022-1082.00077

Dare, W. H., Dennis, S. A., & Paul, R. J. (2015). Player absence and betting lines in the NBA. *Finance Research Letters, 13*, 130-136. doi:https://doi.org/10.1016/j.frl.2015.02.004

Davis, J., & Krieger, K. (2017). Preseason bias in the NFL and NBA betting markets. *Applied Economics, 49*(12), 1204-1212. doi:https://doi.org/10.1080/00036846.2016.1213367

Davis, J., McElfresh, L., Krieger, K., & Fodor, A. (2015). Exploiting week 2 bias in the NFL betting markets. *Journal of Prediction Markets, 9*(1), 53-67. doi:https://doi.org/10.5750/jpm.v9i1.1014

De Bondt, F. M. W., & Thaler, R. (1985). Does the Stock Market Overreact? *The Journal of Finance, 40*(3), 793-805. doi:https://doi.org/10.1111/j.1540-6261.1985.tb05004.x

De Long, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). Noise trader risk in financial markets. *Journal of Political Economy, 98*(4), 703-738. doi:https://doi.org/10.1086/261703

Debreu, G. (1984). La Supériorité de l'Économie Libérale Est Incontestable et Mathématiquement Démontrable. *Le Figaro*.

Deutscher, C., Ötting, M., Schneemann, S., & Scholten, H. (2019). The demand for English premier league soccer betting. *Journal of Sports Economics, 20*(4), 556-579. doi:https://doi.org/10.1177/1527002518794803

Devlin, K. (2010). *The unfinished game: Pascal, Fermat, and the seventeenth-century letter that made the world modern*: Basic Books.

DiFilippo, M., Krieger, K., Davis, J., & Fodor, A. (2014). Early season NFL over/under bias. *Journal of Sports Economics, 15*(2), 201-211. doi:https://doi.org/10.1177/1527002512454544

Direr, A. (2011). Are betting markets efficient? Evidence from European football championships. *Applied Economics, 45*(3), 343-356. doi:https://doi.org/10.1080/00036846.2011.602010

Dixon, M. J., & Coles, S. G. (1997). Modelling Association Football Scores and Inefficiencies in the Football Betting Market. *Journal of the Royal Statistical Society: Series C (Applied Statistics), 46*(2), 265-280. doi:10.1111/1467-9876.00065

Dixon, M. J., & Pope, P. F. (2004). The value of statistical forecasts in the UK association football betting market. *International Journal of Forecasting, 20*(4), 697-711. doi:https://doi.org/10.1016/j.ijforecast.2003.12.007

Dowie, J. (1976). On the Efficiency and Equity of Betting Markets. *Economica, 43*(170), 139-150. doi:https://doi.org/10.1142/9789812819192_0052

Dufwenberg, M., Lindqvist, T., & Moore, E. (2005). Bubbles and experience: An experiment. *American economic review, 95*(5), 1731-1737. doi:https://doi.org/10.1257/000282805775014362

Edelman, D. C., & O'Brian, N. R. (2004). Tote arbitrage and lock opportunities in racetrack betting. *The European Journal of Finance, 10*(5), 370-378. doi:https://doi.org/10.1080/1351847042000199042

Elaad, G., Reade, J., & Singleton, C. (2020). Information, prices and efficiency in an online betting market. *Finance Research Letters, 35*, 101291. doi:https://doi.org/10.1016/j.frl.2019.09.006

Ellison, G., & Ellison, S. F. (2005). Lessons about Markets from the Internet. *The Journal of Economic Perspectives, 19*(2), 139-158. doi:https://doi.org/10.1257/0895330054048632

Ellison, G., & Ellison, S. F. (2009). Search, obfuscation, and price elasticities on the internet. *Econometrica, 77*(2), 427-452. doi:https://doi.org/10.3982/ECTA5708

Engle, R. (2004). Risk and volatility: Econometric models and financial practice. *American economic review, 94*(3), 405-420. doi:https://doi.org/10.1257/0002828041464597

Erikson, R. S., & Wlezien, C. (2008). Are political markets really superior to polls as election predictors? *Public Opinion Quarterly, 72*(2), 190-215. doi:https://doi.org/10.1093/poq/nfn010

European Commission. (2014). *Study on the Coverage, Functioning and Consumer Use of Comparison Tools and Third-Party Verification Schemes for Such Tools*. Retrieved from https://ec.europa.eu/info/sites/info/files/final_report_study_on_comparison_tools_2013_en.pdf

European Commission. (2021). EU Emissions Trading System (EU ETS). Retrieved from https://ec.europa.eu/clima/policies/ets_en

Even, W. E., & Noble, N. R. (1992). Testing efficiency in gambling markets. *Applied Economics, 24*(1), 85-88. doi:https://doi.org/10.1080/00036849200000106

Fair, R. C., & Shiller, R. J. (1990). Comparing information in forecasts from econometric models. *The American Economic Review*, 375-389.

Fama, E. F. (1965). The Behavior of Stock-Market Prices. *The Journal of Business, 38*(1), 34-105. doi:https://doi.org/10.1086/294743

Fama, E. F. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance, 25*(2), 383-417. doi:https://doi.org/10.2307/2325486

Fama, E. F. (1991). Efficient Capital Markets: II. *The Journal of Finance, 46*(5), 1575-1617. doi:10.1111/j.1540-6261.1991.tb04636.x

Fama, E. F. (1998). Market efficiency, long-term returns, and behavioral finance. *Journal of Financial Economics, 49*(3), 283-306. doi:https://doi.org/10.1016/S0304-405X(98)00026-9

Fama, E. F., & French, K. R. (1996). Multifactor explanations of asset pricing anomalies. *The Journal of Finance, 51*(1), 55-84. doi:https://doi.org/10.1111/j.1540-6261.1996.tb05202.x

Fama, E. F., & French, K. R. (2008). Dissecting Anomalies. *The Journal of Finance, 63*(4), 1653-1678. doi:10.1111/j.1540-6261.2008.01371.x

Figlewski, S. (1979). Subjective information and market efficiency in a betting market. *Journal of Political Economy, 87*(1), 75-88. doi:https://doi.org/10.1086/260740

The financial crisis and the role of federal regulators, (2008).

Fodor, A., DiFilippo, M., Krieger, K., & Davis, J. (2013). Inefficient pricing from holdover bias in NFL point spread markets. *Applied Financial Economics, 23*(17), 1407-1418. doi:https://doi.org/10.1080/09603107.2013.829201

Forrest, D., Goddard, J., & Simmons, R. (2005). Odds-setters as forecasters: The case of English football. *International Journal of Forecasting, 21*(3), 551-564. doi:https://doi.org/10.1016/j.ijforecast.2005.03.003

Forrest, D., & Simmons, R. (2008). Sentiment in the betting market on Spanish football. *Applied Economics, 40*, 119–126. doi:https://doi.org/10.1080/00036840701522895

Forsythe, R., Nelson, F., Neumann, G. R., & Wright, J. (1992). Anatomy of an experimental political stock market. *The American Economic Review*, 1142-1161.

Forsythe, R., Palfrey, T. R., & Plott, C. R. (1982). Asset valuation in an experimental market. *Econometrica: Journal of the Econometric Society, 50*(3), 537-567. doi:https://doi.org/10.2307/1912600

Franck, E., Verbeek, E., & Nüesch, S. (2011). Sentimental Preferences and the Organizational Regime of Betting Markets. *Southern Economic Journal, 78*(2), 502–518. doi:https://doi.org/10.4284/0038-4038-78.2.502

Franck, E., Verbeek, E., & Nüesch, S. (2013). Inter-market Arbitrage in Betting. *Economica, 80*(318), 300-325. doi:https://doi.org/10.1111/ecca.12009

Frazzini, A. (2006). The disposition effect and underreaction to news. *The Journal of Finance, 61*(4), 2017-2046. doi:https://doi.org/10.1111/j.1540-6261.2006.00896.x

Gabriel, P. E., & Marsden, J. R. (1990). An examination of market efficiency in British racetrack betting. *Journal of Political Economy, 98*(4), 874-885. doi:https://doi.org/10.1086/261710

Gadanecz, B., Moessner, R., & Upper, C. (2007). Economic derivatives. *BIS Quarterly Review, March*.

Galariotis, E. (2013). Mesdames et Messieurs, momentum performance is not so abnormal after all! *Applied Economics, 45*(27), 3871-3879. doi:https://doi.org/10.1080/00036846.2012.730138

Gambling Commission. (2018). *Gambling participation in 2017: behaviour, awareness and attitudes*. Retrieved from https://www.gamblingcommission.gov.uk/PDF/survey-data/Gambling-participation-in-2017-behaviour-awareness-and-attitudes.pdf

Gandar, J. M., Dare, W. H., Brown, C. R., & Zuber, R. A. (1998). Informed traders and price variations in the betting market for professional basketball games. *The Journal of Finance, 53*(1), 385-401.

Gandar, J. M., Zuber, R. A., & Lamb, R. P. (2001). The home field advantage revisited: a search for the bias in other sports betting markets. *Journal of Economics and Business, 53*(4), 439-453.

Gandar, J. M., Zuber, R. A., O'Brien, T., & Russo, B. (1988). Testing Rationality in the Point Spread Betting Market. *The Journal of Finance, 43*(4), 995-1008. doi:https://doi.org/10.2307/2328148

Geczy, C. C., & Samonov, M. (2015). 215 Years of Global Multi-Asset Momentum: 1800-2014 (Equities, Sectors, Currencies, Bonds, Commodities and Stocks). doi:https://doi.org/10.2139/ssrn.2607730

Georgopoulou, A., & Wang, J. (2016). The Trend Is Your Friend: Time-Series Momentum Strategies across Equity and Commodity Markets*. *Review of Finance, 21*(4), 1557-1592. doi:10.1093/rof/rfw048

Geurkink, Y., Boone, J., Verstockt, S., & Bourgois, J. G. (2021). Machine Learning-Based Identification of the Strongest Predictive Variables of Winning and Losing in Belgian Professional Soccer. *Applied Sciences, 11*(5), 2378. doi:https://doi.org/10.3390/app11052378

Gillen, B. J., Plott, C. R., & Shum, M. (2013). *Inside intel: Sales forecasting using an information aggregation mechanism*. Retrieved from https://www.eco.uc3m.es/temp/IAMField_v2matt.pdf

Gillette, A. B., Stevens, D. E., Watts, S. G., & Williams, A. W. (1999). Price and volume reactions to public information releases: An experimental approach incorporating traders' subjective beliefs. *Contemporary Accounting Research, 16*(3), 437-479. doi:https://doi.org/10.1111/j.1911-3846.1999.tb00590.x

Goel, S., Reeves, D. M., Watts, D. J., & Pennock, D. M. (2010). *Prediction without markets.* Paper presented at the Proceedings of the 11th ACM conference on Electronic commerce.

Goetzmann, W. N., & Huang, S. (2018). Momentum in Imperial Russia. *Journal of Financial Economics, 130*(3), 579-591. doi:https://doi.org/10.1016/j.jfineco.2018.07.008

Goldfarb, A., & Tucker, C. (2019). Digital Economics. *Journal of Economic Literature, 57*(1), 3-43. doi:10.1257/jel.20171452

Golec, J., & Tamarkin, M. (1991). The degree of inefficiency in the football betting market: Statistical tests. *Journal of Financial Economics, 30*(2), 311-323. doi:https://doi.org/10.1016/0304-405X(91)90034-H

Golec, J., & Tamarkin, M. (1998). Bettors love skewness, not risk, at the horse track. *Journal of Political Economy, 106*(1), 205-225. doi:https://doi.org/10.1086/250007

Gorodnichenko, Y., Sheremirov, V., & Talavera, O. (2018). Price setting in online markets: Does IT click? *Journal of the European Economic Association, 16*(6), 1764-1811. doi:https://doi.org/10.1093/jeea/jvx050

Gorodnichenko, Y., & Talavera, O. (2017). Price setting in online markets: Basic facts, international comparisons, and cross-border integration. *American Economic Review, 107*(1), 249-282. doi:https://doi.org/10.1257/aer.20141127

Gramm, M., McKinney, C. N., & Owens, D. H. (2012). Efficiency and arbitrage across parimutuel wagering pools. *Applied Economics, 44*(14), 1813-1822. doi:https://doi.org/10.1080/00036846.2011.554378

Grant, A., Oikonomidis, A., Bruce, A. C., & Johnson, J. E. (2018). New entry, strategic diversity and efficiency in soccer betting markets: the creation and suppression of arbitrage opportunities. *The European Journal of Finance, 24*(18), 1799-1816. doi:https://doi.org/10.1080/1351847X.2018.1443148

Gray, P. K., & Gray, S. F. (1997). Testing Market Efficiency: Evidence from the NFL Sports Betting Market. *The Journal of Finance, 52*(4), 1725-1737. doi:https://doi.org/10.2307/2329455

Griffith, R. M. (1949). Odds adjustments by American horse-race bettors. *The American Journal of Psychology, 62*(2), 290-294. doi:https://doi.org/10.2307/1418469

Grossman, S. J., & Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *The American Economic Review, 70*(3), 393-408. doi:https://doi.org/10.7916/D8765R99

Gruen, A. (1976). An inquiry into the economics of race-track gambling. *Journal of Political Economy, 84*(1), 169-177. doi:https://doi.org/10.1086/260419

Gürkaynak, R. S., Wolfers, J., Carroll, C. D., & Szeidl, A. (2005). *Macroeconomic Derivatives: An Initial Analysis of Market-Based Macro Forecasts, Uncertainty, and Risk [with Comments].* Paper presented at the NBER International Seminar on Macroeconomics.

GVC Holdings. (2019). *Annual Report 2018*. Retrieved from https://gvc-plc.com/wp-content/uploads/2019/05/GVC_Annual-Report-2018-Final.pdf

H2 Gambling Capital. (2019).

Hahn, R. W., & Tetlock, P. C. (2005). Using information markets to improve public decision making. *Harv. JL & Pub. Pol'y, 29*, 213.

Hankins, R., & Lee, A. (2011). *Crowd sourcing and prediction markets.* Paper presented at the CHI.

Hansen, J., Schmidt, C., & Strobel, M. (2004). Manipulation in political stock markets–preconditions and evidence. *Applied Economics Letters, 11*(7), 459-463. doi:https://doi.org/10.1080/1350485042000191700

Hanson, R. (1999). Decision markets. Retrieved from http://mason.gmu.edu/~rhanson/decisionmarkets.pdf

Hanson, R. (2005). The informed press favored the policy analysis market. *George Mason University, 15*. Retrieved from http://mason.gmu.edu/~rhanson/PAMpress.pdf

Hanson, R. (2013). Shall we vote on values, but bet on beliefs? *Journal of Political Philosophy, 21*(2), 151-178. doi:https://doi.org/10.1111/jopp.12008

Hanson, R., & Oprea, R. (2009). A manipulator can aid prediction market accuracy. *Economica, 76*(302), 304-314. doi:https://doi.org/10.1111/j.1468-0335.2008.00734.x

Harris, M. (2016). A History of BetVictor. Retrieved from https://www.ukbookmakers.info/a-history-of-betvictor

Haruvy, E., Lahav, Y., & Noussair, C. N. (2007). Traders' expectations in asset markets: experimental evidence. *American Economic Review, 97*(5), 1901-1920. doi:https://doi.org/10.1257/aer.97.5.1901

Harvey, C. R. (2017). Presidential address: The scientific outlook in financial economics. *The Journal of Finance, 72*(4), 1399-1440. doi:https://doi.org/10.1111/jofi.12530

Harvey, C. R., Liu, Y., & Zhu, H. (2016). … and the cross-section of expected returns. *The Review of Financial Studies, 29*(1), 5-68. doi:https://doi.org/10.1093/rfs/hhv059

Harvey, C. R., & Siddique, A. (2000). Conditional skewness in asset pricing tests. *The Journal of Finance, 55*(3), 1263-1295.

Hausch, D. B., & Ziemba, W. T. (1990a). Arbitrage Strategies for Cross-Track Betting on Major Horse Races. *The Journal of Business, 63*(1), 61-78. doi:https://doi.org/10.1142/9789812819192_0040

Hausch, D. B., & Ziemba, W. T. (1990b). Locks at the racetrack. *Interfaces, 20*(3), 41-48. doi:https://doi.org/10.1287/inte.20.3.41

Havranek, T., Irsova, Z., & Zeynalova, O. (2018). Tuition fees and university enrolment: a meta-regression analysis. *Oxford Bulletin of Economics and Statistics, 80*(6), 1145-1184. doi:https://doi.org/10.1111/obes.12240

Haw Allensworth, R. (2009). Prediction markets and law: A skeptical account. *Harvard Law Review, 122*, 1217.

Hayek, F. A. (1945). The Use of Knowledge in Society. *The American Economic Review, 35*(4), 519-530. Retrieved from www.jstor.org/stable/1809376

Hayek, F. A. (1960). *The constitution of liberty*. London: Routledge & Kegan Paul.

Hayek, F. A., & Caldwell, B. (2014). *The road to serfdom: Text and documents: The definitive edition*: Routledge.

Hickman, D. C. (2020). Efficiency in the madness? examining the betting market for the ncaa men’s basketball tournament. *Journal of Economics and Finance, 44*(3), 611-626. doi:https://doi.org/10.1007/s12197-020-09507-7

Hirshleifer, D. (2015). Behavioral finance. *Annual Review of Financial Economics, 7*, 133-159. doi:https://doi.org/10.1146/annurev-financial-092214-043752

Hirshleifer, D. (2020). Presidential address: Social transmission bias in economics and finance. *The Journal of Finance, 75*(4), 1779-1831. doi:https://doi.org/10.1111/jofi.12906

Hirshleifer, D., Lim, S. S., & Teoh, S. H. (2009). Driven to distraction: Extraneous events and underreaction to earnings news. *The Journal of Finance, 64*(5), 2289-2325. doi:https://doi.org/10.1111/j.1540-6261.2009.01501.x

Hirshleifer, D., & Shumway, T. (2003). Good day sunshine: Stock returns and the weather. *The Journal of Finance, 58*(3), 1009-1032. doi:https://doi.org/10.1111/1540-6261.00556

Hoerl, A. E., & Fallin, H. K. (1974). Reliability of subjective evaluations in a high incentive situation. *Journal of the Royal Statistical Society: Series A (General), 137*(2), 227-230. doi:https://doi.org/2344550

Hong Kong Jockey Club. (2020). Betting Guide. Retrieved from https://special.hkjc.com/racing/info/en/betting/guide_qualifications_pari.asp

Horvat, T., & Job, J. (2020). The use of machine learning in sport outcome prediction: A review. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 10*(5), e1380. doi:https://doi.org/10.1002/widm.1380

Huang, M.-L., & Li, Y.-Z. (2021). Use of Machine Learning and Deep Learning to Predict the Outcomes of Major League Baseball Matches. *Applied Sciences, 11*(10), 4499. doi:https://doi.org/10.3390/app11104499

Hubáček, O., Šourek, G., & Železný, F. (2019). Exploiting sports-betting market using machine learning. *International Journal of Forecasting, 35*(2), 783-796. doi:https://doi.org/10.1016/j.ijforecast.2019.01.001

Huberman, G., & Regev, T. (2001). Contagious Speculation and a Cure for Cancer: A Nonevent that Made Stock Prices Soar. *The Journal of Finance, 56*(1), 387-396. doi:10.1111/0022-1082.00330

Hui, K. W., & Yeung, P. E. (2013). Underreaction to industry-wide earnings and the post-forecast revision drift. *Journal of Accounting research, 51*(4), 701-737. doi:https://doi.org/10.1111/1475-679X.12006

Humphreys, B. R., Paul, R. J., & Weinbach, A. (2014). Understanding price movements in point-spread betting markets: evidence from NCAA basketball. *Eastern Economic Journal, 40*(4), 518-534. doi:https://doi.org/10.1057/eej.2013.10

Humphreys, B. R., Paul, R. J., & Weinbach, A. P. (2013). Consumption benefits and gambling: Evidence from the NCAA basketball betting market. *Journal of Economic Psychology, 39*, 376-386.

Hurley, W., & McDonough, L. (1995). A note on the Hayek hypothesis and the favorite-longshot bias in parimutuel betting. *The American Economic Review*, 949-955.

Interwetten. (n.d.). Über Interwetten. Retrieved from http://www.interwetten.org/en/company/

Jacobs, H., & Müller, S. (2020). Anomalies across the globe: Once public, no longer existent? *Journal of Financial Economics, 135*(1), 213-230. doi:https://doi.org/10.1016/j.jfineco.2019.06.004

Jaffe, J. F., & Winkler, R. L. (1976). Optimal speculation against an efficient market. *The Journal of Finance, 31*(1), 49-61. doi:https://doi.org/10.1111/j.1540-6261.1976.tb03195.x

Jamieson, J. P. (2010). The home field advantage in athletics: A meta-analysis. *Journal of Applied Social Psychology, 40*(7), 1819-1848. doi:https://doi.org/10.1111/j.1559-1816.2010.00641.x

Jegadeesh, N., & Titman, S. (1993). Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. *The Journal of Finance, 48*(1), 65-91. doi:10.1111/j.1540-6261.1993.tb04702.x

Jegadeesh, N., & Titman, S. (2001). Profitability of momentum strategies: An evaluation of alternative explanations. *The Journal of Finance, 56*(2), 699-720. doi:http://doi.org/10.2139/ssrn.166840

Jehue, R., Street, D., & Huizenga, R. (1993). Effect of time zone and game time changes on team performance: National Football League. *Medicine and science in sports and exercise, 25*(1), 127-131. doi:https://doi.org/10.1249/00005768-199301000-00017

Jensen, M. C. (1978). Some anomalous evidence regarding market efficiency. *Journal of Financial Economics, 6*(2), 95-101. doi:https://doi.org/10.1016/0304-405X(78)90025-9

Johnson, T. C. (2002). Rational momentum effects. *The Journal of Finance, 57*(2), 585-608. doi:https://doi.org/10.1111/1540-6261.00435

Jostova, G., Nikolova, S., Philipov, A., & Stahel, C. W. (2013). Momentum in Corporate Bond Returns. *The Review of Financial Studies, 26*(7), 1649-1693. doi:10.1093/rfs/hht022

Kaunitz, L., Zhong, S., & Kreiner, J. (2017). Beating the bookies with their own numbers-and how the online sports betting market is rigged. *arXiv preprint arXiv:1710.02824*.

Keynes, J. M. (1936). *The general theory of employment, interest and money*: London : Macmillan, 1936.

Kirchler, M. (2009). Underreaction to fundamental information and asymmetry in mispricing between bullish and bearish markets. An experimental study. *Journal of Economic Dynamics and Control, 33*(2), 491-506. doi:https://doi.org/10.1016/j.jedc.2008.08.002

Kirchler, M., Huber, J., & Stöckl, T. (2012). Thar she bursts: Reducing confusion reduces bubbles. *American economic review, 102*(2), 865-883. doi:https://doi.org/10.1257/aer.102.2.865

Knight, F. H. (1921). Risk, uncertainty and profit New York. *NY: Harper & Row*.

Kochman, L., & Goodwin, R. (2004). Underdogs are Man's Best Friend: A Test of Football Market Efficiency. *Journal of Sports Economics, 5*(4), 387-391. doi:https://doi.org/10.1177/1527002504264426

Kochman, L., Goodwin, R., & Gilliam, K. (2017). You Win Some and You Lose Some: A Test of the Football Betting Market. *Southern Business & Economic Journal, 40*(1), 56-60.

Kuester, D. D., & Sanders, S. (2011). Regional information and market efficiency: the case of spread betting in United States college football. *Journal of Economics and Finance, 35*(1), 116-122. doi:https://doi.org/10.1007/s12197-009-9113-3

Labour Party. (1942). *The Old World and the New Society: A Report on the Problems of War and Peace Reconstruction*: Labour Party.

Lacey, N. J. (1990). An estimation of market efficiency in the NFL point spread betting market. *Applied Economics, 22*(1), 117-129. doi:https://doi.org/10.1080/00036849000000056

Lawson, R. A. (2019). Capitalism is (Slowly) Winning Around the World. Retrieved from https://www.bushcenter.org/catalyst/capitalism/lawson-capitalism-is-winning-around-the-world.html

LeRoy, S. F. (1989). Efficient Capital Markets and Martingales. *Journal of Economic Literature, 27*(4), 1583-1621. Retrieved from www.jstor.org/stable/2727024

Lesmond, D. A., Schill, M. J., & Zhou, C. (2004). The illusory nature of momentum profits. *Journal of Financial Economics, 71*(2), 349-380. doi:https://doi.org/10.1016/S0304-405X(03)00206-X

Levich, R. M. (1989). Is the foreign exchange market efficient? *Oxford review of economic policy, 5*(3), 40-60.

Levitt, S. D. (2004). Why Are Gambling Markets Organised so Differently from Financial Markets? *The Economic Journal, 114*(495), 223-246. doi:https://doi.org/10.1111/j.1468-0297.2004.00207.x

Levitt, S. D., & List, J. A. (2007). What Do Laboratory Experiments Measuring Social Preferences Reveal About the Real World? *Journal of Economic Perspectives, 21*(2), 153-174. doi:https://doi.org/10.1257/jep.21.2.153

Lewellen, J. (2002). Momentum and Autocorrelation in Stock Returns. *The Review of Financial Studies, 15*(2), 533-564. doi:https://doi.org/10.1093/rfs/15.2.533

Li, J. (2018). Explaining momentum and value simultaneously. *Management Science, 64*(9), 4239-4260. doi:https://doi.org/10.1287/mnsc.2017.2735

Light, R. (2007). The Gambling Act 2005: Regulatory Containment and Market Control. *The Modern Law Review, 70*(4), 626-653. doi:10.1111/j.1468-2230.2007.00655.x

Lin, S., & Rassenti, S. (2012). Are under- and over-reaction the same matter? Experimental evidence. *Journal of Economic Behavior & Organization, 84*(1), 39-61. doi:https://doi.org/10.1016/j.jebo.2012.07.004

Litan, R. E., & Rivlin, A. M. (2001). Projecting the Economic Impact of the Internet. *The American Economic Review, 91*(2), 313-317. doi:https://doi.org/10.1257/aer.91.2.313

Lo, A. W. (2019). Adaptive markets: Financial Evolution at the Speed of Thought. In: Princeton University Press.

Lo, A. W., & MacKinlay, A. C. (1990). Data-Snooping Biases in Tests of Financial Asset Pricing Models. *The Review of Financial Studies, 3*(3), 431-467. doi:https://doi.org/10.1093/rfs/3.3.431

Lünnemann, P., & Wintr, L. (2011). Price stickiness in the US and Europe revisited: evidence from internet prices. *Oxford bulletin of economics and statistics, 73*(5), 593-621. doi:https://doi.org/10.1111/j.1468-0084.2011.00657.x

Malkiel, B. G. (2003). The Efficient Market Hypothesis and Its Critics. *The Journal of Economic Perspectives, 17*(1), 59-82. doi:https://doi.org/10.1257/089533003321164958

Maloney, M. T., & Mulherin, J. H. (2003). The complexity of price discovery in an efficient market: the stock market reaction to the Challenger crash. *Journal of Corporate Finance, 9*(4), 453-479. doi:https://doi.org/10.1016/S0929-1199(02)00055-X

Mandelbrot, B. (1966). Forecasts of future prices, unbiased markets, and" martingale" models. *The Journal of Business, 39*(1), 242-255. doi:https://doi.org/10.1086/294850

Mangold, B., Dooley, M., Flake, G. W., Hoffman, H., Kasturi, T., Pennock, D. M., & Dornfest, R. (2005). The tech buzz game [stock market prediction]. *Computer, 38*(7), 94-97. doi:https://doi.org/10.1109/MC.2005.243

Manski, C. F. (2006). Interpreting the predictions of prediction markets. *Economics Letters, 91*(3), 425-429. doi:https://doi.org/10.1016/j.econlet.2006.01.004

Marshall, B. R. (2009). How quickly is temporary market inefficiency removed? *The Quarterly Review of Economics and Finance, 49*(3), 917-930. doi:https://doi.org/10.1016/j.qref.2009.04.006

Maymin, P. Z. (2019). Wage against the machine: A generalized deep-learning market test of dataset value. *International Journal of Forecasting, 35*(2), 776-782. doi:https://doi.org/10.1016/j.ijforecast.2017.09.008

McLean, R. D., & Pontiff, J. (2016). Does academic research destroy stock return predictability? *The Journal of Finance, 71*(1), 5-32. doi:https://doi.org/10.1111/jofi.12365

Milgrom, P., & Stokey, N. (1982). Information, trade and common knowledge. *Journal of Economic Theory, 26*(1), 17-27. doi:https://doi.org/10.1016/0022-0531(82)90046-1

Mills, B. M., & Salaga, S. (2018). A natural experiment for efficient markets: Information quality and influential agents. *Journal of Financial Markets, 40*(C), 23-39. doi:https://doi.org/10.1016/j.finmar.2018.07.002

Moore, E., & Francisco, J. (2019). The SEC versus the Dow Jones: a profitable betting strategy in NCAA football. *Journal of Prediction Markets, 13*(2), 3-14. doi:https://doi.org/10.5750/jpm.v13i2.1740

Moore, L., & Juh, S. (2006). Derivative Pricing 60 Years before Black–Scholes: Evidence from the Johannesburg Stock Exchange. *The Journal of Finance, 61*(6), 3069-3098. doi:10.1111/j.1540-6261.2006.01012.x

Morgan, J., Orzen, H., & Sefton, M. (2006). An experimental study of price dispersion. *Games and Economic Behavior, 54*(1), 134-158. doi:https://doi.org/10.1016/j.geb.2004.07.005

Morgan, W. D., & Vasche, J. D. (1979). Horseracing demand, parimutuel taxation and state revenue potential. *National Tax Journal*, 185-194. doi:https://doi.org/10.1086/NTJ41862220

Moskowitz, T. J. (2021). Asset Pricing and Sports Betting. *The Journal of Finance, 76*(6), 3153-3209. doi:https://doi.org/10.1111/jofi.13082

Moskowitz, T. J., Ooi, Y. H., & Pedersen, L. H. (2012). Time series momentum. *Journal of Financial Economics, 104*(2), 228-250. doi:https://doi.org/10.1016/j.jfineco.2011.11.003

Moskowitz, T. J., & Vasudevan, K. (2021). What Can Betting Markets Tell Us About Investor Preferences and Beliefs? Implications for Low Risk Anomalies. *Implications for Low Risk Anomalies (May 13, 2021)*. doi:http://dx.doi.org/10.2139/ssrn.3845505

New York Times. (1924). ELECTION BETTING SLOW IN WALL ST.; Plenty of Coolidge Money at 3 to 1, With Few Takers, Commissioners Report. Retrieved from https://www.nytimes.com/1924/09/28/archives/election-betting-slow-in-wall-st-plenty-of-coolidge-money-at-3-to-1.html?searchResultPosition=1

Newall, P. W., Thobhani, A., Walasek, L., & Meyer, C. (2019). Live-odds gambling advertising and consumer protection. *PloS one, 14*(6), e0216876. doi:https://doi.org/10.1371/journal.pone.0216876

Noussair, C. N., & Tucker, S. (2013). Experimental research on asset pricing. *Journal of Economic Surveys, 27*(3), 554-569. doi:https://doi.org/10.1111/joes.12019

Oorlog, D. R. (1995). Serial Correlation in the Wagering Market for Professional Basketball. *Quarterly Journal of Business and Economics, 34*(2), 96-109. Retrieved from www.jstor.org/stable/40473149

Orlov, E. (2011). How does the internet influence price dispersion? Evidence from the airline industry. *The Journal of Industrial Economics, 59*(1), 21-37. doi:10.1111/j.1467-6451.2011.00444.x

Ortner, G. (1998). *Forecasting markets—An industrial application*. Retrieved from http://ebweb.at/apsm/fmaia2.pdf

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

Page, L., & Siemroth, C. (2021). How Much Information Is Incorporated into Financial Asset Prices? Experimental Evidence. *The Review of Financial Studies, 34*(9), 4412-4449. doi:https://doi.org/10.1093/rfs/hhaa143

Palan, S. (2013). A review of bubbles and crashes in experimental asset markets. *Journal of Economic Surveys, 27*(3), 570-588. doi: https://doi.org/10.1111/joes.12023

Pankoff, L. D. (1968). Market Efficiency and Football Betting. *The Journal of Business, 41*(2), 203-214. doi:http://dx.doi.org/10.1086/295077

Paton, D., Siegel, D. S., & Williams, L. V. (2002). A Policy Response to the E-Commerce Revolution: The Case of Betting Taxation in the UK. *The Economic Journal, 112*(480), F296-F314. doi:https://doi.org/10.1111/1468-0297.00045

Paton, D., & Williams, L. V. (1998). Do betting costs explain betting biases? *Applied Economics Letters, 5*(5), 333-335. doi:https://doi.org/10.1080/758524413

Patton, A. J., & Weller, B. M. (2020). What you see is not what you get: The costs of trading market anomalies. *Journal of Financial Economics, 137*(2), 515-549. doi:https://doi.org/10.1016/j.jfineco.2020.02.012

Paul, R. J., & Weinbach, A. P. (2002). Market efficiency and a profitable betting rule: Evidence from totals on professional football. *Journal of Sports Economics, 3*(3), 256-263. doi:https://doi.org/10.1177/1527002502003003003

Paul, R. J., & Weinbach, A. P. (2005a). Bettor Misperceptions in the NBA: The Overbetting of Large Favorites and the “Hot Hand”. *Journal of Sports Economics, 6*(4), 390-400. doi:https://doi.org/10.1177/1527002504266861

Paul, R. J., & Weinbach, A. P. (2005b). Market efficiency and NCAA college basketball gambling. *Journal of Economics and Finance, 29*(3), 403-408. doi:https://doi.org/10.1007/BF02761584

Paul, R. J., & Weinbach, A. P. (2011). NFL bettor biases and price setting: further tests of the Levitt hypothesis of sportsbook behaviour. *Applied Economics Letters, 18*(2), 193-197. doi:https://doi.org/10.1080/13504850903508242

Paul, R. J., Weinbach, A. P., & Humphreys, B. (2011). Revisiting the "hot hand" hypothesis in the NBA betting market using actual sportsbook betting percentages on favorites and underdogs. *Journal of Gambling Business & Economics, 5*(2), 42-56. doi:https://doi.org/10.5750/jgbe.v5i2.569

Paul, R. J., Weinbach, A. P., & Humphreys, B. (2014). Bettor Belief in the “Hot Hand”: Evidence From Detailed Betting Data on the NFL. *Journal of Sports Economics, 15*(6), 636-649. doi:https://doi.org/10.1177/1527002512465414

Paul, R. J., Weinbach, A. P., & Weinbach, J. (2003). Fair bets and profitability in college football gambling. *Journal of Economics and Finance, 27*(2), 236-242. doi:https://doi.org/10.1007/BF02827221

Paul, R. J., Weinbach, A. P., & Wilson, M. (2004). Efficient markets, fair bets, and profitability in NBA totals 1995–96 to 2001–02. *The Quarterly Review of Economics and Finance, 44*(4), 624-632. doi:https://doi.org/10.1016/S1062-9769(03)00033-4

Paul, R. J., Weinbach, A. P., & Wilson, M. (2014). Bettor habits when point spreads and money lines are offered on the same game: the NFL. *Journal of Prediction Markets, 8*(3), 57-74. doi:https://doi.org/10.5750/jpm.v8i3.976

Pennacchi, G. G. (2008). *Theory of asset pricing*: Pearson/Addison-Wesley Boston.

Pennerstorfer, D., Schmidt-Dengler, P., Schutz, N., Weiss, C., & Yontcheva, B. (2020). Information and price dispersion: Theory and evidence. *International Economic Review, 61*(2), 871-899. doi:https://doi.org/10.1111/iere.12443

Plott, C. R., & Chen, K.-Y. (2002). Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. *Social Science Working Paper California Institute of Technology, 1131*. Retrieved from https://resolver.caltech.edu/CaltechAUTHORS:20140317-135547085

Polgreen, P. M., Nelson, F. D., Neumann, G. R., & Weinstein, R. A. (2007). Use of prediction markets to forecast infectious disease activity. *Clinical Infectious Diseases, 44*(2), 272-279. doi:https://doi.org/10.1086/510427

Pope, P. F., & Peel, D. A. (1989). Information, prices and efficiency in a fixed-odds betting market. *Economica*, 323-341. doi:https://doi.org/2554281

Preston, M. G., & Baratta, P. (1948). An experimental study of the auction-value of an uncertain outcome. *The American Journal of Psychology, 61*(2), 183-193. doi:https://doi.org/10.2307/1416964

Reagan, R. (1986). The President's News Conference August 12, 1986. In.

Rhode, P. W., & Strumpf, K. S. (2004). Historical presidential betting markets. *Journal of Economic Perspectives, 18*(2), 127-141. doi:https://doi.org/10.1257/0895330041371277

Robbins, L. (1937). *Economic planning and international order*: Macmillan, London.

Roll, R. (1984). Orange Juice and Weather. *American Economic Review, 74*(5), 861-880. Retrieved from https://EconPapers.repec.org/RePEc:aea:aecrev:v:74:y:1984:i:5:p:861-80

Rosenbloom, E. S., & Notz, W. (2006). Statistical Tests of Real-Money versus Play-Money Prediction Markets. *Electronic Markets, 16*(1), 63-69. doi:10.1080/10196780500491303

Salop, S., & Stiglitz, J. (1977). Bargains and ripoffs: A model of monopolistically competitive price dispersion. *The Review of Economic Studies, 44*(3), 493-510. doi:https://doi.org/2296903

Sandford, J., & Shea, P. (2013). Optimal setting of point spreads. *Economica, 80*(317), 149-170. doi:https://doi.org/10.1111/j.1468-0335.2012.00939.x

Sauer, R. D. (1998). The Economics of Wagering Markets. *Journal of Economic Literature, 36*(4), 2021-2064. Retrieved from http://www.jstor.org/stable/2565046

Sauer, R. D. (2005). The state of research on markets for sports betting and suggested future directions. *Journal of Economics and Finance, 29*(3), 416-426. doi:https://doi.org/10.1007/BF02761586

Sauer, R. D., Brajer, V., Ferris, S. P., & Marr, M. W. (1988). Hold your bets: Another look at the efficiency of the gambling market for National Football League games. *Journal of Political Economy, 96*(1), 206-213. doi:https://doi.org/10.1086/261532

Schnyzer, A., & Hizgilov, A. (2018). The impact of "jet lag" on the AFL point spread wagering market. *Journal of Gambling Business & Economics, 12*(1), 3-15. doi:https://doi.org/10.5750/jgbe.v12i1.1518

Servan-Schreiber, E., Wolfers, J., Pennock, D. M., & Galebach, B. (2004). Prediction markets: Does money matter? *Electronic Markets, 14*(3), 243-251. doi:https://doi.org/10.1080/1019678042000245254

Shank, C. A. (2018). Is the NFL betting market still inefficient? *Journal of Economics and Finance, 42*(4), 818-827. doi:https://doi.org/10.1007/s12197-018-9431-4

Shank, C. A. (2019). NFL betting market efficiency, divisional rivals, and profitable strategies. *Studies in Economics and Finance, 36*(3), 567-580. doi:https://doi.org/10.1108/SEF-11-2018-0354

Shiller, R. J. (1981). Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends? *The American Economic Review, 71*(3), 421-436. Retrieved from http://www.jstor.org/stable/1802789

Shin, H. S. (1992). Prices of state contingent claims with insider traders, and the favourite-longshot bias. *The Economic Journal, 102*(411), 426-435. doi:https://doi.org/2234526

Simon, H. A. (1971). Designing Organizations For An Information-Rich World. In M. Greenberger (Ed.), *Computers, communications, and the public interest* (pp. 37-52). Baltimore: The Johns Hopkins Press.

Sinkey, M., & Logan, T. (2014). Does the hot hand drive the market? Evidence from college football betting markets. *Eastern Economic Journal, 40*(4), 583-603. doi:https://doi.org/10.1057/eej.2013.33

Smith, A. (1776). *An Inquiry into the Nature and Causes of the Wealth of Nations*: McMaster University Archive for the History of Economic Thought.

Smith, M. A., Paton, D., & Williams, L. V. (2006). Market efficiency in person-to-person betting. *Economica, 73*(292), 673-689. doi:https://doi.org/10.1111/j.1468-0335.2006.00518.x

Smith, V. L. (1971). Economic theory of wager markets. *Economic Inquiry, 9*(3), 242-255. doi:https://doi.org/10.1111/j.1465-7295.1971.tb01639.x

Smith, V. L. (1982a). Markets as economizers of information: experimental examination of the "Hayek hypothesis". *Economic Inquiry, 20*(2), 165-179. doi:https://doi.org/10.1111/j.1465-7295.1982.tb01149.x

Smith, V. L. (1982b). Markets as economizers of information: experimental examination of the "Hayek hypothesis". *Economic Inquiry, 20*(2), 165-179. doi:https://doi.org/10.1111/j.1465-7295.1982.tb01149.x

Smith, V. L., Suchanek, G. L., & Williams, A. W. (1988). Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets. *Econometrica, 56*(5), 1119-1151. doi:10.2307/1911361

Snowberg, E., & Wolfers, J. (2010). Explaining the favorite–long shot bias: Is it risk-love or misperceptions? *Journal of Political Economy, 118*(4), 723-746. doi:https://doi.org/10.1086/655844

Snowberg, E., Wolfers, J., & Zitzewitz, E. (2005). Information (in) efficiency in prediction markets. *Information efficiency in financial and betting markets*, 366-386. doi:https://doi.org/10.1017/CBO9780511493614.019

Snowberg, E., Wolfers, J., & Zitzewitz, E. (2013). Prediction markets for economic forecasting. In *Handbook of Economic Forecasting* (Vol. 2, pp. 657-687): Elsevier.

Snyder, W. W. (1978). Horse racing: Testing the efficient markets model. *The Journal of Finance, 33*(4), 1109-1118. doi:https://doi.org/10.1111/j.1540-6261.1978.tb02051.x

Sobel, R. S., & Raines, S. T. (2003). An examination of the empirical derivatives of the favourite-longshot bias in racetrack betting. *Applied Economics, 35*(4), 371-385. doi:https://doi.org/10.1080/00036840110111176

Song, C., Boulier, B. L., & Stekler, H. O. (2007). The comparative accuracy of judgmental and model forecasts of American football games. *International Journal of Forecasting, 23*(3), 405-413. doi:https://doi.org/10.1016/j.ijforecast.2007.05.003

Sowell, T. (2017). Thomas Sowell: 1% shouldn't be dictating to the other 99%. Retrieved from https://www.deseret.com/2007/5/17/20019066/thomas-sowell-1-shouldn-t-be-dictating-to-the-other-99

Stamatopoulos, I., Bassamboo, A., & Moreno, A. (2021). The effects of menu costs on retail performance: Evidence from adoption of the electronic shelf label technology. *Management Science, 67*(1), 242-256. doi:https://doi.org/10.1287/mnsc.2019.3551

Stanley, T. D., & Doucouliagos, H. (2010). Picture this: a simple graph that reveals much ado about research. *Journal of Economic Surveys, 24*(1), 170-191. doi:https://doi.org/10.1111/j.1467-6419.2009.00593.x

Statista. (2019). Revenue of the biggest (Big Five*) European soccer leagues from 1996/97 to 2019/20 (in million euros). Retrieved from https://www.statista.com/statistics/261218/big-five-european-soccer-leagues-revenue/

Stevens, D. E., & Williams, A. W. (2004). Inefficiency in earnings forecasts: Experimental evidence of reactions to positive vs. negative information. *Experimental Economics, 7*(1), 75-92. doi:http://doi.org/10.2139/ssrn.388420

Stigler, G. J. (1961). The economics of information. *Journal of Political Economy, 69*(3), 213-225. doi:http://doi.org/10.1086/258464

Stiglitz, J. (2002). There is no invisible hand. *The Guardian*. Retrieved from https://www.theguardian.com/education/2002/dec/20/highereducation.uk1

Stiglitz, J. (2010). Needed: a new economic paradigm. *Financial Times*. Retrieved from https://www.ft.com/content/d5108f90-abc2-11df-9f02-00144feabdc0

Stöckl, T., Huber, J., & Kirchler, M. (2015). Multi-period experimental asset markets with distinct fundamental value regimes. *Experimental Economics, 18*(2), 314-334. doi:https://doi.org/10.1007/s10683-014-9404-1

Strumpf, K. (2003). Illegal sports bookmakers. *Unpublished manuscript. Available at http://rgco. org/articles/illegal_sports_bookmakers. pdf*.

Stübinger, J., Mangold, B., & Knoll, J. (2020). Machine learning in football betting: Prediction of match results based on player characteristics. *Applied Sciences, 10*(1), 46. doi:https://doi.org/10.3390/app10010046

Suits, D. B. (1979). The elasticity of demand for gambling. *The Quarterly Journal of Economics*, 155-162. doi:https://doi.org/1882605

Summers, L. H. (1985). On economics and finance. *The Journal of Finance, 40*(3), 633-635. doi:https://doi.org/2327785

Summers, L. H. (1986). Does the Stock Market Rationally Reflect Fundamental Values? *The Journal of Finance, 41*(3), 591-601. doi:10.2307/2328487

Sunder, S. (1995). Experimental Asset Markets: A Survey. In J. H. Kagel & A. E. Roth (Eds.), *The Handbook of Experimental Economics* (pp. 445-500): Princeton University Press.

Sung, M., & Johnson, J. E. (2010). Revealing Weak-Form Inefficiency in a Market for State Contingent Claims: The Importance of Market Ecology, Modelling Procedures and Investment Strategies. *Economica, 77*(305), 128-147. doi:https://doi.org/10.1111/j.1468-0335.2008.00716.x

Sung, Y. T., & Tainsky, S. (2014). The National Football League wagering market: simple strategies and bye week–related inefficiencies. *Journal of Sports Economics, 15*(4), 365-384. doi:https://doi.org/10.1177/1527002512466557

Surowiecki, J. (2005). *The Wisdom of Crowds*: Anchor.

Teschner, F., Stathel, S., & Weinhardt, C. (2011). *A prediction market for macro-economic variables.* Paper presented at the 2011 44th Hawaii International Conference on System Sciences.

Tetlock, P. (2004). How efficient are information markets? Evidence from an online exchange. *Social Science Research Network*. doi:http://www.columbia.edu/~pt2238/papers/Tetlock-Efficient_Info_Markets-01_02.pdf

Thaler, R. H., & Ziemba, W. T. (1988). Anomalies: Parimutuel Betting Markets: Racetracks and Lotteries. *Journal of Economic Perspectives, 2*(2), 161-174. doi:https://doi.org/10.1257/jep.2.2.161

Tryfos, P., Casey, S., Cook, S., Leger, G., & Pylypiak, B. (1984). The profitability of wagering on NFL games. *Management Science, 30*(1), 123-132. doi:https://doi.org/10.1287/mnsc.30.1.123

Tziralis, G., & Tatsiopoulos, I. (2012). Prediction markets: An extended literature review. *The Journal of Prediction Markets, 1*(1), 75-91. doi:https://doi.org/10.5750/jpm.v1i1.421

Vandenbruaene, J., De Ceuster, M., & Annaert, J. (2022). Efficient Spread Betting Markets: A Literature Review. *Journal of Sports Economics, 0*(0). doi:https://doi.org/10.1177/15270025211071042

Varian, H. R. (1980). A model of sales. *The American Economic Review, 70*(4), 651-659.

Vergin, R. C. (1977). An investigation of decision rules for thoroughbred race horse wagering. *Interfaces, 8*(1), 34-45. doi:https://doi.org/10.1287/inte.8.1.34

Vergin, R. C. (1998). The NFL pointspread market revisited: anomaly or statistical aberration? *Applied Economics Letters, 5*(3), 175-179. doi:https://doi.org/10.1080/758521377

Vergin, R. C. (2001). Overreaction in the NFL point spread market. *Applied Financial Economics, 11*(5), 497-509. doi:https://doi.org/10.1080/096031001752236780

Vergin, R. C., & Scriabin, M. (1978). Winning strategies for wagering on National Football League games. *Management Science, 24*(8), 809-818. doi:https://doi.org/10.1287/mnsc.24.8.809

Vergin, R. C., & Sosik, J. J. (1999). No place like home: an examination of the home field advantage in gambling strategies in NFL football. *Journal of Economics and Business, 51*(1), 21-31.

Vlastakis, N., Dotsis, G., & Markellos, R. N. (2009). How efficient is the European football betting market? Evidence from arbitrage and trading strategies. *Journal of Forecasting, 28*(5), 426-444. doi:https://doi.org/10.1002/for.1085

Walls, W. D., & Busche, K. (2003). Broken odds and the favourite-longshot bias in parimutuel betting: a direct test. *Applied Economics Letters, 10*(5), 311-314. doi:https://doi.org/10.1080/13504850210147162

Weber, M., & Welfens, F. (2007). How do markets react to fundamental shocks?: An experimental analysis on underreaction and momentum. Retrieved from Available at SSRN: https://ssrn.com/abstract=1022924

Weitzman, M. (1965). *Utility Analysis and Group Behavior: An Empirical Study*. Retrieved from https://EconPapers.repec.org/RePEc:hrv:faseco:3710799

William Hill PLC. (2019). *Annual Report and Accounts 2018*. Retrieved from https://www.williamhillplc.com/media/12906/williamhill_ar18-01032019.pdf

Williams, L. V., & Paton, D. (1997). Why is There a Favourite-Longshot Bias in British Racetrack Betting Markets. *The Economic Journal, 107*(440), 150-158. Retrieved from http://www.jstor.org/stable/2235276

Williams, R. J., & Wood, R. T. (2007). *Internet Gambling: A Comprehensive Review and Synthesis of the Literature* Retrieved from https://opus.uleth.ca/bitstream/handle/10133/432/2007-InternetReview-OPGRC.pdf?sequence=1

Willis, K. E. (1964). Optimum no-risk strategy for win-place pari-mutuel betting. *Management Science, 10*(3), 574-577. doi:https://doi.org/10.1287/mnsc.10.3.574

Wolfers, J. (2006). Point shaving: Corruption in NCAA basketball. *American Economic Review, 96*(2), 279-283. doi:https://doi.org/10.1257/000282806777211757

Wolfers, J., & Zitzewitz, E. (2004). Prediction Markets. *Journal of Economic Perspectives, 18*(2), 107-126. doi:https://doi.org/10.1257/0895330041371321

Wolfers, J., & Zitzewitz, E. (2006). Five open questions about prediction markets. In: National Bureau of Economic Research Cambridge, Mass., USA.

Woodland, B. M., & Woodland, L. M. (1991). The effects of risk aversion on wagering: point spread versus odds. *Journal of Political Economy, 99*(3), 638-653. doi:https://doi.org/10.1086/261770

Woodland, B. M., & Woodland, L. M. (1994). Market Efficiency and the Favorite-Longshot Bias: The Baseball Betting Market. *The Journal of Finance, 49*(1), 269-279. doi:10.1111/j.1540-6261.1994.tb04429.x

Woodland, B. M., & Woodland, L. M. (1997). Efficiency in gambling markets involving spread: a corrected and simplified test. *Applied Economics Letters, 4*(2), 93-95. doi:https://doi.org/10.1080/758526702

Woodland, B. M., & Woodland, L. M. (2000). Testing contrarian strategies in the National Football League. *Journal of Sports Economics, 1*(2), 187-193. doi:https://doi.org/10.1177/152700250000100206

World Federation of Exchanges. (2020). 2020 Market Highlights. Retrieved from https://www.world-exchanges.org/storage/app/media/FH.FY%202020%20Market%20Highlights%20v9.pdf

Yagudin, M., Sempere, N., & Lifland, E. (2021). Prediction Markets in The Corporate Setting. Retrieved from https://forum.effectivealtruism.org/posts/dQhjwHA7LhfE8YpYF/prediction-markets-in-the-corporate-setting#comments

Yergin, D., & Stanislaw, J. (1998). The Commanding Heights: The Battle between Government and Marketplace That Is Remaking the Modern World. *Simon and Shuster, New York*.

Zbaracki, M. J., Ritson, M., Levy, D., Dutta, S., & Bergen, M. (2004). Managerial and customer costs of price adjustment: direct evidence from industrial markets. *Review of Economics and statistics, 86*(2), 514-533. doi:https://doi.org/10.1162/003465304323031085

Zuber, R. A., Gandar, J. M., & Bowers, B. D. (1985). Beating the spread: Testing the efficiency of the gambling market for National Football League games. *Journal of Political Economy, 93*(4), 800-806. doi:https://doi.org/10.1086/261332