Title
Essays on the Economics of Suspense, Surprise, Superstars, and Soda

Author
Kaplan, Scott Meyer

Publication Date
2021

Peer reviewed|Thesis/dissertation

Essays on the Economics of Suspense, Surprise, Superstars, and Soda

by

Scott M. Kaplan

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor David Zilberman, Co-chair
Professor James Sallee, Co-chair
Professor Benjamin Handel
Professor Sofia B. Villas-Boas

Spring 2021

Essays on the Economics of Suspense, Surprise, Superstars, and Soda

Abstract

Essays on the Economics of Suspense, Surprise, Superstars, and Soda

by

Scott M. Kaplan

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor David Zilberman, Co-chair

Professor James Sallee, Co-chair

Economists have always been interested in understanding and estimating demand for different goods. Demand is affected by a number of factors, but one that is very important is information. In this work, I estimate the impacts of both instrumental and non-instrumental information on demand for important and commonly consumed products. Instrumental information is information received by an agent that leads to a contingent action. For instance, receiving information about the negative health effects of consuming sugar sweetened-beverages may affect an agent's decision about how many sugar-sweetened beverages they consume. On the other hand, non-instrumental information is information consumed for the sake of entertainment and attention-capture. For example, information disseminated through the news or a novel may not lead to any direct actions, but is enjoyable and welfare improving nonetheless. My dissertation aims to explore the economic implications of demand in response to information.

Chapter one, titled *Entertainment Utility from Skill and Thrill*, uses revealed preference methods to estimate demand for non-instrumental information in entertainment. I do this by examining the "thrill" associated with the trajectory of an event, which includes both suspense and surprise, and the "skill" of performers in an event. I apply the theory presented in Ely et al. (2015) to conduct an empirical analysis that examines the effect of thrill on consumer attention. I extend the Ely et al. (2015) framework by examining spectator preferences for characteristics of the performers themselves, which I call "skill." I use game-specific, high-temporal frequency television ratings data from the National Basketball Association (NBA) to measure spectator responses to skill and thrill. First, I find that a doubling of skill present in a game leads to an approximately 11% increase in initial viewer turnout, while the expected thrill of a game has no statistically significant impact. Next, I show that thrill during a game increases viewership by 7-30%, while a doubling of skill

on the court during a specific portion of a game leads to a 1.9-2.4% increase in viewership, depending on specification. Interestingly, I find a negative interactive effect between suspense and skill, suggesting that heightened suspense leads to differentially higher viewership with lower skill on the court. The findings suggest that skill of information-conveying agents primarily impacts viewership on the extensive margin (across games), while thrill is highly time-dependent and primarily impacts viewership on the intensive margin (within games). These findings have important implications for entertainment media companies, including leagues and television broadcasters, and advertisers.

Chapter two, titled *The Economic Value of Popularity: Evidence from Superstars in the National Basketball Association*, estimates spectator willingness-to-pay for superstars in the National Basketball Association. Using microdata from an online secondary ticket marketplace and exogenous player absence announcements, I find 4-16% ($7-$42) reductions in prices when superstars are announced to miss games. Additionally, LeBron James and Stephen Curry exhibit even larger impacts in away game absences–21% ($75/ticket) for LeBron and 18% ($55/ticket) for Curry. The results suggest popularity is a more significant determinant of WTP than ability, and in line with existing superstar literature, popularity predicts price impacts convexly. This paper provides a novel methodology to estimate superstar value, generating implications for the entertainment industry.

Chapter three, titled *Soda Wars: The Effect of a Soda Tax Election on University Beverage Sales*, examines how soda sales changed due to the campaign attention and election outcome of a local excise tax on sugar-sweetened beverages (SSB), commonly referred to as a soda tax. Using panel data of beverage sales from university retailers in Berkeley, California, we estimate that soda purchases relative to control beverages significantly dropped immediately after the election, months before the tax was implemented in the city of Berkeley or on campus. Supplemental scanner data from off-campus drug stores reveal this result is not unique to the university setting. The findings suggest soda tax media coverage and election outcomes can have larger effects on purchasing behavior than the tax itself.

Contents

List of Figures

List of Tables

Acknowledgments

David Zilberman provided me with an opportunity to develop a love and passion for research. He gave me a chance as a second year undergraduate, and has been an invaluable mentor in all dimensions of my life ever since. He inspired me to get a Ph.D., supported me both personally and financially throughout my entire graduate school career, and always gave me the freedom and flexibility to pursue research ideas and projects I was interested in. He involved me as a collaborator, gave me increasing levels of responsibility in projects, and helped me develop an extensive network of other colleagues. I am forever indebted to him, and consider him a lifelong mentor and friend. He and his wife, Leorah, have always made me feel like I have a second family in Berkeley.

Jim Sallee provided invaluable guidance on all of my dissertation research. He always made thoughtful comments on my work, challenged me to explore all dimensions of a problem, and supported my interests and overall development as a researcher. The quality of my research has been massively improved because of Jim's influence. His creativity, work ethic, and attention to detail are characteristics I hope to carry forward in my future career as a researcher and teacher. Sofia Villas-Boas was a mentor to me during my first year of graduate school, and generously involved me in several different projects. She helped me learn how to write an empirical paper. She is and will continue to be a valuable collaborator and friend. I would not have made it through the job market without her. Ben Handel was always willing to talk and supported both my research and career interests. His ability to simultaneously conduct high quality research, convey difficult course material intuitively, and provide unending and genuine support to his students are a testament to his impact as an advisor.

Carmen Karahalios was instrumental in my success and happiness as a researcher in the Department of Agricultural and Resource Economics. She was always someone I could reach out to on just about anything. She and Diana Lazo helped me seamlessly navigate the Berkeley bureaucracy, always having my best interests in mind. My graduate cohort provided a fun and stimulating learning and research environment. My research has been substantially improved because of their advice and willingness to be constructively critical. I cannot imagine having gone through graduate school with a better group of peers.

I also want to thank all of the teachers I have had throughout my life. Their passion for and love of learning was contagious, and I am excited to continue to develop and pursue a joint career in research and teaching.

I would not be where I am today without the support of my family. My mom, Connie Tanquary, fostered a love of learning when I was a child. She taught me to remain humble and work hard, and to let the hard work lead me to what I was passionate about. My dad, Karl Kaplan, showed me the impact of higher education on the pursuit of a career. My brother, Grant Kaplan, has always been interested in my work and has celebrated my successes. My grandma and grandpa, Helen and Bob Tanquary, have always supported me and encouraged me to pursue my dreams. I am so lucky to have such a loving family.

Last but not least, I want to thank the University of California, Berkeley. I am so fortunate to have spent my entire adult life to date as a part of the Berkeley community. I have learned and taken so much from this experience. Berkeley has taught me how to think inside and outside of the classroom. I have learned the importance of close community and embracing diversity. I have made lifelong friends. I will always consider Berkeley a place I call home.

Chapter 1

Entertainment Utility from Skill and Thrill[1]

1.1 Introduction

Access to information is a crucial component of an economic agent's decision-making process. Information leading to such contingent actions is defined as instrumental. Instrumental information applies to the entire spectrum of economic decisions, for instance how gas prices influence which type of car to buy, how a sugar-sweetened beverage tax impacts soda consumption, or how wages in a certain industry impact whether or not to change jobs. In particular, this information provides additional certainty about a subsequent decision, which leads to welfare-improving actions, and it is often the case that agents are willing to pay a premium for such information because of the additional certainty it offers. In contrast, non-instrumental information does not have direct consequences for economic decision-making under constraints, but provides utility nonetheless. For instance, individuals may be attentive to the performance of candidates in a political debate, how a television series will play out, or which team will prevail in a sporting event. In situations featuring non-instrumental information, uncertainty over an outcome is itself a source of pleasure for individuals.

Most sources of non-instrumental information are found in entertainment settings, since the uncertainty associated with the information is not associated with a financial stake. The

[1]Thanks to Cheenar Gupte for excellent research assistance. Special thanks to my core advising team: Ben Handel, Jim Sallee, Sofia Villas-Boas, and David Zilberman. I'd also like to thank Nola Agha, Anocha Aribarg, Max Auffhammer, Giovanni Compiani, Kwabena Donkor, Karl Dunkle-Werner, Claire Duquennois, Gabe Englander, Alexander Frankel, Hal Gordon, Andy Hultgren, Benjamin Krause, Megan Lang, Brian Mills, Kate Pennington, Jeff Perloff, Dan Putler, Gordon Rausser, Leo Simon, Avner Strulov Shlain, Steve Tadelis, Dmitry Taubinsky, and Miguel Villas-Boas for thoughtful comments on this research. I'd also like to thank participants in the United States Naval Academy Economics Seminar, UC Berkeley IO Seminar Series, the North American Association of Sports Economics Annual Meeting, the Marketing Summer Research Series in the Haas School of Business, the ERE Seminar Series at UC Berkeley, and the NYU Stern Marketing Seminar Series. Finally, special thanks to Ryan Davis for help in acquisition of the play-by-play data, and Todd Whitehead for useful feedback.

global entertainment media industry exceeds $2 trillion, and has grown 60% over the last 10 years (PWC 2019). Entertainment in its current form does not exist without well-crafted and targeted information updating that attracts and keeps consumers' attention. Additionally, provision of non-instrumental information in certain entertainment settings has important social implications. The ability to retain consumers through media outlets allows them to remain informed about important, economically consequential issues.

One can think of the outlay of non-instrumental information as the "thrill" associated with an event. Thrill refers to adjustments in a spectator's belief state as a result of new information about an outcome. Ely et al. (2015) define two primary characteristics of thrill: suspense and surprise. Higher suspense is defined as higher variance in future beliefs over an outcome, and higher surprise is defined as a larger difference in current beliefs about an outcome compared to previous beliefs. For instance, suppose a golfer is entering the final nine holes of a tournament in second place. There is clear suspense over whether or not the golfer will prevail—beliefs are going to update relatively soon given the approaching finality of the event. But on the 13th hole, the golfer drives the tee shot into the water! This constitutes a significant change in the belief state about the golfer's chances to win.

Additionally, the extent to which thrill is meaningful may depend on the "skill" of the golfer, tennis player, or political candidate conveying the information. A political debate between Joe Biden and Donald Trump likely garners much higher overall attention than a debate between candidates for a local election. I define "skill" using characteristics of the performers themselves, which includes measures of productivity and popularity.

This paper uses revealed preference methods to explore and quantify demand for non-instrumental information in entertainment, examining the "thrill" associated with the trajectory of an event, and the "skill" associated with the performers involved. I take the theory presented in Ely et al. (2015) to conduct an empirical analysis that examines the effect of thrill on consumer attention. I extend the Ely et al. (2015) framework by examining spectator preferences for skill. I employ game-specific, minute-by-minute television ratings data from the National Basketball Association (NBA) during the 2017-18 and 2018-19 seasons to measure viewership in response to skill and thrill. Thrill is the suspense and surprise experienced during the course of a game, as defined in Ely et al. (2015). I define the skill of a specific player as the total number of fan All-Star votes they receive in a given season. While there are many different avenues of entertainment to study non-instrumental information, live sports is a natural application since (i) the skill of players is directly observed and publicly available, (ii) outcomes are plausibly random conditional on an initial information state, unlike a book or movie, and (iii) because of the size of and value generated by the industry.

I rely on two different empirical strategies to measure two separate dimensions of these impacts. First, I estimate *initial* viewership turnout in response to the presence of total skill and the *expected* thrill of a game. Next, I utilize within-game play-by-play data at the second-of-game level, where I observe the level of skill on the court, score differential, and real-time win probabilities for each team, to assess television viewership responses to skill and thrill as they evolve *during* a game. I use these within-game estimates to understand the viewership impact of counterfactual game structures.

The findings suggest that skill and thrill play important, but different, roles in generating viewership. First, I find that a doubling of skill present in a game leads to an approximately 11% increase in initial viewer turnout, while expected thrill has no statistically significant impact. For context, the presence of LeBron James, who led all players in All-Star votes during the 2017-19 seasons, corresponds to approximately 120% of the average aggregate number of All-Star fan votes of *all* players in a game. Thus, the presence of LeBron alone results in an approximately 13.5% increase in initial TV viewership. These results are remarkably similar to those found in Kaplan (2020), which uses secondary ticket marketplace data to assesses the impact of a superstar absence announcement for a specific game on listed prices, finding that the absence of LeBron James leads to a 13% ($42/ticket) average reduction in ticket prices.

Next, I use the evolution of absolute score differential over the course of a game to measure viewership responses to thrill. This analysis uses a more observable measure of thrill than the structural definitions from Ely et al. (2015). I find that a one-point decrease in the absolute score differential does not impact viewership in the first or second quarters, but increases viewership by 0.6% and 1.2% in the third and fourth quarters, respectively, strongly supporting the hypothesis that viewers relish thrilling games, not just games that are close. Contextualizing these results further, second half viewership is 8.2-20.5% lower on average for games with a 14+ score differential compared to a 0-8 differential, while these differences are 12.0-29.6% when only examining the fourth quarter. I extend this analysis to look at absolute score differential during a game *in reference to the closing point spread*, similarly finding that viewership declines are starker towards the end of games. I find that for every one-point increase in the score differential from the closing spread, viewership declines by 0.1-0.9%, with larger decreases found in later stages of a game. This suggests that a one-standard deviation change in score differential in reference to the spread during the final quarter segment (9.3 points) exhibits an economically meaningful impact on viewership (6.4-7.3% reduction), which amounts to roughly half the size of the impact of thrill over the absolute score differential.

Finally, I jointly assess within-game viewership impacts from suspense, surprise, and skill, directly implementing the structural definitions of suspense and surprise from Ely et al. (2015). I find that a doubling of suspense during a game increases viewership by 0.4-0.6%, and a doubling of surprise by 0.6-1.0%, not accounting for additional or differential impacts associated with skill. While these magnitudes are seemingly small, suspense and surprise can take on an extremely large range of values. For instance, in the last segment of the fourth quarter, a 0-2 point game averages 18 times more suspense than a 14+ point game. In this case, viewership would be approximately 6.8-10.8% higher through suspense alone. On the other hand, the range of surprise exhibited over the course of a game is slightly lower than that of suspense. Specifically, in the fourth quarter a 0-2 point game features 14 times more surprise on average than a 14+ point game, which would translate into a viewership increase of 9.0-13.4%.

In the fully interacted model, I find that a doubling of skill on the court *at a given time* *during the game* leads to a 1.9-2.4% increase in viewership. The comparison of responses to skill and thrill using initial versus within-game viewership suggest that viewers respond to skill primarily on the extensive margin (across games), while responses to thrill take place primarily on the intensive margin (within games). In other words, viewers are much more likely to be interested in a game prior to it starting because of the skill of the players involved, but are more likely to respond within a game as thrill evolves. Interestingly, I also find a negative interactive effect between suspense and skill, suggesting that heightened suspense leads to differentially higher viewership with lower skill on the court, supporting the traditional notion that spectators may only turn on games featuring lesser-known players (or teams) if they're nearing the end and exhibiting sufficiently high suspense. I find no evidence for an interactive effect between skill and surprise, and in fact when conducting the joint estimation featuring both thrill and skill, I find little evidence to suggest surprise impacts viewership in this entertainment setting.

The remainder of this paper proceeds as follows. First, I review the bodies of literature this work is motivated by and contributes to in section 1.2. Next, I develop a model of spectator utility from entertainment in section 1.3. I then overview the data, develop the set of empirical strategies used in estimating viewership responses to skill and thrill, and present the results of the analysis in section 1.4. Section 1.5 contextualizes and provides an economic interpretation of the results, while also presenting a counterfactual analysis assessing viewership responses to an alternative game structure. This section also discusses the implications of contest design in generating additional thrill. Finally, section 1.6 concludes.

1.2 Literature Review

This research contributes to several notable bodies of literature. First and foremost, there is a small existing literature on suspense and surprise. Ely et al. (2015) provide the definitions of suspense and surprise used in this analysis and is the primary existing study on this topic. They determine the optimal suspense and surprise information policies that maximize expected utility. Their study incorporates practical examples from entertainment and socially-relevant settings, including novels, political races, and live sports. I expand on their work by examining the quality of the performers themselves, and how the presence of these agents can affect responses to suspense and surprise in an event. Preceding studies have also examined modified versions of suspense and surprise in a theoretical manner and in various settings, including live sports (Bryant et al. 1994, Su-lin et al. 1997), game shows (Chan et al. 2009), and in the context of the Hangman's Paradox (Geanakoplos et al. 1989, Geanakoplos et al. 1996; Borwein et al. 2000). An adjacent literature uses laboratory experiments to measure physiological responses to suspense and surprise, emphasizing that animals are genetically driven to respond to such occurances (Itti and Baldi 2009; Ranganath and Rainer 2003; Fairhall et al. 2001; Ebstein et al. 1996).

To the best of my knowledge, there have been two peer-reviewed, empirically-oriented studies to date using the suspense and surprise framework developed in Ely et al. (2015).

Bizzozero et al. (2016) examine television viewership responses to suspense and surprise over the course of tennis matches, finding that surprise, and to a lesser extent, suspense, generate positive but relatively small viewership impacts. In particular, they find that a one standard deviation increase in suspense (surprise) raises audience viewership by 1,260 (2,630) viewers per minute, which combine to cause a 3.65% viewership increase. They implement two separate, but similar, methodologies to measure impacts of suspense and surprise: a Markov method and a "betting odds" method, which uses live betting odds between each point during a match to dictate outcome probabilities. Buraimo et al. (2020) examine television viewership in response to suspense and surprise using the European professional football market. They also introduce "shock," at each portion of a match, which is defined as the difference between current outcome probabilities and expected probabilities prior to the start of a match. Their findings also suggest relatively small impacts of suspense and surprise on viewership; a one standard deviation in both suspense and surprise increase audience viewership by 1.2%. Two recent working papers have assessed viewership responses to suspense and surprise in esport tournament streams (Simonov et al. 2020) and professional baseball (Liu et al. 2020).

This paper aims to extend the suspense and surprise literature in several key ways. First, I explore a broader question that includes how the quality of agents performing in these events affects viewership, providing evidence of the relative magnitude impact of skill and thrill. Second, I examine viewership responses to thrill over alternative game outcomes, which may be unrelated to the final outcome of who wins or loses. In particular, I explore suspense and surprise with respect to the closing point spread of a game, finding statistically significant and economically meaningful impacts. Third, I use my estimates of viewership responses to thrill to assess the impact of a counterfactual game structure that leads to higher levels of thrill by construction. Finally, I examine an entirely different sport and geographic market: professional basketball in the United States. There are notable differences between the structure of the event and the types of spectators watching games, which may partially explain magnitude differences in television viewership responses in this paper compared to previous work.

The second body of literature focuses on information preferences, which includes the theory of addictive goods, and outcome resolution, formalizing the notion that individual taste preferences are consistent with utility-maximizing behavior and may change over time (Stigler and Becker 1977; Becker and Murphy 1988; Kreps and Porteus 1978; Caplin and Leahy 2001). I aim to expand on this work by discussing and evaluating preferences for non-instrumental information, especially in the context of outcome resolution. In particular, evaluating the psychological and emotional attributes of entertainment is important in understanding the types of information individuals desire (Fowdur et al. 2009). For instance, studies have shown that story "spoilers" have large impacts on demand for entertainment goods, even suggesting that they have the potential to increase consumer enjoyment (Leavitt and Christenfeld 2011; Johnson and Rosenbaum 2015; Levine et al. 2016; Ryoo et al. 2020). Naturally, there has also been significant research assessing the impact of outcome uncertainty on demand for live sports (Rottenberg 1956; Knowles et al. 1992; Humphreys and Miceli 2019; Alavy et al. 2010; Forrest et al. 2005).[2] I extend this research by more closely examining the evolution of beliefs over the course of an event, using random variation in event trajectories to assess attention-based responses. This is particularly important as audiences increasingly explore real-time gambling in live sports, which is likely to depend heavily on information relayed throughout the course of an event (Kaplan and Garstka 2001; Haugh and Singal 2020; Salaga and Tainsky 2015).

The third relevant body of literature is in hedonic pricing. Rosen (1974) provides a theoretical framework that describes the total value of a good as a combination of the values of its attributes, which has led to a rich body of literature applying the concept to a wide range of products (Busse et al. 2013; Sallee et al. 2016; Currie and Walker 2011; Chay and Greenstone 2005; Luttik 2000). This work focuses on two primary attributes of entertainment goods–the skill of the performers and the thrill of the event itself. Television ratings data is a natural avenue to explore impacts of these characteristics on consumer attention, as there has been other work examining viewership responses to well-defined programming characteristics (Fournier and Martin 1983; Anstine 2001; Livingston et al. 2013). Furthermore, there is existing work using hedonic pricing methods in entertainment to understand the value of star performers (Scully 1974; Kahn 2000; Rosen 1981; Hausman and Leonard 1997; Krueger 2005; Chung et al. 2013, Grimshaw and Larson 2020; Kaplan 2020). To the best of my knowledge, there is no existing research jointly measuring the impact of skill and thrill on demand.

The fourth and final body of literature is on the economics of advertising and consumer attention. Many forms of entertainment rely on advertising as a large source of revenue, and advertisers themselves pay for the quantity and types of consumers the entertainment attracts (Becker and Murphy 1993; Wilbur 2008; Bertrand et al. 2010; Hartmann and Klapper 2018). The stakes for advertisers are quite high – analyzing time-use survey data, Aguiar et al. (2013) finds that the average American spends about 20% of their time consuming some form of entertainment. The skill of performers and evolution of thrill during the course of an event is paramount in generating spectator attention, and this work aims to assess the extent to which each contributes to recruitment and retention of viewers. Furthermore, the type of information content used by advertisers in entertainment settings is important for generating meaningful engagement with potential customers (Resnik and Stern 1977; Bagwell 2005). In particular, there is a clear differentiation between informative content, which corresponds characteristics like prices and deals, and emotional content, which corresponds to characteristics like humor, slang, and emojis. Studies have shown that provision of emotional content leads to higher levels of consumer engagement (Aaker 1997; Lee et al. 2018). In fact, Madrigal and Bee (2005) find that the use of suspense as an advertising tactic is an important driver of consumer attention. Using revealed preference methods to understand how consumer attention responds to skill and thrill is important in understanding how to

[2] It is important to note that while thrill and outcome uncertainty are related, they characterize different processes. Outcome uncertainty examines probabilities of different outcomes happening at different times, while thrill looks more fundamentally at the variance in the evolution of beliefs over the course of an event.

better engage audiences with different advertising strategies.

1.3 Model of Utility from Entertainment

This section presents a conceptual framework to understand consumer demand for skill and thrill. While the definition of player skill is straightforward (the number of fan All-Star votes received), thrill requires a more structured definition of the specific characteristics that lead to its evolution during a game. Specifically, I separate thrill into two distinct components: suspense and surprise. I rely on structure from Ely et al. (2015) to formally develop a mathematical interpretation of suspense and surprise that can be used in conjunction with skill to assess spectator preferences.

Defining Suspense and Surprise

Suppose there are two teams, A and B. Teams can be thought of as sports teams, political candidates, or characters in a movie, book, or play. Suppose each team i is defined by their strength, $V_i \in \mathbb{R}^+$. Denote Team A's strength as V_A and Team B's strength as V_B. Team A and B compete in an event lasting T periods, where the outcome is fully resolved in period T when a winner is declared. Let $P_t(A)$ denote the probability at time t that Team A wins, and $1 - P_t(A)$ the probability at time t that Team A does not win, where the set $P_t = \{P_t(A), 1 - P_t(A)\}$ represents an outcome probability pair at time t. For the specific case of $t = 0$, $P_0(A) = \frac{V_A}{V_A + V_B}$ and $P_0(B) = 1 - P_0(A)$, representing the prior belief that each respective team will emerge victorious at $t = T$.

Let beliefs about future outcome probabilities take the following structure. At time t there is a belief martingale $\tilde{\mu} = (\tilde{\mu}_t)_{t=0}^T$, which is a sequence of beliefs about future outcome probability pairs believed at time t. Assume now that $\tilde{\mu}$ evolves as a first-order Markov process over $t = 1, ..., T$. Namely, $\mathbb{E}[\tilde{\mu}_{t+1}|\mu_0, ..., \mu_t] = \mu_t$ for all $t \in \{1, ..., T\}$. It is important to note that with this structure, there must be a sequence representing realized outcome probability pairs observed at each t, P_t. Denote this sequence $\mu = (P_t)_{t=0}^T$. Additionally, beliefs about some period $t + n$ while based at time t are written as $\mathbb{E}[\tilde{\mu}_{t+n}|\mu_0, ..., \mu_t] = \mu_t$ where $n \in [1, T - t]$. With this setup, I define suspense at time t, X_t, as follows:

$$X_t = \mathbb{E}_t \left[(\tilde{\mu}_{t+1} - \mu_t)^2 \right] \tag{1.1}$$

Thus, suspense at time t is higher when there is higher variance in beliefs about the difference in the probability pair at time $t + 1$ and the realized probability pair at time t. In words, the larger the potential swings in the probability pair between period t and $t + 1$, the higher the suspense. Due to the Markovian nature of the setup, $\mathbb{E}_t[P_{t+1}] = P_t$.

Surprise, on the other hand, is a backward-looking (*ex-post*) belief. An agent only experiences surprise in response to something that has already transpired. Using the framework above, surprise Y_t is defined as follows:

$$Y_t = (\mu_t - \mu_{t-1})^2 \tag{1.2}$$

Thus, there is higher surprise at time t the larger the variance in the realized probability pair between time t and time $t-1$. In words, the larger the "swing" in realized outcome probabilities between $t-1$ and t, the higher the surprise in t. It should be noted that surprise is tightly interlinked with suspense—an event with a large amount of surprise may lead to a more or less suspenseful state at time t.

Model of Entertainment Utility

In this section, I introduce a novel model of spectator utility derived from entertainment, which includes suspense, surprise, and skill. I begin by developing a framework that generates utility for individual i from an entertainment event j. Denote the total skill of all players in a game as S_j, which I assume to be time invariant and continuous, and the thrill during a specific portion of a game as $H_j(r) = X_j(r) + Y_j(r)$, where r is a continuous measure of time remaining in an event. Thus, the expected thrill of an entire event can be written $\int_{r=R}^{0} \mathbb{E}[H_j(r)]dr$ where R represents the length of an entire event.[3]

Additionally, assume the cost of watching $C(t, X_j)$ to be a function of time spent watching t and all time-invariant, event-specific costs X_j. I assume $\frac{\partial C(t,X_j)}{\partial t} > 0$ and $\frac{\partial^2 C(t,X_j)}{\partial t^2} > 0$, and $t = T$ denotes the maximum time that can be spent watching an event. Thus, the utility for individual i from game j can be written as follows:

$$U_{ij} = B_j - C(t, X_j) + \psi\left[S_j * t\right] + \int_0^t \phi \mathbb{E}[H_j(r)]dr + \theta\left[(S_j * t) * \int_0^t \mathbb{E}[H_j(r)]dr\right] + \xi_i + \epsilon_{ij} \tag{1.3}$$

where ψ and ϕ are average marginal utilities from skill and thrill, respectively. I also allow for an interactive effect of skill and thrill on utility, where θ represents the average marginal utility associated with this interaction. For example, an event with large levels of skill and thrill may exhibit differentially higher (or lower) utility than the additive components of skill and thrill alone. I assume here that an individual experiences the skill in an event linearly with time spent watching, although this assumption will be relaxed.[4] B_j represents some baseline average utility from event j, ξ_i an individual utility shifter, and ϵ_{ij} an i.i.d. residual term.

There are two choices an individual must make in their decision to watch an event: a choice of the amount of time to spend watching, t^*, and how to allocate t^* across a game.

[3] In contrast to the definitions of suspense and surprise, I rely on continuous time notation in the development of the model in order to derive solutions analytically. However, the implications are analogous for a setup using discrete time.

[4] This assumption will be relaxed in different ways. For instance, skill in an event can actually be measured as "skill-minutes," which accounts for the length of time players of different skill spend actually playing in a game, and how that overlaps with time spent watching.

Here, I rely on the assumption that $\frac{\partial \mathbb{E}_{r=R}[H_j(r)]}{\partial r} < 0$, which suggests that expected thrill (at time remaining R) is monotonically increasing over the course of an event. With this assumption in place, an individual making an *ex-ante* decision about how much time to spend watching a game should choose to allocate their time beginning with the end of an event, working backwards.[5] With this structure, t^* is the solution to the following:

$$\arg\max_t U_{ij}(t) = -C(t, X_j) + \psi(S_j * t) + \int_0^t \phi \mathbb{E}[H_j(r)]dr + \theta \left[(S_j * t) * \int_0^t \mathbb{E}[H_j(r)]dr \right] \quad (1.4)$$

$$\frac{\partial C(t, X_j)}{\partial t} = \psi S_j + \phi \mathbb{E}[H_j(t^*)] + \theta \left[S_j * \int_0^{t^*} \mathbb{E}[H_j(r)]dr + (S_j * t^*) * \mathbb{E}[H_j(t^*)] \right] \quad (1.5)$$

This result suggests that the optimal time spent watching is determined by setting the marginal opportunity cost of time spent watching equal to the marginal benefit of from watching. The marginal benefit of watching is the sum of the marginal benefit from additional skill, additional thrill, and the interaction of skill and thrill.[6] The empirical analysis will directly estimate the marginal utilities from skill, thrill, and their interaction.

Linking Theory to the Empirical Application

The theoretical framework provides structural definitions of suspense and surprise, and a general utility function for spectators of entertainment that incorporates skill. Now, I apply the theoretical framework to examine viewership responses to both skill and thrill (1) before a game begins and (2) as a game evolves.

To measure thrill, I use real-time outcome probabilities computed and observed at the second-of-game level to calculate suspense and surprise at each second of a game. In particular, I use a pre-defined forward (backward) looking window W, which is measured in seconds-of-play within a game, to compute suspense (surprise). I compute the variance of observations found within W to obtain a second-of-play measure of suspense and surprise. To measure skill, I use the total number of fan All-Star votes received by each player playing within a specific game.

I take a three-fold approach to measuring viewership responses to skill and thrill. First, I measure how *initial* viewership (i.e., the number of individuals tuning in for the start of a game) responds to total skill present in a game, and *ex-ante* thrill expected in a game.

[5] In reality, an individual may not make an *ex-ante* decision about how much time to spend watching an event. However, they may have a well-formed prior about how much time they plan to watch given expected thrill, and then adjust using some stochastic process depending on the progression of an event. Thus, it may be the case that the expression is not monotonic for $r < R$. Section **??**3.3 presents evidence of the validity of this assumption.

[6] Note that with the assumption $\frac{\partial \mathbb{E}_{r=R}[H_j(r)]}{\partial r} < 0$, time spent watching t and time remaining r are identical.

I measure *ex-ante* expected thrill using variation from the initial point spread given for a specific matchup.[7] Next, I estimate viewership responses over the course of a game to "observable" thrill. Observable thrill is defined as the absolute point differential at different points of a game. Although observable thrill does not directly implement the definitions of suspense and surprise presented in Ely et al. (2015), it provides a proxy for thrill that is more directly observed by the viewer (since the viewer is not likely to observe second-by-second real-time outcome probabilities). Finally, I jointly implement the structural definitions of suspense and surprise, as well as skill, to estimate viewership responses within a game. In particular, I use variation in skill present on a court *at specific times of a game*.[8] With this framework, I am also able to determine the interactive effect of skill and thrill on viewership.

A Note on Alternative Outcomes and Viewership Response Mechanisms

As constructed, the model assumes that thrill manifests itself with respect to beliefs about the final outcome of an event, which takes place at time T. However, suspense and surprise can be generalized to refer to beliefs about a state *within* an event. For instance, instead of $P_t(A)$ referring to the probability at time t that Team A will emerge victorious, $P_t(A)$ could refer to the probability that Team A makes a highly improbable shot (i.e. a half-court buzzer beater) at time t. Specifically, suppose that such a shot takes place at time $t-1$ and the outcome of the shot at time t. The same definitions of suspense and surprise would apply: an agent would experience suspense during $t = \{0, ..., t-1\}$ as to whether or not such a shot will go in, which may be more suspenseful if Team A has a player known for taking and making these types of shots. An agent would experience some amount of surprise at time t depending on whether or not the shot goes in at t. This stylistic example is important in explaining why agents may experience suspense and surprise with respect to moments during an event that have little to no bearing on the event's final outcome. In the context of the empirical analysis, this generalization will be useful in examining viewership responses to suspense and surprise over alternative outcomes.

It is also important to expand upon mechanisms for within-game viewership responses to suspense, surprise, and skill. There are two primary ways viewership may respond: through *viewer recruitment* and *viewer retention*. In the case of suspense, both viewer recruitment and viewer retention are likely to occur. For example, a potential viewer who is not currently watching may be alerted in some way about a game reaching some level of suspense, and decide to tune in. A more naive viewer may be channel surfing and determine a game has a necessary threshold of suspense to stop and tune in. A game that becomes more suspenseful is also more likely to retain viewers who were already tuned in before suspense increased. In the case of surprise and skill, viewer retention is a more likely mechanism for increased viewership than viewer recruitment. To be surprised, a viewer must have been watching at both $t-1$ and t, and so a viewer may not be inclined to enter an event

[7]In other words, I measure initial viewership in response to *ex-ante* expected thrill *only* through variation in the initial point spread.

[8]For instance, skill changes within a game when a player is playing versus off the court.

because something surprising took place. On the other hand, surprise witnessed by viewers already watching is likely to lead to significant viewership retention. Additionally, a viewer is not likely to respond to a superstar re-entering the game, rather is more likely to turn off the game when a superstar is substituted out. While the viewership data I have access to does not allow me to separately identify these mechanisms, empirical estimates can be interpreted under this general framework. Future work can directly assess the relevance of each of these mechanisms using household-level viewership data as well as complementary data from information-providing applications (e.g. Twitter).

1.4 Television Viewership Responses to Skill and Thrill

This section presents an overview of the data used in the analysis, the empirical strategy, and the estimation results. The empirical approach is three-fold. First, I examine initial viewership responses to skill and expected thrill, where expected thrill is measured via the relationship between cumulative observed thrill and the initial point spread of a game. I then estimate viewership responses within a game to "observable" thrill, as measured by changes in the absolute score differential over the course of a game. Finally, I jointly estimate the viewership response to skill and thrill using the structurally defined parameters laid out in Section 1.3.

Overview of Data

There are three primary sets of data used in the analysis: (i) fixed game characteristics data providing time-invariant information about each analyzed game, (ii) second-of-game play-by-play data indicating detailed information about each moment of the game, including the real-time outcome probability for each team, and (iii) high temporal frequency television viewership data from The Nielsen Company.[9]

Game Characteristic Data

Game characteristics data, which includes time-invariant information about each game in the sample, was collected from `NBA.com`, `fivethirtyeight.com`, and `Basketball Reference` for all NBA games (regular season and playoffs) during the 2017-18 and 2018-19 NBA seasons. Most important of these characteristics include the home and away teams, time-of-day, network (local or the specific nationally-televised network), the initial point spread, and an extensive list of team- and player-specific characteristics associated with each matchup.

[9]Data granted from The Nielsen Company (US), LLC. The conclusions drawn from the Nielsen data are those of the researchers and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

Play-by-Play Data

Play-by-play data characterizes every "meaningful" action within a game, and is provided at a second-of-play level. A non-exhaustive list of common occurrences warranting an observation include a made or missed basket, turnover, foul, out-of-bounds stoppage, or timeout. Most importantly, this data characterizes the real-time score and win probability at each second of play a game, as well as a "wall clock" variable representing the time-of-day associated with each observation.[10] The last component is crucial, since it allows for accurate and precise merging of the play-by-play data with the TV ratings data, which are denoted in time-of-day units.

Television Ratings Data

The final dataset used in this analysis is TV ratings data acquired from The Nielsen Company.[11] The data includes 15-minute interval ratings for every nationally televised NBA game from the 2017-18 and 2018-19 seasons (including playoffs). The relevant metric for this analysis is the projected total number of individuals watching during any given 15-minute interval.

Summary Statistics

Table 1.1 presents conventional summary statistics for the data used in the viewership analysis. The table is broken down into two separate parts: (i) characteristics that are static and do not adjust over the course of a matchup (fixed-game characteristics) and (ii) characteristics that dynamically change during a matchup (within-game characteristics). One can see that there are 477 different games analyzed in this study, and nearly 1.4 million unique "plays," as given by the play-by-play data. In the fixed-game characteristics, there is good observed variation in the expected competitiveness of matchups, as given by the distribution of the "Point Spread" variable. The within-game data includes the primary characteristics used in the analysis of suspense, surprise, and stardom. There is substantial variation in "Total Viewership," where the least-viewed games attract hundreds of thousands of viewers, and the most-viewed games receive tens of millions of viewers.

Figures 1.1 and 1.2 show the average viewership trajectories by absolute score differential quintile and initial point spread, respectively. Interestingly, there is a nearly monotonic upward trend in viewership during a game. This may reflect several possible dynamics,

[10] According to Inpredictable, the real-time win probability is a function of game time, point differential, possession, and the closing point spread. A locally-weighted logistic regression is performed at each second of the game, where the smoothing window shrinks as the game progresses. For the final few seconds of the game, regression is abandoned in favor of a decision tree approach. There are additional complexities associated with "non-possession states," which account for times during the game when neither team discretely possesses the ball. The `locfit` package in `R` was used to perform the analysis.

[11] Data granted from The Nielsen Company (US), LLC. The conclusions drawn from the Nielsen data are those of the researcher and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

Table 1.1: Summary Statistics

		Mean	SD	Min	Max	N
Fixed-Game Characteristics						
	Cum. All-Star Votes (1,000s)	6,710.16	3,893.84	372.69	17,035.61	477
	Point Spread	4.88	3.56	0	18	477
	Cum. PER	313.85	34.74	226.80	430.60	477
	Total Points Scored	218.32	21.29	158	301	477
	Number of Scoring Events	111.98	11.83	85	153	477
Within-Game Characteristics						
	Total Viewership (1,000s)	2,683.29	2,460.62	265	20,956	4,771
	Score Differential	8.14	7.02	0	53	1,383,209
	Underdog Margin	-2.79	10.38	-53	38	1,383,209
	Consecutive Points	3.34	2.17	0	30	1,383,209
	Real-Time Diff. in Win Prob.	49.86	31.65	0	100	1,383,209

including individual time constraints preventing the viewing of a full game, or that the end of games typically feature more thrill. Another important insight from Figure 1.2 is that initial point spread appears to predict viewership, at least in the final stages of a game, suggesting that games expected to end with a close margin are, on average, more suspenseful and surprising.

Figure 1.1: Average Television Viewership Over Time by Score Differential Quintile

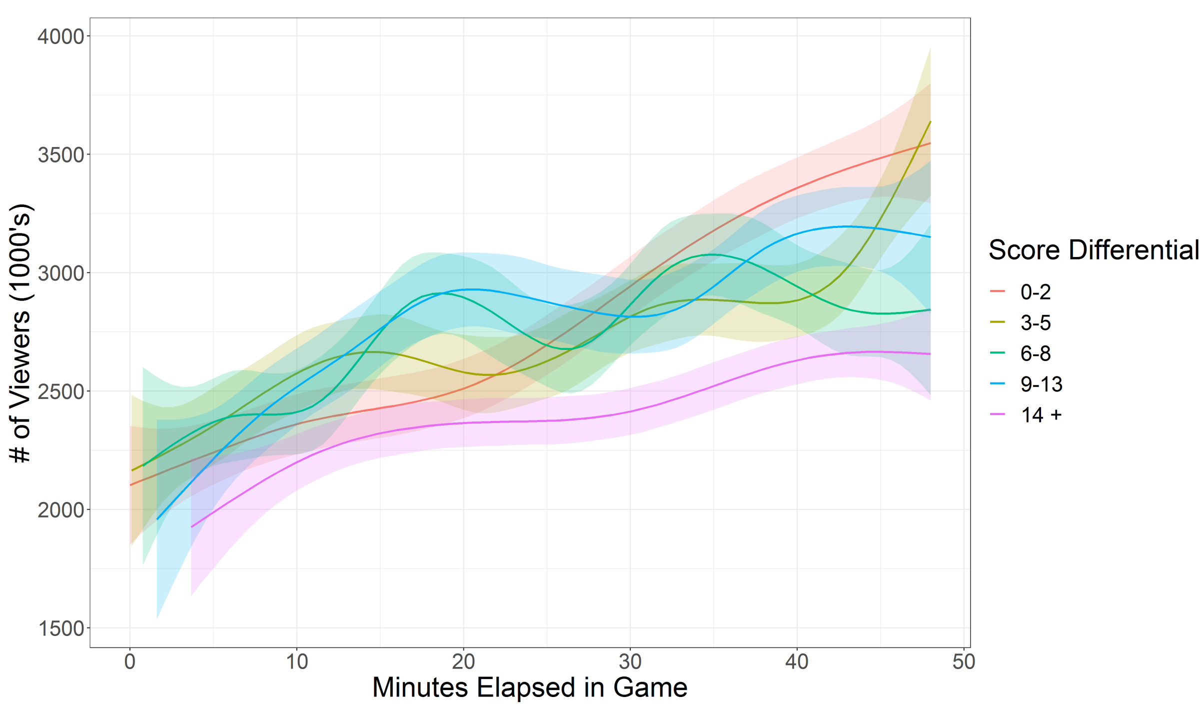

Figure 1.2: Average Television Viewership Over Time by Initial Point Spread Quintile

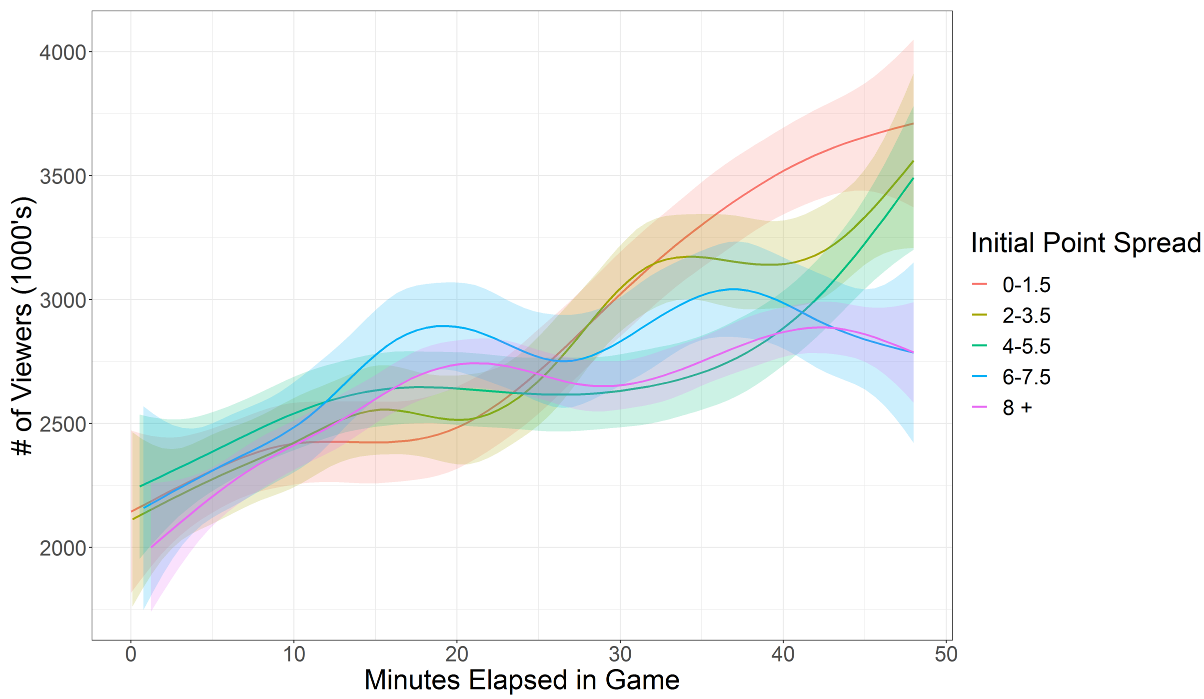

Figure 1.3 presents the correlation between the difference in real-time win probabilities of two competing teams and the absolute score differential, computed at the quarter segment level. While the definitions of suspense and surprise rely on beliefs, which are directly represented by the real-time outcome probabilities, there are a couple of key reasons score differential can also be useful in understanding viewership impacts from these characteristics. First, score differential is immediately observable to the viewer (unlike real-time win probability), and so it is likely the case that transparency of score differential is driving thrill-induced viewership responses. Second, because the television viewership data is observed over 15-minute intervals, it is difficult to pick up specific "spikes" in real-time win probability changes that are likely to occur in a thrilling game. Score differential is a smooth metric that still enables the capture of changes in suspense and surprise.

While the correlation is generally quite high (> 0.85), it is lower at the start and end of a game. This is intuitive—at the beginning of a game, absolute score differential is likely to be relatively low, yet real-time win probabilities remain heavily dependent on the initial expectation state over which team is likely to emerge victorious. On the other hand, at the end of a game win probabilities can fluctuate dramatically, even for small changes in the absolute score differential. This is precisely the effect I set out to measure–higher suspense (surprise) is associated with higher variance in future (past) beliefs about an outcome period-to-period.

Figure 1.3: Correlation between Difference in Real-Time Win Probability and Absolute Score Differential

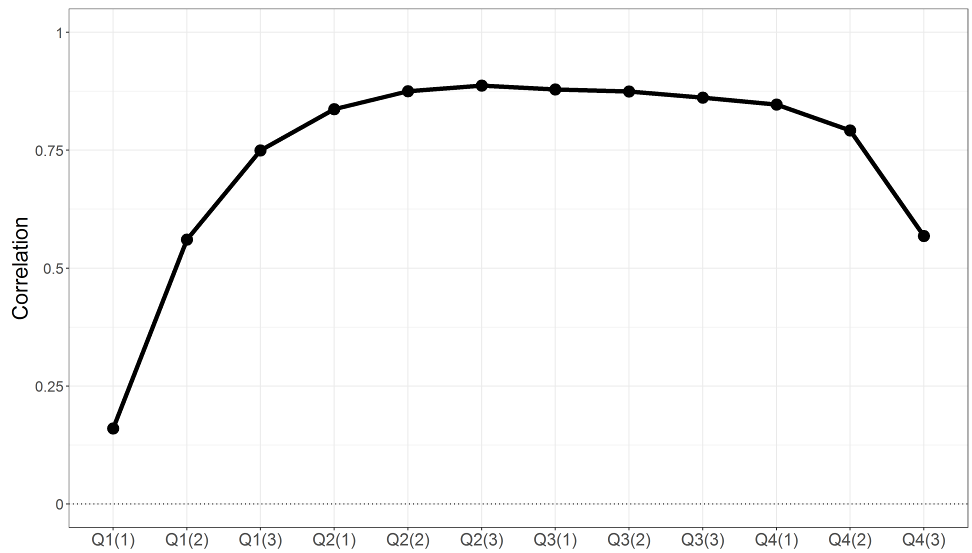

Finally, Figures 1.4 and 1.5 visually depict the nature of suspense and surprise over the course of a game. Each of the figures relies on the real-time win probability data provided at the second-of-game level, although these figures present the average suspense and surprise at the quarter segment level. Figure 1.4 examines the one-period forward-looking variance in real-time win probability differential between the two competing teams, which corresponds to suspense, while Figure 1.5 examines the one-period backward looking variance, corresponding to surprise. Figures 1.4 and 1.5 use a forward-(backward-)looking window of three minutes of game time (180 seconds). In subsection 3.5 presenting the results, I examine viewership impacts for one and three minute windows, but the findings are robust to reasonable window sizes.

In Figure 1.4, we see that games generally become more suspenseful as they progress, and then trail off slightly at the very end when outcome resolution begins to take place. Additionally, suspense is increasing at a faster rate and earlier on for close games. In terms of magnitude, a 0-2 point game experiences anywhere from 0-18 times more suspense than a 14 + point game, depending on the stage of the game.

Figure 1.5 depicts the trend associated with surprise. Games generally become more surprising as they progress, as is the case with suspense, but high score differential games exhibit greater surprise earlier in games. This is also intuitive—for there to be high variance in backward-looking win probability differential, there must be large leads occurring. Furthermore, while close games may not be surprising initially, they tend to feature more surprise in later stages, since even marginal score differential changes lead to relatively large

Figure 1.4: Variance in One-Period Forward-Looking Real-Time Win Probability Differential by Score Differential (Suspense)

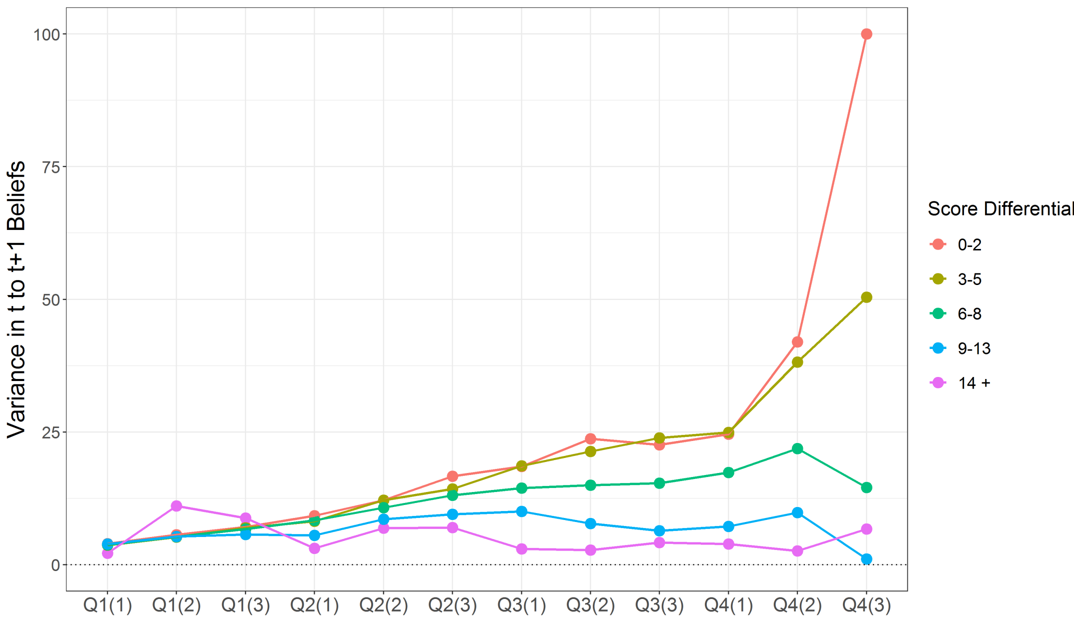

swings in real-time win probability differences.[12] Section 1.4 will present the viewership implications associated with suspense and surprise using both the observed absolute score differential and the structural variance parameters visualized in Figures 1.4 and 1.5.

Empirical Strategy

The empirical analysis in this section attempts to understand the impact of skill and thrill on television viewership. First, I develop a model to estimate *initial* viewership responses to skill and expected thrill. Second, I construct a model to estimate television viewership responses to "observable" thrill as it evolves within a game. Finally, I provide a framework to jointly estimate viewership responses to suspense, surprise, and skill within a game.

Initial Viewership Responses to Skill and Expected Thrill

The first component of the viewership analysis is to examine initial television viewership in response to skill and *expected* thrill. While there is full information before a game starts about the amount of skill that will be present (i.e. which players will be playing and their associated skill, measured by the number of All-Star fan votes they receive), thrill evolves

[12]While Figures 1.4 and 1.5 appear quite similar, the correlation between suspense and surprise at the second-of-game level is 0.26.

Figure 1.5: Variance in One-Period Backward-Looking Real-Time Win Probability Differential by Score Differential (Surprise)

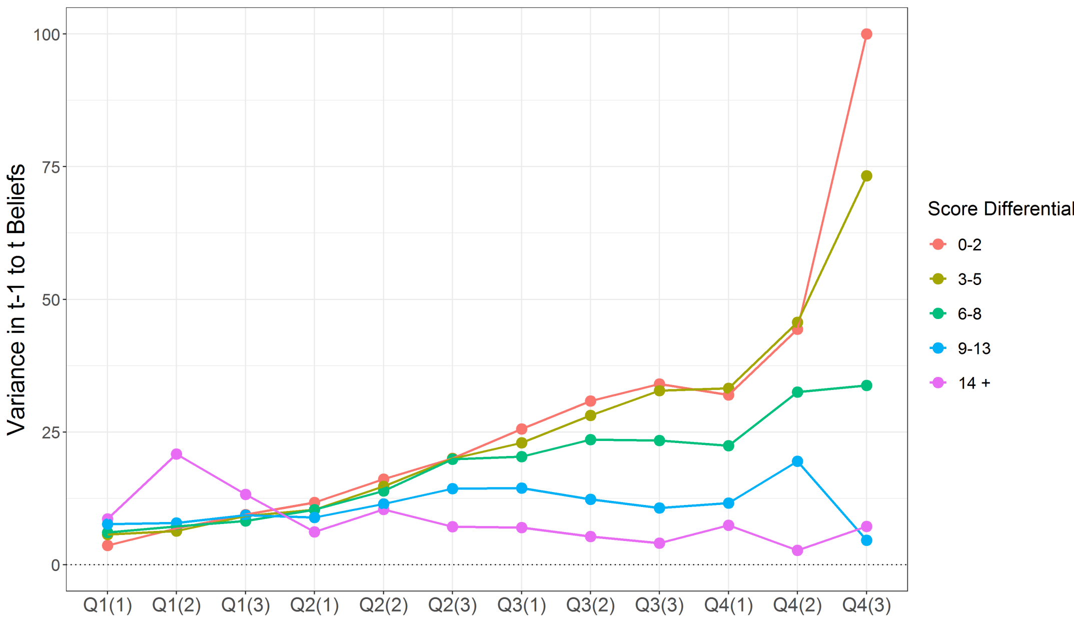

quasi-randomly during a game and so is not known beforehand. However, thrill may be correlated with the expected competitiveness of a game, where one may believe a more competitive game induces greater levels of thrill. The expected competitiveness of a game is measured by the initial point spread, and is directly observable prior to a game starting.

To measure expected thrill, I perform the following estimation, which resembles a two-stage least squares procedure.

$$\text{CumulativeThrill}_j = \delta \text{APS}_j + \lambda \text{Skill}_j + \mathbf{X}_j \Delta + \epsilon_{j,t=0} \tag{1.6}$$

$$\text{Viewership}_{j,t=0} = \gamma \widehat{\text{CumulativeThrill}}_j + \beta \text{Skill}_j + \mathbf{X}_j \Gamma + \epsilon_{j,t=0} \tag{1.7}$$

Where $\text{CumulativeThrill}_j$ is the summation of all instantaneous suspense and surprise that occurs during game j, $\widehat{\text{CumulativeThrill}}_j$ are the fitted values from the first-stage, and APS_j is the closing absolute point spread observed prior to a game. Note that the estimates are not meant to be interpreted as causal, rather as a descriptive relationship between initial point spread and cumulative thrill in the first stage, and the relationship between skill and expected thrill (as predicted by the initial point spread) on initial viewership in the second stage.

Observable Thrill

Measuring viewership responses to instantaneous suspense and surprise, which evolve during the course of a game, requires richer data and a different modeling strategy. Section 1.3 characterizes suspense and surprise in a structural way using the definitions from Ely et al. (2015), relying on outcome probabilities at a granular level that are not directly observed by spectators. However, I first analyze viewership responses to thrill using a directly observable game characteristic: absolute score differential at each point during a game. Absolute score differential is the primary metric by which a viewer internalizes thrill with respect to the final outcome of a game. While it is inherently difficult to separate the notions of suspense and surprise using this metric (since score differential at a given point can reflect both forward- and backward-looking beliefs), it provides an intuitive understanding of how viewership responds to thrill over the course of a game.

As implied by the definitions in Section 1.3, suspense and surprise are heavily dependent on time remaining in an event, since this impacts the extent to which beliefs can change across periods. Equation 1.8 provides a general empirical model to measure viewership impacts in response to observed absolute score differential and time remaining in an event.

$$V_{jt} = (C_{jt} * \mathbf{Q_{jt}})\mathbf{\Lambda} + \gamma S_{jt} + \alpha_j + \eta_t + \epsilon_{it} \tag{1.8}$$

V_{jt} represents total viewership for game j at time-of-game t. C_{jt} denotes the specific game characteristic impacting thrill (e.g. absolute score differential), and Q_{jt} is a time-of-game indicator (e.g. a minute of a game). $\mathbf{\Lambda}$ represents a vector of time-varying coefficients that reflect the impact of C_{jt} on viewership. S_{jt} represents the cumulative number of fan All-Star votes of all players on the court in game i at time t of the game (i.e. a measure of cumulative skill). It is critical to control for time-varying skill, since it is likely to be correlated with thrill during certain portions of a game. α_j and η_t represent game and quarter-segment fixed effects, respectively.

One important distinction to make is the difference between a close game and a thrilling game. A game featuring a low score differential in the first quarter would be characterized as close, but not thrilling, since the variance in beliefs about the outcome probabilities in the next period is low (suspense), and there was likely low variance in the evolution of beliefs prior to this point (surprise).[13] On the other hand, a low score differential in the fourth quarter would be considered both close and thrilling. Intuitively, the differential viewership impacts across the horizon of a game for similar score differentials is the variation used to separate the impact of thrill on viewership versus the impact of a close game.

An important assumption needed to interpret these estimates as plausibly causal, and the reason a live sporting event is a desirable setting to examine suspense and surprise, is *path-independence* of outcomes.

[13] See Figure 1.4 for a visual depiction of this.

Assumption 1: Path-Independence. The realized absolute score differential in period $t+1$, $|D_{t+1}|$, is random conditional on the score differential at time, $|D_t|$, and fixed information known prior to a game, μ_0.

$$|D_{t+1}| \sim \mathcal{N}(|D_t|, \sigma^2 \mid \mu_0) \tag{1.9}$$

This assumption states that the absolute score differential evolves randomly, conditional on the score differential in the previous time period and fixed information known prior to a game that may impact the evolution of the score differential (e.g. the initial point spread). Essentially, the evolution of the absolute score differential is a first-order Markov process, accounting for dependence on the initial state of the game μ_0.

Observable Thrill over Alternative Outcomes

Individuals may also experience suspense or surprise with respect to an outcome unrelated to which team wins the game. Examples include which team covers the point spread, total points scored over/unders, and other within-matchup propositions. In order to make the analogy to absolute score differential, I assume that an agent who cares about these outcomes maintains the same utility function from thrill as seen in equation 1.3.[14] Here, the alternative outcome I will examine is the point spread set before a game begins, which is one of the most common measures gambled on by bettors. In this case, it is not the absolute score differential that determines thrill, rather the absolute score differential *in reference to the point spread.*

The point spread is defined as the number of points P_{jT} such that $V_{jA} + F(P_{jT}) = V_{jB}$, where $F(\cdot)$ is a one-to-one function mapping points to strength.[15] I index by T since point spreads typically refer to $\mathbb{E}[D_T]$. Using this setup, the absolute score differential in reference to the closing point spread can be defined:

$$|D'_{jt}| = |D_{jt} + P_{jT}| \tag{1.10}$$

where both D_{jt} and P_{jT} use the same team as the reference point for scoring. For instance, if the home team is always used as the reference point, $D_{jt} > 0$ implies the home team is leading, and $P_{jT} > 0$ implies the home team is an underdog. To understand the application of this outcome empirically, take the following concrete example. Suppose there is a matchup featuring the Cleveland Cavaliers and Boston Celtics, where the Cavaliers are the home team. If the closing point spread was -7, and the score at the end of the third quarter was 85 - 82 favoring Cleveland, then the absolute score differential from the spread would be equal to four. However, if the score was 85 - 82 in favor of Boston, the absolute score differential from the spread would be equal to ten.

[14] This may be a strong assumption if individuals that care about these outcomes have an explicit financial stake, and thus suspense is endogenously chosen.

[15] Note that I index strength here at the matchup level, allowing for strength for a specific team to differ across matchups.

To measure thrill from this outcome, I rely on the methodology used in Salaga and Tainsky (2015), who study television viewership for all PAC-12 football games from 2009-15. They examine the impact of score differential during a game in reference to the closing point spread on average television viewership for a game (they do not measure viewership changes over time within games). The authors note that it is important to de-confound estimates from viewership corresponding to the actual game outcome, represented by the raw score differential. To try and account for this, the authors subset their analysis sample to i) the second half of games, ii) games with the absolute score differential above some threshold level $G_{t=0.5T}$ at halftime, and iii) games whose absolute score differential does not fall below some threshold $G_{t>0.5T}$ during the second half of a game.

One important difference in my approach is that I use real-time win probability estimates for each game instead of absolute score differential to determine the subsample to study. This is because a uniform score differential threshold may correspond to significantly different win probabilities in different games. I set $G_{t=0.5T} = 0.6$ and $G_{t>0.5T} = 0.4$. Games must meet the criteria where at halftime, the difference in win probabilities of each team winning is ≥ 0.6, and over the course of the second half, that difference does not fall below 0.4. The results are not sensitive to restrictions reasonably close to these bounds.

Applying this approach, I estimate a model of viewership in response to suspense over the absolute score differential in reference to point spread as follows:

$$V_{jt} = (|D'_{jt}| * \mathbf{Q_{jt}})\mathbf{\Lambda} + (|D_{jt}| * \mathbf{Q_{jt}})\mathbf{\Gamma} + \alpha_j + \eta_t + \epsilon_{jt} \tag{1.11}$$
$$\text{s.t.} \quad |P_{t=halftime}(A) - P_{t=halftime}(B)| > G_{t=0.5T} \quad \& \quad |P_{t>halftime}(A) - P_{t>halftime}(B)| > G_{t>0.5T}$$

where all the terms maintain their previous definitions.

Joint Model of Suspense, Surprise, and Skill

This section presents an empirical model to jointly estimate the impacts of skill and thrill on viewership. I rely on the structural definitions of suspense and surprise given in Section 1.3. The general form of the estimating equation is as follows:

$$V_{jt} = \mu X_{jt} + \rho Y_{jt} + \psi S_{jt} + \lambda(X_{jt} * S_{jt}) + \nu(Y_{jt} * S_{jt}) + \mathbf{Z_{jt}\Gamma} + \alpha_j + \eta_t + \epsilon_{jt} \tag{1.12}$$

V_{jt} represents total viewership for game j at time-of-game t. X_{jt} denotes the structurally defined suspense parameter, Y_{jt} the structurally defined surprise parameter, and S_{jt} a cumulative measure of skill.[16] $\mathbf{Z_{jt}}$ includes a set of controls that evolve within-game, and matchup and time-of-game fixed effects are denoted as α_j and η_t, respectively.

[16] Using the definitions of suspense and surprise presented in section 1.3, it is necessary to define the length of the period-to-period interval in which they can occur. I use several different bandwidths in the estimations presented in section 1.4, including a 1-minute of play time window (i.e. 60 seconds of game clock time, not real time), 3-minute, and 5-minute window.

There are several advantages of this estimation approach. First, it allows for suspense and surprise to be included as separate terms in a single estimation so their impacts on viewership can be separately identified. Second, it allows for the inclusion of a time-variant measure of observable skill, since it relies on within-game variation in the cumulative skill of all players playing at a given point in a game. Finally, it allows for the inclusion of interaction terms between skill and suspense, and skill and surprise, which are useful in understanding differential viewership impacts to thrill depending on the presence of skill. The following section will present the results from each of the empirical models discussed here.

Results

This section presents estimation results from the empirical models of skill and thrill posed in the previous section.

Initial Television Viewership

Table 1.2 presents the impact of skill, as measured by cumulative All-Star fan votes for all players playing in a game, and expected cumulative thrill on initial TV ratings for nationally-televised games. Cumulative thrill is the sum of instantaneous suspense and surprise, as defined in Section 1.3, over the course of an entire game. Column (1) presents the first-stage estimation examining the relationship between initial point spread and cumulative thrill. Column (2) presents the second-stage estimation, examining initial viewership in response to skill and expected thrill, where expected thrill only relies on variation from the initial point spread. Each of these specifications controls for "combined current win percentage" to account for the average quality of the two teams playing in a game, as well as an "aggregate team value" continuous control variable to account for the number of people that may be expected to watch a specific team independent of other important factors.[17] These team values are calculated each year by Forbes, and are a good indicator of the total size of each team's fanbase (Badenhausen and Ozanian 2019). Other important controls are listed at the bottom of the table, which include month-of-year, day-of-week, and time-of-day fixed effects.

Intuitively, column (1) shows there is a negative and statistically significant relationship between the initial point spread and cumulative thrill, suggesting that for games featuring higher initial absolute point spreads (i.e. games that are expected to be less competitive), we observe less total thrill. The effect is sizable – for every one point increase in the absolute point spread, cumulative thrill falls by approximately 2.18%.[18] However, column (2) suggests that initial viewership is not affected by the expected thrill of a game (or the initial point

[17]Since these are nationally televised games, a home team fixed effect does not make as much sense in these specifications as it does in the context of the ticket price analysis (since there are geographic preferences). Including a dummy for each team present in a matchup leads to insignificant point estimates for all variables, likely because of the insufficient power associated due to the relatively low number of nationally-televised games.

[18]The distribution of absolute point spreads within the sample of data is presented in Table 1.1.

Table 1.2: Impact of Skill and Expected Thrill on Initial TV Ratings

	Dep. Var: log(Cumulative Thrill)	Dep. Var: log(1000's of Initial Viewers)
Abs. Point Spread	-0.0218^{***}	
	(0.0082)	
$\log(\widehat{\text{Cum. Thrill}})$		0.0580
		(0.2416)
log(Ag. All-Star Votes)	0.0835^{*}	0.1124^{***}
	(0.0462)	(0.0362)
log(Avg. Current Win PCT)	-0.4468^{*}	0.2822^{*}
	(0.2677)	(0.1607)
log(Ag. Team Value)	-0.1354^{**}	0.1830^{**}
	(0.0644)	(0.0891)
Month FE	Yes	Yes
Day-of-Week FE	Yes	Yes
Time-of-Day FE	Yes	Yes
Streak FE	Yes	Yes
TV Network FE	Yes	Yes
Dbl Header FE	Yes	Yes
Holiday FE	Yes	Yes
Playoff Gm FE	Yes	Yes
Clustered Robust SEs (Home + Away)	Yes	Yes
Observations	477	477
R^2	0.1386	0.7441
Adjusted R^2	0.0420	0.7154
Note:	*p<0.1; **p<0.05; ***p<0.01	

spread), suggesting that individuals are no more likely to tune in for the start of a game that features a higher expected level of thrill. On the other hand, the skill present in a game significantly impacts initial viewership. For a 100% increase in the skill of players present, initial viewership increases by approximately 11.2%. One can also see that "combined current win percentage" and "aggregate team value" have larger impacts on initial viewership than skill and expected thrill. This is not surprising considering that the quality of the two competing teams and their market sizes are likely to be primary determinants of initial viewership.

It is clear that skill is a significant driver causing individuals to tune into games. For context, LeBron James obtained over 4.6 million fan votes during the 2018-19 season, which corresponds to approximately 120% of the average aggregate number of All-Star fan votes of all players in a matchup (3.8 million). In other words, LeBron's average fan All-Star vote total is just above the total number of All-Star votes of all players in an average game. Using the results from analysis in Table 1.2, the presence of LeBron alone results in an approximately 13.5% increase in initial TV ratings. These results are remarkably similar to those found in Kaplan (2020), which uses secondary ticket marketplace data to assesses the impact of a superstar absence announcement for a specific game on listed prices. The analysis finds that the absence of LeBron James leads to a 13% average reduction in ticket prices.

Observable Thrill

The primary observable characteristic of thrill in these matchups is the absolute score differential in matchup i at time t, D_{it}. Table 1.3 shows four separate estimations. Columns (1) and (3) present the "naive" estimations, namely the average impact of absolute score differential on log viewership. This specification is meant to capture viewership in response to the close game effect, which can be measured uniformly over a game (i.e. a 2-point game in the first quarter is just as close as a 2-point game in the fourth quarter).[19] Column (3) differs from column (1) only in that it controls for time-varying skill. Columns (2) and (4) present the time-varying impacts of absolute score differential on viewership, which corresponds to equation 1.8. Again, column (4) differs from column (2) in that it controls for time-varying skill.

Table 1.3: Impact of Absolute Score Differential on TV Ratings

	Dependent Variable: log(Total Proj. Viewers Watching)			
Absolute Score Diff.	-0.0056^{***}	0.0014	-0.0047^{***}	0.0017
	(0.0013)	(0.0018)	(0.0009)	(0.0018)
Absolute Score Diff. * Q2		-0.0012		-0.0015
		(0.0018)		(0.0018)
Absolute Score Diff. * Q3		-0.0059^{**}		-0.0060^{**}
		(0.0024)		(0.0024)
Absolute Score Diff. * Q4		-0.0115^{***}		-0.0105^{***}
		(0.0025)		(0.0023)
log(Ag. All-Star Votes)			0.0324^{***}	0.0210^{**}
			(0.0093)	(0.0072)
Game FE	Yes	Yes	Yes	Yes
Quarter Segment FE	Yes	Yes	Yes	Yes
Observations	1,384,623	1,384,623	1,382,923	1,382,923
R^2	0.9450	0.9466	0.9464	0.9476
Adjusted R^2	0.9450	0.9466	0.9464	0.9475
Note:	*p<0.1; **p<0.05; ***p<0.01			

[19] As mentioned previously, it is important not to conflate the effect of a close game versus suspense and surprise on viewership.

One can see that on average across an entire game, a one point increase in the absolute score differential reduces television viewership by 0.47-0.55%, and so close games are important in raising viewership. Columns (2) and (4) break out the impacts of absolute score differential by quarter of the game. There is a clear relationship between time remaining in the game and the impact of score differential on viewership – a one point increase in absolute score differential in the fourth quarter leads to an approximately 1.1-1.2% drop in viewership, compared to a drop in the first two quarters that is not significantly different from zero. This is strong evidence in support of the impact of thrill on viewership – marginal score differential changes lead to higher viewership impacts when they lead to a larger variance in beliefs, either forward- or backward-looking. As shown in Table 1.1, which depicts summary statistics of the play-by-play data, the mean and standard deviation of absolute score differential are 8.14 and 7.02, respectively, suggesting that viewership responses are quite sensitive to changes in absolute score differential. Examining both columns (3) and (4), it is clear that skill impacts viewership within a game, albeit to a lesser extent than thrill. A doubling (one SD) increase in skill on the court during a game increases viewership by approximately 2.1-3.2% (1.2-1.9%).

Figure 1.6: Household Viewership Results by Score Differential Bin by Quarter Segment (% Change)

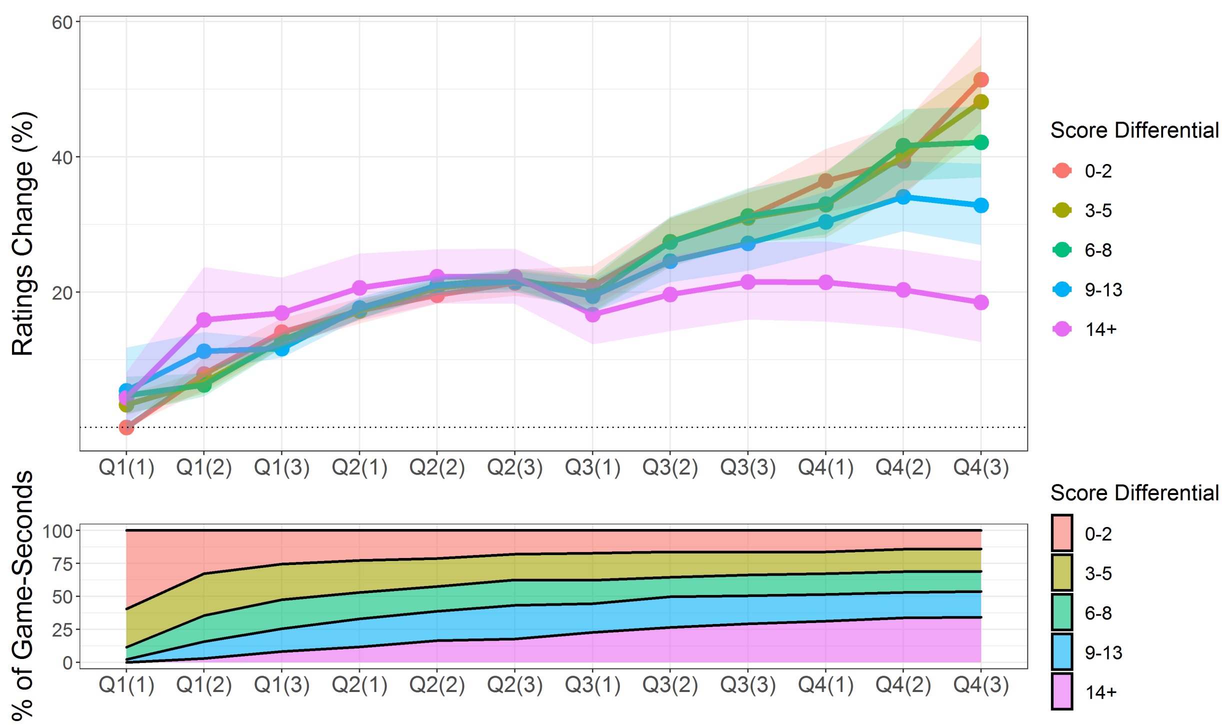

Figure 1.6 depicts thrill impacts using absolute score differential at an even more granular level. I split each game into twelve equally long quarter segments, and absolute score differential is divided into five bins using the quintiles of the distribution of score differential in the data. All points in Figure 1.6 represent coefficients from an estimation taking the form of equation 1.8, and can be interpreted as relative to the omitted score differential bin-by-quarter segment (the 0-2 bin in the first quarter segment, Q1(1)). First, this graph confirms that average viewership over the course of a game is generally increasing, as shown in Figures 1.1 and 1.2. It is also clear that there are heterogeneous impacts of absolute score differential on viewership as a game progresses to its later stages. While in the first half there are no significant differences between each of the score differential bins and viewership changes, in the second half viewership flattens out for the higher score differential bins compared to the lower bins. In particular, a game in the closest absolute score differential quintile (0-2 points) features 8.2-20.5% lower viewership in the second half compared to a game in the largest absolute score differential quintile (14 + points), with the difference increasing monotonically as a game approaches its finality.

It is clear that the 14 + absolute score differential bin exhibits the most stark impacts on viewership. Figure 1.7 examines these effects more closely, looking at the tails of the distribution of absolute score differential. Here, impacts appear to be much more sensitive than those in the primary support of the score differential distribution, where marginal increases in absolute score differential when the differential is already quite high are much more impactful on viewership than marginal increases when the differential is quite low. This may suggest a non-linear response to thrill during a game. Estimations using alternative binning structures as well as level (instead of log) changes in viewership on are presented in the Appendix.

Observable Thrill in Reference to the Point Spread

Table 1.4 presents results depicting the effect of absolute score differential in reference to the closing point spread on viewership. Columns (1) and (2) present results of the naive estimation, which measures average viewership impacts associated with games close-to versus far-from the initial point spread, while columns (3) and (4) show thrill-driven impacts. Additionally, columns (2) and (3) control for the average impact of the raw absolute score differential on viewership, while column (4) controls for differential impacts of the raw absolute score differential on viewership by time of game.

In the naive model, the hypothesized sign of the coefficient on absolute score differential from the spread is negative, namely the further the absolute score differential gets from the point spread, the lower viewership becomes. One can see from columns (1) and (2) that controlling for absolute score differential is important, since it is likely correlated with absolute score differential from the point spread and also has negative impacts on viewership. Column (2) suggests there are no statistically significant viewership impacts associated with a close game in reference to the spread in a national audience. Columns (3) and (4) provide the thrill-driven impacts of score differential from the spread on viewership. While there does

Figure 1.7: Household Viewership Results by Score Differential Bin by Quarter Segment (Tails)

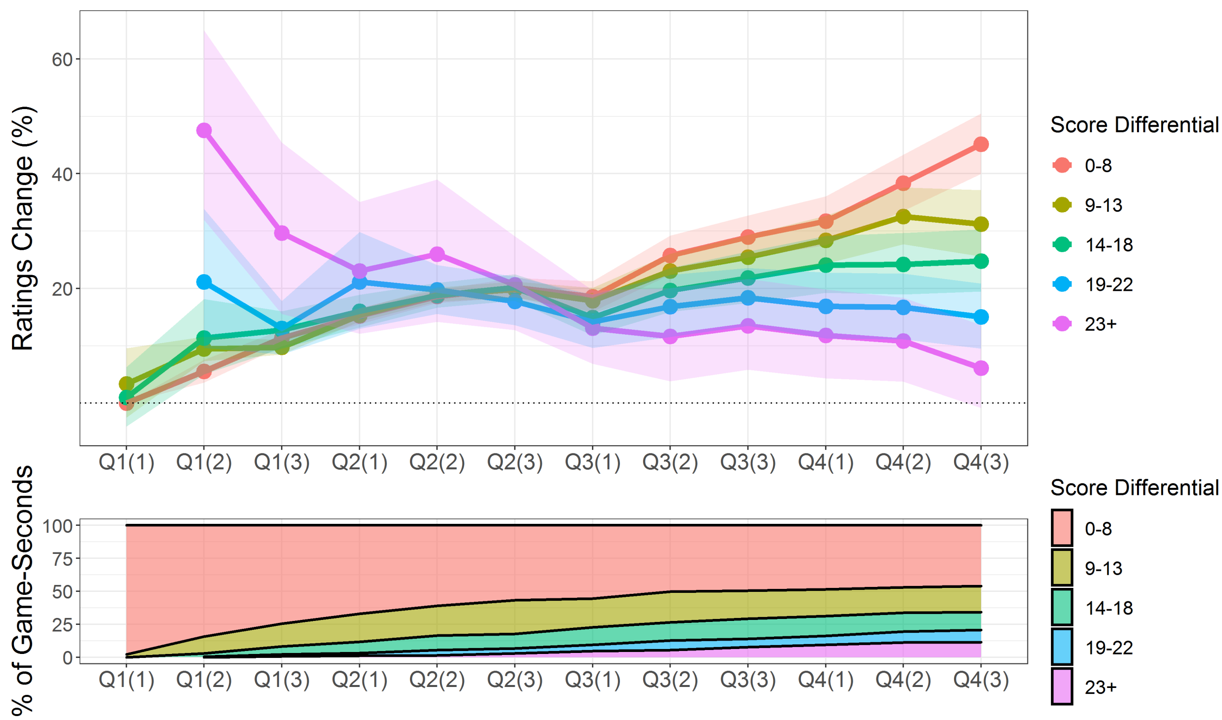

not appear to be significantly different effects from zero until the end of the third quarter, it is clear that as the game progresses, a higher absolute score differential from the spread leads to larger decreases in viewership. This is result has an identical explanation to the results found in Table 1.3 and Figures 1.6 and 1.7. Since the omitted period is Q3(1), the true effect of the score differential in reference to the point spread on viewership in the final quarter segment is -0.0079 in specification (3) and -0.0069 in specification (4), suggesting that for every one-point increase in the score differential from the spread, viewership declines by approximately 0.79% and 0.69%, respectively. As expected, these results are approximately half the magnitude of the impact of raw absolute score differential on viewership. However, given these estimates, a one-standard deviation change in score differential in reference to the initial point spread during the final quarter segment (9.3 points) can still have an economically meaningful impact on viewership (6.4-7.3% reduction).

Joint Estimation of Suspense, Surprise, and Skill

The next set of estimations assesses the joint impact of suspense, surprise, and skill on television viewership. This analysis relies on the structurally defined suspense and surprise parameters, as well as within-game variation in the level of skill on the court, providing

Table 1.4: Impact of Thrill with Respect to Point Spread on TV Ratings

	Dependent Variable: log(Total Proj. Viewers Watching)			
Absolute Score Diff. From Spread	-0.0035^{*}	-0.0041	0.0021	0.0025
	(0.0016)	(0.0024)	(0.0021)	(0.0020)
Absolute Score Diff. From Spread * Q3(2)			-0.0003	-0.0010^{**}
			(0.0002)	(0.0004)
Absolute Score Diff. From Spread * Q3(3)			-0.0022^{**}	-0.0023^{**}
			(0.0006)	(0.0007)
Absolute Score Diff. From Spread * Q4(1)			-0.0031^{**}	-0.0021
			(0.0010)	(0.0014)
Absolute Score Diff. From Spread * Q4(2)			-0.0061^{***}	-0.0055^{**}
			(0.0012)	(0.0015)
Absolute Score Diff. From Spread * Q4(3)			-0.0100^{***}	-0.0094^{***}
			(0.0013)	(0.0018)
Score Differential Control	No	Yes	Yes	Yes
Score Differential x Quarter Segment Control	No	No	No	Yes
Game FE	Yes	Yes	Yes	Yes
Quarter Segment FE	Yes	Yes	Yes	Yes
Observations	40,588	40,588	40,588	40,588
R^2	0.9821	0.9821	0.9857	0.9859
Adjusted R^2	0.9821	0.9821	0.9857	0.9859
Note:				*p<0.1; **p<0.05; ***p<0.01

high-temporal frequency changes that can be separated from time invariant, game-specific factors and general viewership trends over the course of a game. Table 1.5 presents the results of six separate estimations: columns (1) - (3) use a one-minute forward-(backward-)looking window to calculate suspense (surprise) at each second-of-play during a matchup, while columns (4) - (6) use a three-minute window.[20] Columns (2) - (3) and (5) - (6) also control for the average impact of absolute score differential during a game on viewership, so as to account for potential correlation between a suspenseful or surprising game and a "close" game. Finally, columns (3) and (6) include the interactive effect of suspense and surprise with skill with the goal of measuring differential viewership responses to thrill under varying levels of skill.

There are several notable takeaways from Table 1.5. First, the estimates suggest that a doubling of suspense increases viewership by 0.38-0.60%. However, as exhibited in Figure

[20] The results are not sensitive to reasonable window size adjustments.

Table 1.5: Impact of Suspense, Surprise, and Starpower on TV Ratings

	Dependent Variable: log(1000's of Total Viewers)					
	1 Minute Window			3 Minute Window		
log(Surprise)	0.0075***	0.0064***	0.0074	0.0096***	0.0086***	−0.0178
	(0.0011)	(0.0010)	(0.0146)	(0.0020)	(0.0019)	(0.0224)
log(Suspense)	0.0056***	0.0038**	0.0424**	0.0060***	0.0039***	0.0462**
	(0.0015)	(0.0013)	(0.0153)	(0.0012)	(0.0011)	(0.0184)
log(All-Star Votes)	0.0236**	0.0237**	0.0203***	0.0200**	0.0202**	0.0187***
	(0.0082)	(0.0078)	(0.0052)	(0.0073)	(0.0070)	(0.0051)
Absolute Score Diff.		−0.0021**	−0.0023**		−0.0022**	−0.0024**
		(0.0010)	(0.0010)		(0.0010)	(0.0010)
log(Surprise) * log(All-Star Votes)			−0.0002			0.0017
			(0.0010)			(0.0014)
log(Suspense) * log(All-Star Votes)			−0.0026**			−0.0029**
			(0.0010)			(0.0012)
Game FE	Yes	Yes	Yes	Yes	Yes	Yes
Quarter Segment FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,381,973	1,381,973	1,381,973	1,381,973	1,381,973	1,381,973
R^2	0.9470	0.9471	0.9474	0.9471	0.9473	0.9474
Adjusted R^2	0.9470	0.9471	0.9473	0.9471	0.9473	0.9474
Note:	*p<0.1; **p<0.05; ***p<0.01					

1.4, suspense can take on an extremely large range of values. For instance, in the last segment of the fourth quarter, a 0-2 point game averages 18 times more suspense than a 14+ point game. In this case, viewership would be approximately 6.84-10.80% higher through suspense alone. On the other hand, the magnitude of the surprise effect depends greatly on the specification, particularly when comparing the fully interacted estimations (columns 3 and 6) with the other specifications. In the non-fully interacted estimations, a doubling of surprise increases viewership 0.64-0.96%. In the fourth quarter, a 0-2 point game features 14 times more surprise on average than a 14+ point game, which would translate into a viewership increase of 8.96-13.44%. However, specifications (3) and (6) exhibit largely different impacts of surprise on viewership. In particular, when interacting suspense and surprise with skill, the impact of surprise on viewership is no longer statistically significant. It is fairly intuitive that for a sport like basketball, which features frequent scoring and smooth updating in outcome probabilities, surprise would have a lower impact on viewership.

There are a couple of interesting and intuitive takeaways regarding the impact of skill on

Table 1.6: Within-Game Superstar Viewership Impacts

Player	All-Star Fan Votes	Avg. Total Votes when Off Court	Viewership Impact (%)
LeBron James	4,620,809	3,310,401	2.83
Giannis Antetokounmpo	4,375,747	2,652,937	3.35
Luka Doncic	4,242,980	1,286,935	6.94
Kyrie Irving	3,881,766	2,711,703	2.91
Stephen Curry	3,861,038	4,680,438	1.67
Kawhi Leonard	3,580,531	3,388,249	2.15
Derrick Rose	3,376,277	2,742,202	2.50
Paul George	3,122,346	3,505,377	1.81
Kevin Durant	3,150,648	5,280,465	1.21
James Harden	2,905,488	2,578,477	2.29
Joel Embiid	2,783,833	3,155,921	1.79
Anthony Davis	2,520,728	1,893,613	2.70
Dwyane Wade	2,208,598	1,833,541	2.45
Kemba Walker	1,395,330	1,190,191	2.38
Dirk Nowitzki	394,622	2,463,873	0.33

viewership. First, examining the average impact of skill while holding suspense and surprise constant, all specifications suggest that a doubling in the number of All-Star fan votes on the court at a given time during the game leads to a 1.87-2.37% increase in viewership. While these estimates are substantially lower than those found in the initial TV viewership analysis, the source of variation, and therefore interpretation of the coefficient magnitudes, is different. While the aforementioned analyses look at the time invariant impact of skill on viewership, this estimation relies on within-game changes in the level of skill on the court at any given time, and thus they complement one another in interpreting the impact of skill on spectator demand. One can think of the skill-induced viewership changes estimated in Table 1.5 as occurring on the intensive margin (within games), while the larger estimates found in the initial TV viewership analysis as occurring on the extensive margin (across games). It may be the case that spectators face two different decisions with respect to skill: the likelihood of watching a game *at all* because of the aggregate skill of all players playing, and whether to continue watching a game when it features changes in skill on the court at a given time.

Table 1.6 translates the coefficient on skill from specification (3) to the corresponding within-game viewership impact for some of the most skilled players. I compare each individual superstar's total All-Star fan vote tally to the average cumulative number of All-Star fan votes when each player is *off* the court. For instance, LeBron James received 4.6 million All-Star fan votes in the 2018-19 season. The total average number of All-Star fan votes on the court in games where he is playing but not on the court is 3.3 million. Thus, LeBron's average skill impact translates to a 140% increase in on-court popularity, increasing within-game viewership by 2.83% when he is on the court playing.

The second important takeaway from Table 1.5 related to skill are the results from the fully interacted estimations in columns (3) and (6). Most interesting is the relationship between skill and suspense. While suspense (holding skill constant) continues to have a significant and meaningful impact on viewership, the interactive effect between skill and suspense is negative and of smaller magnitude importance than the individual coefficients on suspense and All-Star votes. In words, suspense leads to differentially *higher* viewership with *lower* skill on the court. While an individual may watch a game featuring LeBron James no matter the level of suspense, they may exhibit a more sensitive and heightened response to suspense in games featuring less skill. This supports and quantifies the traditional idea that spectators only turn on games featuring lesser-known players if they're nearing the end and exhibiting sufficiently high suspense.

1.5 Discussion

This section discusses important implications of the empirical findings. First, I contextualize the results from the television viewership analyses by comparing and contrasting their effects with other important drivers of demand for NBA games. In doing this, I also provide revenue implications for the NBA associated with demand for skill and thrill. Next, I propose a counterfactual game structure that introduces more "finality" to an event, which would lead to larger levels of thrill over the course of a game, and assess the viewership implications associated with its implementation.

Effect Sizes and Revenue Implications

The empirical analysis assessing skill and thrill finds effect sizes on viewership demand for skill of approximately 11%, and viewership demand for thrill between 7-30%, depending on specification. In particular, an increase in viewership of 7-30% corresponds to viewership increases of 187,830 - 804,987 individuals during a 15-minute programming interval, and the 11% initial viewership increase associated with a doubling of skill corresponds to approximately 295,162 additional viewers. For comparison, a one standard deviation increase in 15-minute level viewership is approximately 92% of 15-minute interval mean viewership (2.46 million individuals). Playoff games experience approximately 93% higher viewership than regular season games, and holiday games experience 54% higher viewership than non-holiday games. Additionally, viewership increases by approximately 45% on average from the start to the end of a game.

Examining the effect of other characteristics on television viewership provides further evidence that both skill and thrill are highly important economic factors in driving demand for NBA games. Next, I assess league revenue implications associated with skill and thrill. Table 2.7 presents the projected season-level value in both ticket sales and television viewership settings for the highest skill NBA players. The impacts in column 3 are based on results

Table 1.7: Season-Level Ticket Price and TV Viewership Player Impacts

Player	All-Star Votes	WTP Impact (Millions of $)	Initial WTW Impact (Millions of $)
LeBron James	4,620,809	69.06	24.51
Giannis Antetokounmpo	4,375,747	-1.57	23.21
Luka Doncic	4,242,980	30.13	22.51
Kyrie Irving	3,881,766	8.09	20.59
Stephen Curry	3,861,038	48.18	20.48
Kawhi Leonard	3,580,531	10.76	19.00
Derrick Rose	3,376,277	2.09	17.91
Kevin Durant	3,150,648	7.50	16.72
Paul George	3,122,346	29.03	16.56
James Harden	2,905,488	12.98	15.41
Joel Embiid	2,783,833	13.99	14.77
Anthony Davis	2,520,728	5.73	13.37
Dwyane Wade	2,208,598	41.86	11.72
Kemba Walker	1,395,330	23.74	7.40
Dirk Nowitzki	394,622	32.19	2.09

Note: These estimates represent the season-level monetary impacts each player had based on the difference-in-differences (DID) estimations in Kaplan (2020) for secondary marketplace ticket data (column 3) and initial viewership estimations for the TV viewership data (column 4). For the DID estimates from Kaplan (2020), this meant multiplying by 20,000 people on average per arena and 82 games over the course of a season. The initial viewership estimates were estimated via a log-log specification. So, the season-level impacts were determined using the player-specific % total of the average cumulative number of All-Star votes present for all players in a specific game (3.8 million), the approximate total value of television broadcasting for the NBA during a season (∼\$2.7 billion; Sports-Illustrated 2014), and the total number of regular season games each team plays (82).

presented in Kaplan (2020), which uses a difference-in-differences approach to examine the ticket price reductions on a secondary marketplace associated with superstar player absence announcements for specific games, extrapolated over an entire NBA season. The impacts in column 4 are using the results from Table 1.2, again extrapolated over an entire season. One can see that from ticket sales alone, the impacts associated with the presence of superstars range from millions to tens of millions of dollars over the course of a season, with LeBron James leading at \$69 million. The viewership impacts are slightly smaller on average, but still on the order of millions to tens of millions of dollars over the course of a season, with a maximum player value (which once again corresponds to LeBron James) of \$24.5 million. The differences in magnitudes between ticket sales and television ratings may be due to a number of factors, but one likely contributor is the heightened allure of skill when watching in person.

To assess estimated revenues from thrill, I use changes in advertising revenues associated with differences in thrill across games. Figure 1.8 presents viewership change estimates associated with observable thrill, and is identical to Figure 1.6. Assuming a cost-per-thousand (CPT) viewership minutes estimate of \$25 (Fou 2014; Friedman 2017), 20% of programming time during a game spent on advertisements (Statista 2014), and a 15-minute average length of each of the 12 quarter segments, the difference in advertising revenue between a 0-2 point game versus a 14+ point game in the final quarter segment is approximately $50,000.[21] Aggregating this difference over the course of an entire second-half, I find that advertising revenues are $130,000 higher for 0-2 point games compared to 14+ point games. While these revenue differences are economically sizeable, they are likely to underestimate the true welfare associated with thrill since they do not account for increases in consumer surplus of inframarginal viewers due to enhanced thrill.

Figure 1.8: Estimated Advertising Revenues from Thrill

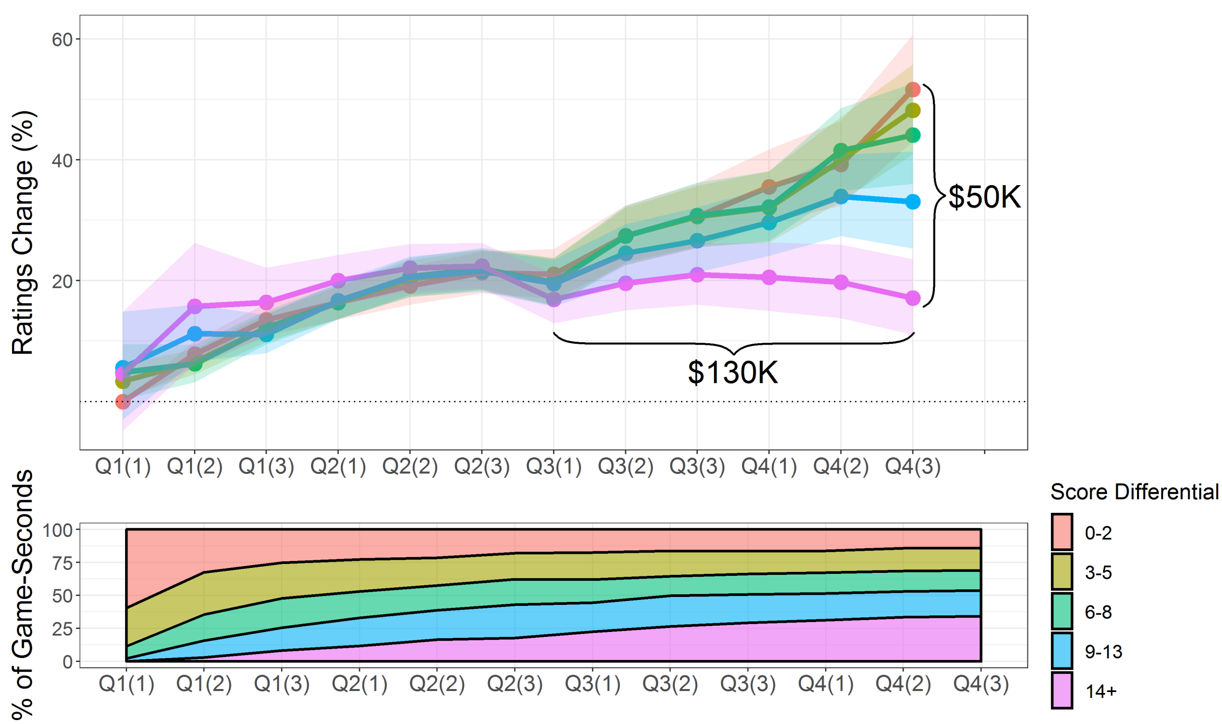

Counterfactual Game Structures

Sports leagues are always discussing and considering different measures and rule changes that have the potential to enhance the fan experience and increase market size. Understanding how counterfactual structures may affect viewership is important to better understand the economic implications of proposed adjustments. In this section, I present and analyze a counterfactual game structure that introduces more "finality" into a game. Examining Table

[21]The five absolute score differential bins presented represent the quintiles of the distribution within the data. Thus, approximately 20% of game-seconds within a game experience a 0-2 point score differential, and 20% of game-seconds experience a 14+ point score differential.

1.3 and Figure 1.6, it is clear that much of the thrill that takes place during a game occurs towards the end. This is quite intuitive – on average, the impending final outcome of an event generates larger swings in the outcome probability than earlier stages of a game. Figure 1.9 shows this explicitly – in the left-pane, I plot the average, 75th, and 95th percentiles of thrill (suspense + surprise) within each quarter segment. It is very apparent that thrill is increasing monotonically over the course of a game at these points in the thrill distribution.[22] The right-pane presents a visual depiction of a scenario enhancing the finality of an event. For instance, suppose instead of a single meaningful outcome in a game (i.e. whichever team wins and loses the game), each quarter of a game represents a meaningful outcome. Under this scenario, each game would include four meaningful outcomes. The right-pane of Figure 1.9 portrays the distribution of thrill over the course of a game *extrapolated from the fourth quarter distribution of thrill*. I take the distribution of thrill from the fourth and final quarter presented in the left-pane, and extrapolate it to the first three quarters.

Figure 1.9: Full-Game Thrill Trajectories (left) and Fourth-Quarter Extrapolated Thrill Trajectories (right) by Percentile

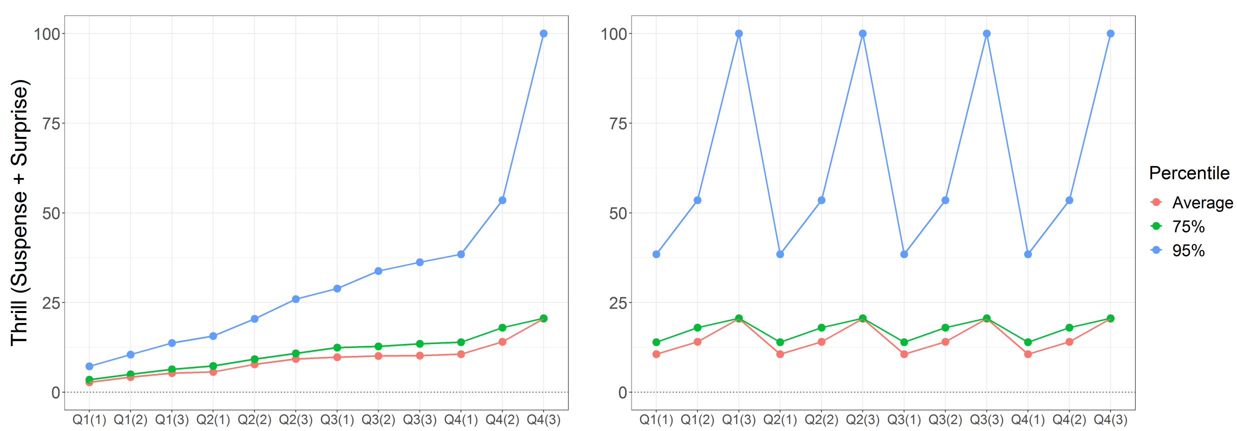

One can see that this exercise generates a great deal of additional thrill in a game, especially for the 95th percentile of thrill observed in my sample of games. Table 1.8 presents changes in viewership-minutes associated with this counterfactual scenario. I perform two different extrapolations: (i) extrapolating the thrill trajectory from the fourth quarter to the previous three quarters, and (ii) extrapolating the thrill trajectory from the second half to the first half. In the first scenario, I increase the number of final outcomes in an event from one to four, while in the second I increase it from one to two. I assess changes in viewership-minutes within a game (i.e. at the quarter or half level, depending on the scenario), at the

[22]Lower percentiles of the thrill distribution do not exhibit this monotonic pattern. This is due to the fact that games featuring less thrill experience the peak of their thrill trajectory earlier in a game. Figures 1.4 and 1.5 show this clearly for the bottom quintiles of the absolute score differential distribution.

Table 1.8: Viewership Changes from Increased "Finality" (1000's of Viewership-Minutes)

			Level of Thrill		
			Average	75^{th} Percentile	95^{th} Percentile
Quarters					
	Game-Level		1,207.86 (0.60)	1,159.12 (0.57)	1,791.94 (0.87)
		First Quarter	632.18 (1.33)	629.28 (1.33)	868.20 (1.80)
		Second Quarter	367.21 (0.68)	359.92 (0.67)	589.74 (1.09)
		Third Quarter	208.47 (0.37)	169.93 (0.30)	333.99 (0.58)
		Fourth Quarter	–	–	–
	Season-Level		1,485,671 (0.60)	1,425,716 (0.57)	2,204,084 (0.87)
Halves					
	Game-Level		818.04 (0.41)	841.96 (0.42)	1,168.05 (0.58)
		First Half	818.04 (0.82)	841.96 (0.85)	1,168.05 (1.15)
		Second Half	–	–	–
	Season-Level		1,006,195 (0.41)	1,035,616 (0.42)	1,436,698 (0.58)
Full Game			–	–	–

Note: Percent-changes in viewership-minutes are indicated in parentheses.

level of a game, and at the level of an NBA season, which encompasses 1,230 total regular season games (not including playoff games).

Looking at Table 1.8, one can see that when extrapolating the thrill trajectory from the fourth quarter to the first three quarters, viewership-minutes increases between 1.16 - 1.79 million minutes over the course of an entire game, which corresponds to a 0.60-0.87% increase. At a season-wide level, which corresponds to 1,230 total games played during the regular season, increases in viewership-minutes range between 1.43 - 2.20 billion. Viewership-minutes increase differentially more in the first quarter, since the typical thrill trajectory is at its lowest during the earliest portions of a game.[23] Comparing the second-half extrapolation and fourth quarter extrapolation scenarios, viewership minutes increase by a lower amount in the second-half case, which is intuitive given that we are introducing less finality in this scenario.

Contextualizing the effect sizes observed from this counterfactual, Figures 1.5 and 1.6 show that thrill can enhance viewership by up to 30% in the fourth quarter, and 5-10% over the course of an entire game. In the counterfactual, thrill enhances viewership between 0.4-1.8%, which is significantly lower. It is clear from this comparison that the evolution

[23]There is an implicit assumption here that viewers are treating each quarter as an independent outcome, and that there are no differential responses to thrill based on the timing of when the outcome takes place. For instance, if four quarters are played in a game and each counts equally as an outcome, viewers may still respond more to thrill in the final quarter, as it represents the last remaining outcome in a game.

of thrill within a game outweighs a structural modification to a game's structure (without compromising the integrity of the outcome) from the standpoint of increasing viewership. It also provides important insight into the nature of thrill itself – people enjoy thrill because of it's stochasticity within a game, and the magnitude of viewership increases associated with games featuring more versus less thrill reflects this.

This research suggests a deeper set of questions about how to optimally design contests, and how to balance parity or rules within a game to maximize consumer attention. One natural extension of this is to think about different designs of game endings to induce additional thrill – an example of this is the difference in overtime formats between football in the NFL and NCAA. Future work should aim to carefully assess additional proposed changes in these leagues, as well as address implications for other forms of entertainment, particularly ones that feature elements of skill and thrill.

1.6 Conclusion

This paper uses revealed preference methods to explore and quantify demand for non-instrumental information in entertainment, examining the thrill associated with the trajectory of an event, and the skill associated with information-conveying agents. Using the theory presented in Ely et al. (2015), I perform an empirical analysis assessing the effect of suspense and surprise on consumer attention, introducing an additional element assessing spectator preferences for the skill of agents involved.

I observe three primary findings. First, skill is an important driver of a viewer's initial decision to watch a game, while expected thrill has no significant impact. In particular, for a doubling of skill present in a game, initial viewership increases by approximately 11%. For context, the presence of LeBron James alone results in an approximately 13.5% increase in initial TV viewership. These results are remarkably similar to those found in Kaplan (2020), which uses secondary ticket marketplace data to assesses the impact of a superstar absence announcement for a specific game on listed prices. The analysis finds that the absence of LeBron James leads to a 13% ($42/ticket) average reduction in ticket prices.

Second, I measure thrill using the evolution of absolute score differential during a game. I find that a one-point decrease in the absolute score differential does not impact viewership in the first or second quarters, but increases viewership by 1.2% in the fourth quarter, strongly supporting the idea that viewers demand thrilling games, not just games that are close. Contextualizing these results further, second half ratings are 8.2-20.5% lower on average for games with a 14+ score differential margin compared to a 0-8 margin, while these differences are 12.0-29.6% when only examining the fourth quarter. I extend this analysis to look at absolute score differential during a game in reference to the initial point spread. I find that a one-standard deviation increase in score differential in reference to the spread during the final quarter segment (9.3 points) causes a 6.4-7.3% reduction in viewership.

Third, I directly implement the structural definitions of suspense and surprise from Ely et al. (2015). I find that a doubling of suspense during a game increases viewership by

0.4-0.6%, and a doubling of surprise by 0.6-1.0%, not accounting for any additional impacts associated with skill. For additional context, in the last segment of the fourth quarter, a 0-2 point game averages 18 times more suspense and 14 times more surprise than a 14+ point game. In this case, viewership would be approximately 6.8-10.8% higher only through suspense and 9.0-13.4% only through surprise. Additionally, I find that a doubling of the number of All-Star fan votes on the court at a given time during the game leads to a 1.9-2.4% increase in viewership. Interestingly, I find a negative interactive effect between suspense and skill, suggesting that heightened suspense leads to differentially higher viewership with lower skill on the court, supporting the traditional notion that spectators may only turn on games featuring lesser-known players (or teams) if they are nearing the end and exhibiting sufficiently high suspense.

There are several avenues of future work based on the findings and implications presented here. First, micro-level data on individual viewership that can be linked to demographic information would provide a rich assessment of heterogeneity in viewership patterns in response to entertainment characteristics, particularly in response to within-game information updating. A complementary experiment could measure consumer attentiveness to advertisements in response to skill and thrill at a specific point of a game, and assess purchasing conversion rates for advertised products and services. Individuals' opportunity cost of leisure time can also be measured using exogenous variation in the thrill of games experienced in different locations around the world at different times of the day.

A different set of analyses should examine the outlay of non-instrumental information in different domains. For example, individuals gain enjoyment from being informed about important local, state, and national elections, often using different entertainment platforms to garner information. In the case of politics, many individuals are consuming politically-relevant information simply for the sake of entertainment. But there is an interesting dimension of civic engagement that may result depending on the thrill of the information. Specifically, how might social media engagement, donations, and even voting activity respond to the thrill of an election? Understanding these impacts has never been more important.

Chapter 2

The Economic Value of Popularity: Evidence from Superstars in the National Basketball Association[1]

2.1 Introduction

Understanding the interest in and impact of superstars began with Rosen (1981), which developed a model to explain how certain talented individuals in a specific occupation are able to differentiate themselves from the rest of a pool of individuals, and obtain differentially higher salaries as a result. Superstars are prevalent across all types of activities and industries – for instance, the impact Steve Jobs made on the technology sector, Henry Ford in automobile manufacturing, or Rosen himself on the field of economics. Rosen (1981) emphasized two common elements describing superstars: "a close connection between personal reward and the size of one's own market," and "a strong tendency of market size and reward to be skewed towards the most talented people in the activity."

A natural place to examine superstars is in sports, where the word "superstar" is actually used to describe individuals that fit Rosen's description. While different leagues have varying degrees of superstar influence, the National Basketball Association (NBA) is widely regarded as "superstar-driven" (Heindl, 2018; Knox, 2012). More than other sports, the NBA's surge in popularity can be largely attributed to the fame and marketability of its top players (Morris 2018; Adgate 2018). In addition, the observable nature of productivity, popularity, and resulting compensation makes it an appealing empirical laboratory to measure and evaluate the economic impacts of these individuals. While several studies have examined the impact of player performance on compensation (as laid out in Rosen and Sanderson

[1] I would like to thank Vaibhav Ramamoorthy, Cheenar Gupte, and Amit Sagar for excellent research assistance. Thanks to Deepak Premkumar, Jim Sallee, Sofia Villas-Boas, Joshua Wilbur, Hal Gordon, and David Zilberman for thoughtful comments. Special thanks to participants at the *2019 MIT Sloan Sports Analytics Conference* for useful comments and conversations.

2001), there is limited understanding of how different individual characteristics, including both player ability (i.e. the actual contribution to the performance of their team) and popularity, translate into a compensation distribution, and whether or not these relationships support the theory of superstars. Furthermore, there has been limited empirical research comprehensively examining the returns to a superstar's marketability, and how the pure popularity of an individual drives demand for their services.

This paper attempts to understand the impact of individual superstar ability and popularity on consumer demand, and discuss implications for individual compensation. In particular, what is the overall premium in terms of television viewership and ticket prices associated with watching superstar players, and to what extent is this premium driven by player ability versus popularity? More specifically, what is the loss in value, as measured by listed price changes on a secondary ticket marketplace, associated with the announcement of a specific superstar's absence for a game? This analysis provides a novel framework to assess the value of superstars, particularly in sports, which can be applied to other select industries. The findings also have significant ramifications for leagues in general, including policies on announcement timing of player absences and player compensation schemes based on popularity. They also inform team decision-making with respect to player personnel and potential implications of introducing dynamic pricing in primary ticket marketplaces. Finally, the findings can help inform national TV networks' strategies about which games to televise, and in the case of superstar absences, whether to "flex" to a different matchup.

What is a superstar player? Superstars are not necessarily the "best" players from a statistical standpoint – they are defined as much by their ability as they are by their popularity (Adler 1985). To encapsulate both of these factors in the analysis, and attempt to distinguish between them, I analyze all players who made the NBA All-Star team as a starter at least once across the 2017-18 and 2018-19 seasons. This criteria provides a great comparison of popularity and ability – the starters for the All-Star rosters are based on a weighted average of votes between fans (50%), players (25%), and select media members (25%), while the reserves are selected using a vote by head coaches. Of course, there are other metrics that may be useful to rank player popularity and ability; jersey sales for each player may indicate their relative popularity, and a player's "efficiency rating" (PER) may indicate their ability. I use the All-Star criteria as a cutoff to determine which players to analyze, as it has been used in previous studies and incorporates notions of both popularity and ability (Berri and Schmidt 2006; Yang and Shi 2011; Jane 2016).[2] One of the primary goals of this paper is to isolate independent variation in popularity and ability across superstars so that their impacts on a player's economic value can be compared.

The empirical approach taken is two-fold and relies on data from the 2017-18 and 2018-19 NBA seasons. First, I use matchup-level data to estimate the impact of player popularity and ability on ticket prices and television ratings. I find that a 1% increase in the popularity

[2] Along with all players selected to the 2017-18 and 2018-19 All-Star teams, we also include players that would have made the All-Star team had the fan vote counted 100%. This includes Manu Ginobili, Luka Doncic, and Derrick Rose. We also include Dwyane Wade and Dirk Nowitzki, who were additions to the 2018-19 All-Star team made by Commissioner Adam Silver.

of a matchup (as measured by the cumulative number of All-Star votes of all players playing) leads to a 0.12-0.13% increase in ticket prices and TV ratings, while cumulative player ability (as measured by cumulative player-efficiency rating of all players playing in a matchup) has no statistically significant impact. These results provide evidence that the superstar allure for fans is primarily associated with a player's popularity.

The next component of the analysis relies on within-matchup, temporal changes in ticket prices. I take advantage of exogenous variation in a superstar's availability for specific matchups, where players may miss games for unforseen reasons. Additionally, superstar player absences have been an especially relevant point of discussion with respect to the NBA, since absences are trending upwards as a result of teams choosing to "load manage" (purposefully rest) players earlier in the season and more often (Whitehead 2017). Using difference-in-differences (DID) and event-study methodologies, I examine ticket price impacts when superstar players are announced out of specific matchups. The findings suggest statistically significant price declines for the most popular stars, including LeBron James, Stephen Curry, and Dwyane Wade, among others, ranging from 4-16% ($7-$42) per ticket. In addition, I analyze absences in home vs. away games, finding that the away effects for LeBron James and Stephen Curry are even larger, at 21% ($75) per ticket for LeBron and 18% ($55) per ticket for Curry. The findings from the two sets of analyses are largely consistent both qualitatively and quantitatively – the most popular stars lead to the largest impacts on prices and ratings, the relationship between popularity and impact on prices is convex, and these impacts are on the order of 4-16%.

The paper will proceed as follows. First, I review the relevant literature this paper contributes to. Next, I discuss the data collection strategy and presents relevant summary statistics. Then, I overview the empirical strategy and assumptions for identification. The following section showcases the results. Next, I discuss and synthesize the findings. Finally, the paper concludes.

2.2 Literature Review

This work falls into several important strands of literature. First, there has been substantial research in hedonic pricing, which attempts to value specific, nonmarket attributes of goods. It also contributes to the literature on dynamic pricing and strategic interactions among buyers and sellers in secondary ticket marketplaces. Finally, several papers have examined the impact of superstars in different labor contexts, including sports, suggesting that quality and popularity of players are important factors for spectators. This body of literature examines superstar athlete impacts on a variety of metrics, including attendance, player salaries, and broadcast audiences. I extend all of this literature by (1) using a novel and well-identified methodology to esimate consumer willingness-to-pay to watch superstars by looking at ticket price movements in a secondary ticket marketplace, (2) testing heterogeneous, matchup-specific factors that may impact the value associated with a superstar, and (3) leveraging unique, high temporal frequency microdata on ticket prices for all NBA games for the 2017-18 and 2018-19 seasons.

Hedonic Pricing and Player Value

The literature on hedonic pricing aims to understand and estimate the relative value of each attribute of a good. The theory of hedonic pricing was developed in Rosen (1974), which was the first paper to describe the total value of a good as a combination of the values of its attributes. There have been numerous empirical papers attempting to price attributes in different settings, from vehicles (Busse et al. 2013; Sallee et al. 2016) to air quality (Currie and Walker 2011; Chay and Greenstone 2005) to real estate (Luttik 2000). These papers use data on similar products with varying attributes of interest in an attempt to estimate the marginal value of these attributes. Additionally, Scully (1974) was the first paper to examine the marginal revenue product of athletes, comparing how much they are paid with how much they contribute to their team's success, finding that player salary relative to their contribution to winning was still lower than 50%. Kahn (2000) provides a seminal overview examining the key relationship between athlete productivity and pay, how players are allocated across a league, and how league market structures affect player salaries. The research presented here contributes to this literature by being the first to utilize rich microdata with substantial variation in confounding factors (e.g. competitiveness of opponents, market size, etc.) to perform a well-identified, plausibly exogenous estimation of the economic value of superstars.

Dynamic Pricing in Secondary Ticket Marketplaces

The second relevant strand of literature includes work on dynamic pricing in primary and secondary marketplaces, including event tickets, hotels and home-sharing (e.g. AirBnB), and airline tickets (Jiaqi Xu et al. 2019; Williams et al. 2017; Sweeting 2012; Blake et al. 2018; Levin et al. 2009; Oskam et al. 2018; Mills et al. 2016; Courty and Davey 2020). Early research on dynamic pricing examined how in airline ticket markets, consumers often learn new information about their demands over time, which may be an important reason for the existence of both primary and secondary ticket marketplaces (Courty 2003a). Additionally, the dynamic pricing nature of secondary ticket marketplaces allows for real-time updating of preferences of both consumers and producers, which may lead to real-time price changes in response to realized information about an event (Courty 2003b). The research presented here differs substantially from much of the previous theoretical work on pricing in these marketplaces, in particular ticket marketplaces, in that it relies on changes in the quality of attributes of an event to determine individuals' value for those attributes (i.e. their value for watching a specific superstar play).

While this research builds on many of the theoretical aspects of ticket pricing, it takes a primarily empirical approach. The seminal empirical paper in this field explaining dynamic pricing patterns using secondary ticket marketplace microdata is Sweeting (2012), which examines Major League Baseball games and develops a game-theoretic framework to discuss the dynamics of buyer-seller interactions on secondary marketplaces as a matchup gets closer. Similar to this research, Sweeting (2012) finds that much of the buying and selling activity in marketplaces, including price adjustments, occurs in the few days before an event. Clarke (2016) uses microdata from a secondary ticket marketplace to assess seller dynamics on ticket resale markets, finding that there is a great deal of heterogeneity in seller pricing strategies. Most notably, Clarke finds that 40% of sellers have a "negative scrap value" (i.e. if their ticket does not sell, they have a zero or negative value associated with attending the game) and 20% of sellers value their tickets above the franchise's face value. Thus, negative price effects associated with the announcement of a superstar absence may reflect a lower bound (in absolute value terms) because sellers who do not adjust still have a weakly negative value associated with this announcement, but may face transaction costs that are too high or fall victim to the "sunk cost fallacy."

Economics of Superstars

Rosen was the first to understand the economic notion of superstars in Rosen (1981), which was later expanded upon in Rosen and Sanderson (2001). Their work developed a model to explain how certain talented individuals in an occupation are able to differentiate themselves from the rest of a pool of individuals, and obtain differentially higher salaries as a result. An expansion of this work attempts to differentiate between the popularity and ability of a star performer; namely there may be a premium for watching a player with average talent, but who is quite popular for other reasons (Adler 1985). Krueger (2005) examines "rockstars" in the music industry, creatively using the number of millimeters of print columns in *The Rolling Stone Encyclopedia of Rock and Roll* to measure musician prominence, showing that a 200 millimeter increase in print leads to a 5-15% increases in prices. Interestingly, he also finds the superstar effect (measured by concert ticket prices) nearly tripled between 1981-2003. The research presented here synthesizes nicely with these findings, finding that (i) there is a convex relationship between a superstar's presence and willingness-to-pay, and (ii) superstar popularity is a more meaningful factor in ticket price and TV ratings adjustments than superstar ability.

Other papers have examined superstar effects in the context of sports. In professional golf, Brown (2011) assesses the impact of Tiger Woods' presence on the performance of competing golfers, finding an adverse effect on their performance when competing in the same tournaments as Woods. Other studies have examined brand alliances between companies and superstars, determining the extent to which these partnerships drive value (Yang et al. 2009; Chung et al. 2013). In professional soccer, there have been several empirical studies that attempt to identify television audience demand for superstar talent (Buraimo and Simmons 2015), and quantify the characteristics that affect superstar wages, including both on-field performance and popularity (Scarfe et al. 2020; Bryson et al. 2014; Bryson et al. 2013). In German professional soccer, Lehmann and Schulze, 2008 regresses salary proxies of 359 players on indicators of talent and performance, finding neither explains salaries for the upper 95^{th} percentile of players. Franck and Nüesch (2012) find opposite evidence, namely that both player talent and popularity increase the market value of star players. In addition, they find that marginal returns to productivity in terms of player salaries are much larger among stars than average players. Hausman and Leonard (1997) was the first empirical paper to analyze the effect of superstar players in the NBA, examining their effect on attendance and television viewership. They find substantial impacts for these players, especially in the case of away games, where fans in those markets were enthusiastic to watch these superstars when they came to town. More comprehensive analyses have estimated the impact of All-Star votes on fan attendance, finding that top vote-getters can lead to thousands of additional tickets sold (Berri and Schmidt 2006; Jane 2016). This paper expands on previous analyses by measuring the relative impacts of popularity and ability on willingness-to-pay using ticket price data, and examining plausibly exogenous price changes to causally identify superstar value under heterogeneous conditions.

2.3 Overiew of Data Collection and Characteristics

This project leverages unique, high temporal frequency microdata from a large online secondary ticket marketplace, as well as data on exact timing of injury announcements for different players. The analysis is supplemented with television ratings data from The Nielsen Company.[3] This section (i) describes the data collection and organization methodology for each source of data, and (ii) presents high-level summary statistics.

Overview of Data

Secondary Ticket Marketplace

An integral component of this project was collecting ticket-listing data from a large, online secondary ticket marketplace that offers tickets for events ranging from concerts to sporting events. The analysis relies on the use of such a marketplace since sellers and buyers can react instantaneously to announcements about player absences.

This data was accessed by routinely querying a REST (Representational State Transfer, a protocol built on-top of the standard web protocols) service provided by the secondary ticket marketplace every 30 minutes (or a total of 48 collections per day) for every remaining NBA matchup in the season.[4] For each ticket listing, metadata on the corresponding NBA game was collected (e.g. home and away teams as well as date and time of matchup), data on the listing characteristics (listing price, quantity available, and a listing identifier), and

[3]Data granted from The Nielsen Company (US), LLC. The conclusions drawn from the Nielsen data are those of the researchers and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

[4]A REST service is an HTTP-backed protocol that defines a set of rules for querying, updating, adding, and deleting data on a website. The REST protocol is how a website can securely expose its database without giving everyone unlimited control over the data.

identifiers for the time of data collection. This data provided high granularity snapshots for observing price changes before and after superstar absence announcements.

The analysis presented in this paper relies on a sample of ticket prices within 3 days of a matchup, primarily because this is when the majority of single-game superstar absence announcements occur, and because most of the buyer/seller activity on the secondary market occurs during this timeframe. Additionally, ticket buyers and sellers may exhibit different types of responses (in terms of timing) depending on the amount of time between the announcement and the affected game. Announcements impacting games in this three-day window are likely to exhibit more immediate changes, and thus make for clearer analysis of price impacts.[5]

Absence Announcements

Absence announcements were collected from a popular fantasy basketball website, which provides detailed injury information and other reports for all players. This website provides regular updates on announcements from teams regarding player absences. Since all announcements are documented and accessible going back several years, I examined announcements pertaining to each All-Star player for the 2017-18 and 2018-19 NBA seasons. Because of the complex nature of many of the announcements and their timing, I manually examine every announcement pertaining to each of these players to determine which corresponded to missed games, and the exact time an announcement was made. When announcements were vague about the expected duration of missed time for a player (for example, if a player was announced to be out for "several weeks"), a very conservative lower bound horizon of the expected number of missed games was used. Once all relevant announcements were classified, I matched the time of announcement applicable to a specific game and player to ticket prices at that time for the relevant game.

Television Ratings

A supplemental analysis relies on television ratings data for all nationally televised games during the 2017-18 and 2018-19 seasons from the Nielsen Company. The primary metric I examine is the projected number of total households watching across the United States at the start of each matchup. This data, as well as the ticket price data (at a more aggregated level), is used to observe the "total superstar impact" of a game (as measured by the cumulative number of All-Star votes across all players suiting up for a game) on initial TV ratings and ticket prices.

[5] An interesting and important avenue of future research is to examine ticket price movements in response to longer-term absence announcements.

Game Characteristics

To perform the panel analysis analyzing ticket prices and TV ratings at the matchup-level, I rely on a rich dataset of matchup-specific characteristics collected from several different sources, including `NBA.com`, `fivethirtyeight.com`, and `Basketball Reference`. These characteristics include state variables corresponding to each matchup (i.e. date and time information, absolute point spread, aggregate number of All-Star votes of all players playing, aggregate player-efficiency rating of all players playing, average winning percentage of the two teams, etc.). Section B.3 lays out the set of covariates included and provides relevant summary statistics.

Summary Statistics

Secondary Ticket Marketplace Data

A summary of relevant variables collected from the secondary ticket marketplace microdata is presented in Table 2.1. It should be noted that these are the primary summary statistics of the per-game *averages* for the continuous variables (listing price and quantity per listing), and the per-game *counts* for the count variables (number of observations, listing IDs, section IDs, and collection IDs). The data spans 2,330 NBA matchups, corresponding to 95% of the total number of regular games played over two NBA seasons (2,460).[6] The "Listing Price" refers to the price posted by a seller for a specific listing. The "Quantity per Listing" denotes the number of seats available in a specific listing posted by a seller. The "Listing ID" is a unique listing-specific identifier, the "Collection ID" is a unique identifier corresponding to when the data was collected (i.e. each 30 minute collection gets a unique identifier), and "Section ID" corresponds to the section of the arena the listing is located in. Finally, "Number of Observations" corresponds to the number of unique listing-by-collection ID data points for each matchup.

Table 2.1: Across-Game Ticket Data Summary Statistics (2,330 Total Matchups)

Data Characteristic	Mean	Std. Dev.	Min.	Max.
Num. Obs.	37,660.88	26,648.85	70	215,346
Listing Price	$157.12	$107.06	$12.75	$995.01
Quantity per Listing	3.39	0.73	1.92	5.69
Listing IDs	826.80	682.10	28	5,357
Collection IDs	114.31	29.43	4	139
Section IDs	113.30	35.66	18	228

[6] Reasons for missing data for certain matchups include server restarts and changes of event-names mid-season on the secondary ticket marketplace that were not automatically identified by the data collection program.

Table 2.1 shows that there is an average of 114.31 collection times for each matchup, which corresponds to approximately 57.16 hours prior to each matchup. There is an average of 826.80 unique listings per matchup across an average of 113.30 different arena sections. The average per-matchup listing price is $157.12 with a quantity per listing of 3.39.

Because of the high temporal frequency of this microdata, I can observe the time trends of average listing price and quantity of tickets posted to a secondary marketplace for each matchup. Figure 2.1 presents three different quantity time trends in terms of "hours to game": the top pane presents the average total quantity of tickets available on the secondary marketplace for each matchup, the second pane presents the average number of tickets *added* (i.e. posted by sellers) to the marketplace per matchup, and the third pane the average number of tickets *sold* on the marketplace per matchup. I assume the disappearance of a listing on the marketplace implies that this listing was sold, either to a buyer or to the "seller" of the listing who decided to go themselves.[7] One can see that the quantity of tickets available for a given matchup declines as the matchup approaches. This is intuitive, as these tickets represent a "perishable good" and have no value once a matchup is completed. Interestingly, the average number of tickets posted (added) to the marketplace is somewhat uniform in terms of hours to game (with the exception of dips during night-time hours when most sellers and buyers are asleep), but the average number of tickets sold spikes in the five or so hours before a game.

Figure 2.1: Per-Game Average Number of Active, Added, and Sold Listings by Hours to Game[8]

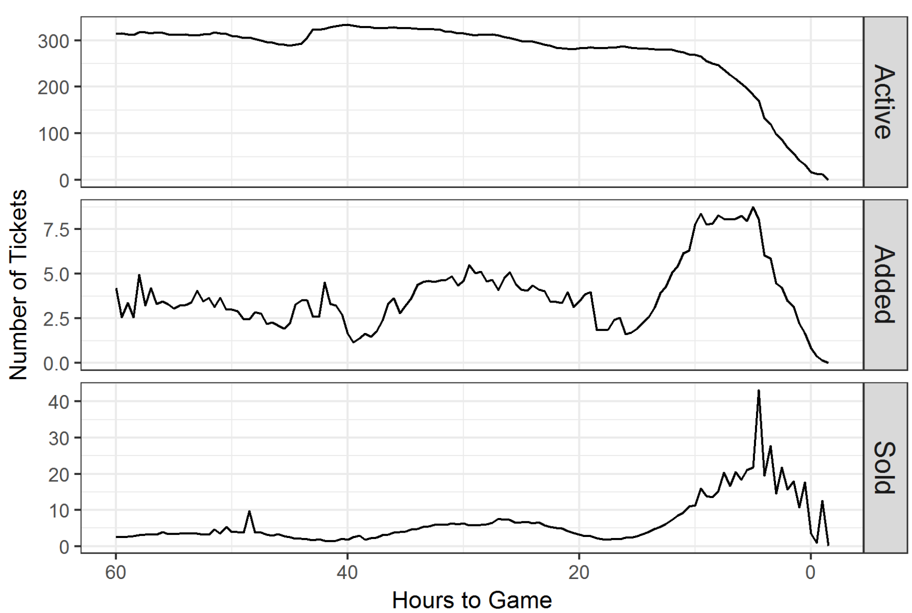

[7]In other words, a seller may have their tickets purchased by another buyer, or decide to "purchase" their own tickets (i.e. remove the listing and go to the game themselves).

Additionally, Figure 2.2 plots the average listing price across all matchups by hours to game. There is generally a downward trend in prices as a matchup approaches, decreasing from around \$145/ticket two days before a matchup to around \$100/ticket just before game-time. It is also clear that the volatility in prices substantially increases closer to the game. This may be attributed to an increase in activity on the marketplace – there are matchups where sellers are trying to unload tickets and continue to lower prices, and other matchups where buyers are trying to obtain tickets, causing remaining sellers to increase their prices. This heightened activity appears to occur beginning about 10 hours before a game starts.

Figure 2.2: Average Listing Price by Hours to Game

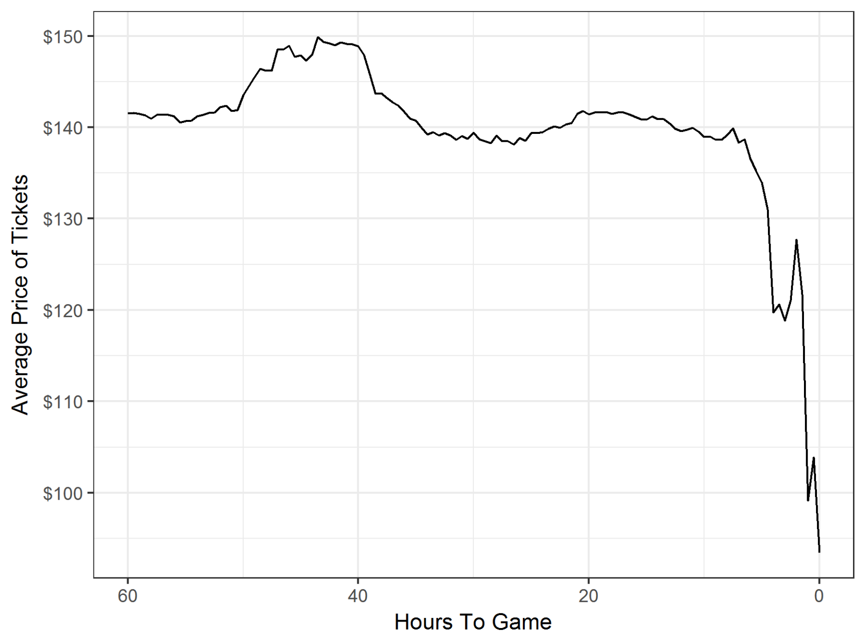

Player Absence Announcements

The second set of the collected data is the timing of player absence announcements. Figure 2.3 presents the distribution of announcements for all qualifying All-Star players across the 2017-18 and 2018-19 NBA seasons in terms of hours to game. In the case of announcements referring to multiple games, I only include observations corresponding to announcements within three days of a game to maintain consistency with the chosen time window.[9] In Figure 2.3, there are 192 announcement-matchup pairs falling within three days of a matchup.

[8]Please note the different y-axis scale for each pane.

[9]Table 2.2 provides both the "total number of games missed" (not just the most immediate game corresponding to a given announcement) for each All-Star player corresponding to all documented announcements, as well as the "total number of games analyzed" in our analysis.

Figure 2.3: Distribution of Player Absence Announcements by Hours to Game

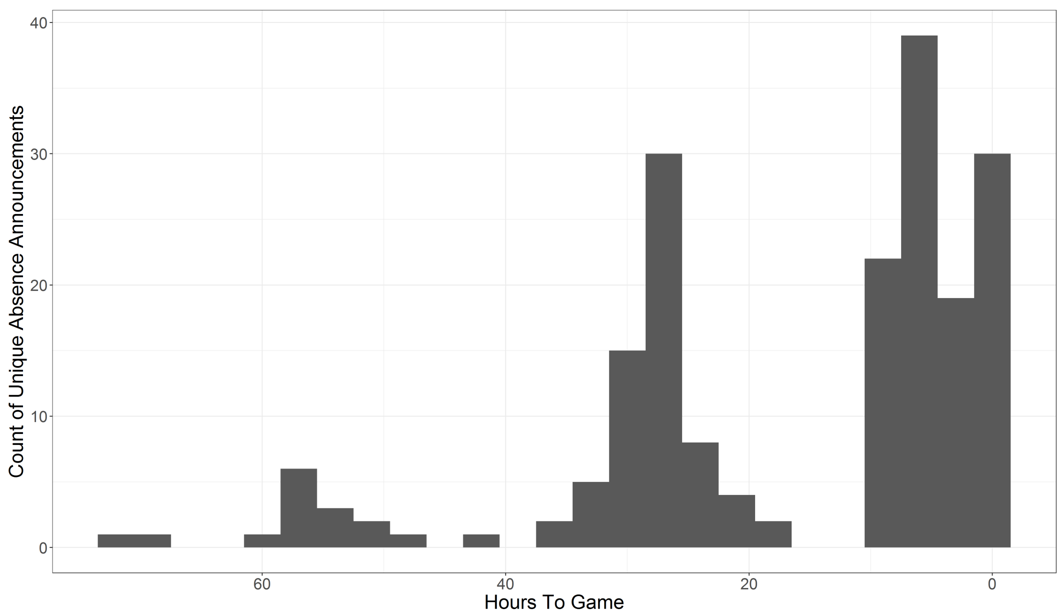

One can see that most of these announcements occur within 12 hours of a game, some coming as close as a few minutes beforehand. This inherently limits the sample size of games that can be analyzed, since there needs to be an adequate timeframe pre- and post-announcement to witness ticket price changes. Many announcements also occur approximately 24 hours prior to a game, which may be the result of a player experiencing an injury during the first game of a back-to-back, or an injury that does not require a "game-time decision." There are also noticeable dips in announcement counts 12-20 hours prior to a game because these times often fall during the middle of the night. Rarely do announcements for a player absence for a specific matchup occur more than 36 hours prior to a game.[10]

Table 2.2 presents the names of each starting All-Star player (or players that would have been voted a starter had the fan vote counted for 100%), how many "qualifying" games they missed (i.e. an explicit announcement for a matchup indicating the exogenous nature of a player's absence) as a result of injury, rest, or "other" reasons, the total number of games, and the number of games for each player that was included in the analysis on ticket price changes. For each listed player, I am able to analyze most, if not all, of the qualifying games they were absent for. Reasons for not being able to analyze certain qualifying games include if the announcement occurred "too close" to the matchup, "too far" from a matchup (since I only analyze announcements within three days of the corresponding matchup), missing ticket price data as a result of event-name changes on the secondary marketplace, or if another

[10]As mentioned previously, our analysis does not consider the effect of a long-term injury announcement on games more than 3 days into the future.

Table 2.2: Count (by Reason) of Qualifying "Missed-Games" for each Starting-Caliber All-Star Player

Player	Injury	Rest	Other	Total	Total Analyzed
Anthony Davis	26	3	1	30	21
DeMar DeRozan	4	3	0	7	0
DeMarcus Cousins	35	6	0	41	0
Giannis Antetokounmpo	17	0	0	17	16
James Harden	10	2	0	12	6
Joel Embiid	12	6	0	18	13
Kemba Walker	2	0	0	2	2
Kevin Durant	17	1	0	18	16
Kyrie Irving	35	1	1	37	27
Paul George	8	0	0	8	7
Stephen Curry	42	1	0	43	20
Luka Doncic	10	0	0	10	9
Dwyane Wade	3	0	7	10	8
Dirk Nowitzki	2	2	0	4	3
LeBron James	22	3	0	25	16
Kawhi Leonard	6	14	1	21	21
Derrick Rose	32	0	0	32	19

I did not analyze games in which DeMar DeRozan (row 2) or DeMarcus Cousins (row 3) missed, as both of them did not make the All-Star Team during the 2018-19 season (despite being All-Star starters during the 2017-18 season). The criteria for a player to be analyzed was that they were an All-Star during both seasons, a starter during at least one of the two seasons, or would have been voted an All-Star starter with 100% weight on the fan vote at least one of the two seasons. Also note that Manu Ginobili is not present, as he did not miss any qualifying games during the 2017-18 season in which he would have been voted an All-Star starter with a 100% weighted fan vote.

superstar was announced as out for that qualifying game as well.

Matchup Characteristics, Television Ratings, and All-Star Votes

A summary of relevant game characteristics data, which was collected from `Basketball Reference`, `fivethirtyeight.com`, and `NBA.com` for all NBA games (regular season and playoffs) during the 2017-18 and 2018-19 seasons, is presented in Table 2.3. One can see the average number of cumulative All-Star votes in a matchup is just over 3.8 million. For context, LeBron James received 4.6 million votes and Stephen Curry 3.8 million for the 2018-19 season, suggesting that each of these players *alone* generate just as much popularity as the average NBA game.

Table 2.4 summarizes these same game characteristics as well as the total projected number of viewers from the television ratings data. Note that this data comes from the

Table 2.3: Game Characteristics Summary Statistics (2,624 Total Matchups)

Data Characteristic	Mean	Std. Dev.	Min.	Max.
Aggregate # of All-Star Votes (1,000's)	3,856.87	3,279.58	31.10	18,347.76
Absolute Point Spread	5.84	4.28	0	26
Aggregate Player Efficiency Rating	302.25	32.48	169.50	431.90
Avg. Final Win %	0.50	0.10	0.22	0.76
Aggregate Market Size (1,000's of people)	3,530.50	1,051.86	2,025	7,700
Attendance	18,056.58	1,964.23	10,079.00	22,983.00

sample of all *nationally televised* games during the 2017-18 and 2018-19 seasons, of which there were 480 in total (332 non-playoff games and 148 playoff games). Table 2.4 exhibits a couple of interesting characteristics of this data. First, average viewership is more than 2 million for a nationally televised game. Furthermore, when separating the sample between playoff and non-playoff games, average viewership increases from 1.5 million for non-playoff games to nearly 3.5 million for playoff games. Next, the range of aggregate number of All-Star votes found in a national TV matchup is quite large. On average, there are nearly 6.7 million All-Star votes across all players in a matchup, but this can be as little as 372,000 and as high as 18.35 million. A game featuring LeBron James or Stephen Curry alone would include more than an order of magnitude more All-Star votes than the lowest total All-Star votes game from this sample!

Table 2.4: TV Ratings and Game Characteristics Summary Statistics (480 Total Matchups)

Data Characteristic	Mean	Std. Dev.	Min.	Max.
Projected # of Viewers (1,000's)	2,134.92	1,645.06	265	11,151
Aggregate # of All-Star Votes (1,000's)	6,669.58	3,897.47	372	18,347.76
Absolute Point Spread	4.93	3.56	0	18
Aggregate Player Efficiency Rating	313.96	34.72	226.80	430.60
Avg. Final Win %	0.60	0.08	0.27	0.75
Aggregate Market Size (1,000's of people)	3,910.41	1,105.17	2,125	7,199

Finally, Figure 2.4 visually depicts the cumulative distribution of average All-Star fan votes across all 659 eligible players over the 2017-18 and 2018-19 seasons. One can see that a sizable majority of players receive a negligible number of votes, and that 75% of all votes are concentrated within the 95^{th} percentile of players (approximately 33 players). This is a product of the concentration of popularity towards the top several players in the league.

Figure 2.4: Cumulative Distribution of All-Star Votes by Player Rank

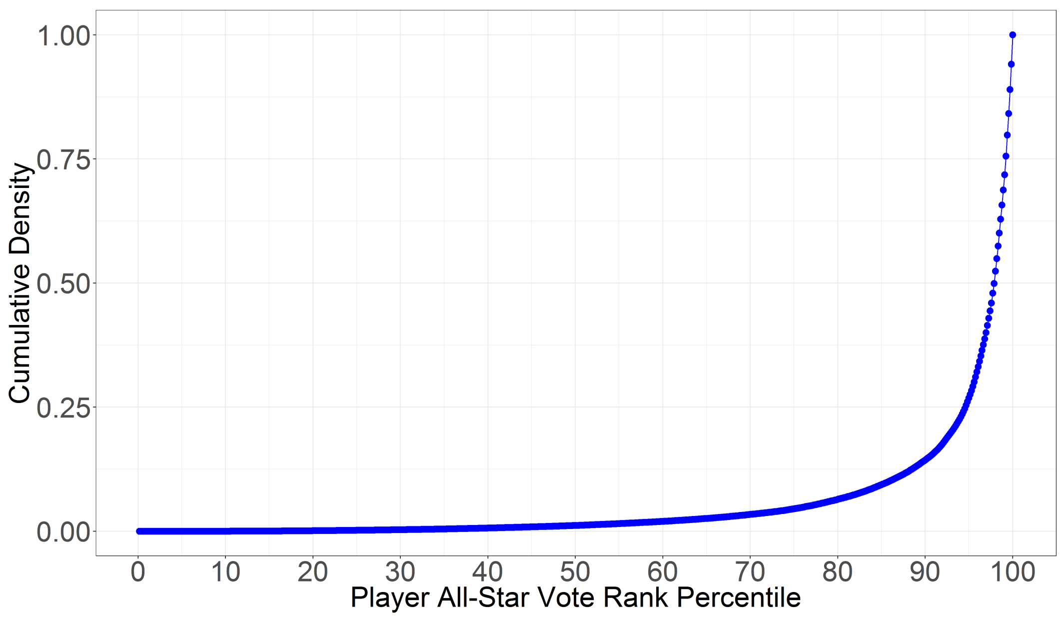

2.4 Empirical Methodology

The empirical methodology in this paper is two-fold. First, I use a fixed-effects panel regression approach to estimate the impact of player popularity and ability, among other factors, on ticket prices and TV ratings. I rely on a quasi-LASSO framework (using motivation from Athey and Levin 2001) to determine the relationship between residualized popularity, ability, and team quality on residualized ticket prices/TV ratings over the entire support of the data using a rich set of controls with flexibility in the functional form. Next, difference-in-differences (DID) and event study frameworks are used to identify the causal effect of a specific superstar's absence on ticket prices within a certain matchup. Under important assumptions regarding identification, these estimates represent the per-ticket value of each superstar's presence to fans in attendance. This framework relies on a plausibly exogenous "announcement" of a player's absence for an upcoming game, at which point ticket prices for that game should respond according to the missing player's value. I then conduct heterogeneity analyses to determine how these values differ for home vs. away absences and based on the franchise value of the home team.

Panel Analysis

The initial analysis examines ticket prices and TV ratings at the matchup-level. First, using a simple fixed effects model, I estimate the impact of player popularity, as measured by the cumulative number of All-Star votes of all players in a specific matchup, player ability, as measured by the cumulative PER of all players in a specific matchup, and team quality, as measured by the current average win percentage of the two teams in the matchup, on ticket prices and TV ratings. One feature of the NBA regular season schedule that allows for rich variation in these variables is that each team plays each other no less than two times and no more than four times. The estimating equation is written below:

$$y_i = \gamma WinPCT_i + \eta AllStarVotes_i + \theta PER_i + \mathbf{X_i}\beta + \epsilon_i \tag{2.1}$$

where y_i represents the outcome variable for matchup i. I examine two separate analyses with two different outcome variables: (i) weighted average ticket price on the secondary marketplace for matchup i, and (ii) starting TV rating (as measured by projected viewership) for matchup i. $\mathbf{X_i}$ represents a rich set of matchup-specific controls.[11]

To more flexibly understand the impacts of popularity, ability, and team quality on ticket prices and TV ratings, I conduct a "quasi-LASSO" reduced form analysis that performs separate kernel-density (LOESS) regressions for each of the residualized independent variables (popularity, ability, and team win percentage) on each of the residualized outcome variables (following Athey and Levin 2001). This procedure allows for the estimation of a "smooth" relationship between each independent variable and either prices or initial TV ratings, while accounting for an extremely rich set of controls with flexible functional forms. It is particularly useful when the independent variables are highly correlated with one another, which is the case with player popularity, player ability, and team quality.

There are three sets of estimating equations needed to conduct this analysis. First, equation (2.2) regresses independent variable $x_i \in \{WinPCT_i, AllStarVotes_i, PER_i\}$ on a rich set of controls, which includes flexible 5^{th} order polynomials for all controls $\neq x_i$ represented by $g(V_i)$ and $z(P_i)$. Additionally, a rich set of interactions of the controls is included in Γ_i.

$$x_i = g(V_i) + z(P_i) + \Gamma_i\eta + \epsilon_i \tag{2.2}$$

Equation 2.2 is estimated six times for each $x_i \in \{WinPCT_i, AllStarVotes_i, PER_i\}$ and whether the outcome variable is TV ratings or ticket prices.[12] So, in the estimating equation for $x_i = WinPCT_i$, $g(V_i)$ and $z(P_i)$ represent $g(AllStarVotes_i)$ and $z(PER_i)$, respectively.

Next, equation (2.3) regresses the weighted average ticket price (or initial TV rating) for matchup i on the same right-hand side as equation (2.2), namely:

$$y_i = g(V_i) + z(P_i) + h(W_i) + \Gamma_i\beta + \nu_i \tag{2.3}$$

where equation (2.3) is estimated when y_i denotes the weighted average ticket price for matchup i in one specification, and the initial TV rating for matchup i in a separate specification. This is, again, done for each $x_i \in \{WinPCT_i, AllStarVotes_i, PER_i\}$.

[11]Results Tables 2.5 and 2.6 denote the control variables used in each of the two analyses.

[12]These controls are the same as those found in Table 2.5 for ticket price as the outcome variable, and Table 2.6 for TV ratings as the outcome variable.

I then take the residuals from equations (2.2) and (2.3) and estimate a LOESS (kernel-density) regression of the vector of residualized y_i, denoted $\tilde{y}_i$, on the vector of residualized x_i, denoted $\tilde{x}_i$. The estimating equation for this analysis is as follows:

$$\tilde{y}_i = f(\tilde{x}_i) + \lambda_i \tag{2.4}$$

where $f(\cdot)$ is the kernel estimated for a LOESS regression (Cleveland 1979).

Difference-in-Differences and Event-Study Analyses

To obtain a plausibly causal effect of ticket price responses to a player's absence, I construct a counterfactual group that models ticket price movements without a player's absence, and compare those movements to the "treated" games, where a specific superstar player is announced to be out. This is important because there are underlying trends in ticket prices for NBA games that may bias the estimate of a player's absence if not controlled for with an appropriate counterfactual. There are several different ways of doing this – for example, I could use ticket listings from all other games on the same day and compare their price movements to ticket listings for the treated game on that day, which I denote the *same day counterfactual*. There are pros and cons to this method. On one hand, I am comparing games that occur during the same point in the season. On the other hand, there are a different number of games each day, which could limit the size of the counterfactual group, as well as completely different teams and markets involved each day.

A second way of constructing a counterfactual is through the *same team counterfactual*. This counterfactual compares games for the team of a specific superstar where that superstar was absent, to other games of that specific team not confounded by *any* superstar absences. For example, Golden State Warriors guard Stephen Curry missed the game on December 6, 2017 against the Charlotte Hornets in Charlotte. The same team counterfactual would consist of a subset of other Golden State Warriors games where Stephen Curry played *and no other superstar players were announced to be absent.*[13] So, the game where Warriors forward Kevin Durant was announced out due to injury against the Brooklyn Nets in Brooklyn on November 19, 2017 would not be included in this subset of potential counterfactual games. This counterfactual is preferred for the analysis since it controls for team-specific trends of ticket prices and their movements that may be common across many of their games, which is likely more valuable than the controls allowed by the same day counterfactual.

Primary Estimating Equations

This part of the analysis uses difference-in-differences (DID) and event-study estimations for each superstar player. Using the same-team counterfactual, the DID estimating equation is

[13]This only includes "qualifying games" for other superstars, as defined in the Data Characteristic section. Namely, I <u>*do*</u> include games that another superstar may have missed, but that weren't explicitly announced (for example, if another superstar was known to be out for the rest of the season prior to the treated game being analyzed).

written as follows:

$$\ln(Price_{ish}) = \beta_1 Absence_i + \beta_2 PostAnn_h + \beta_3 (Absence * PostAnn)_{ih} + \alpha_{is} + \alpha_h + \epsilon_{ish} \quad (2.5)$$

where $Price_{ish}$ represents the average listed price for tickets in section s for matchup i at hours-to-game h. So, an observation for the left-hand side variable would be the average listed price of tickets in section 201 for the Golden State Warriors vs. Houston Rockets matchup on October 17, 2017 listed on October 17, 2017 four hours before the game. $Absence_i$ is a binary variable = 1 if there was a superstar absence for matchup i, and $PostAnn_h$ is a binary variable = 1 if the announcement had already been made at hours-to-game h. Hours-to-game is used as the measure of time since matchups occur at different times during the day (e.g. 7:30pm EST or 10:30pm EST) and across days (e.g. October 16th vs. October 17th). Additionally, average ticket price trajectories are heavily dependent on the number of hours before gametime, as quantity of tickets available and prices on the secondary marketplace are very time-dependent (see Figures 2.1 and 2.2). Thus, for the Golden State Warriors @ Brooklyn Nets matchup on November 19, 2017, Kevin Durant was announced out of the game at 8:49am EST, which would correspond to 6 hours and 11 minutes to the game (which was at 3:00pm EST). The DID treatment coefficient is represented by β_3, which approximately represents the percentage change in ticket prices associated with a superstar absence, and is the primary coefficient of interest. Finally, α_{is} represents arena section-by-matchup fixed-effects, and α_h is an hours-to-game fixed effect. A log-level specification is preferred since prices cannot fall below zero, and thus the distribution of prices is censored.

Because I am attempting to determine the causal impact of a superstar absence on ticket prices, I estimate an event study to i) confirm parallel pre-trends in ticket prices for the treatment and counterfactual matchups, and ii) to determine the effect of a superstar absence on ticket prices in each time-period following the announcement (instead of just the post-announcement versus pre-announcement average effect that is obtained by the DID in equation (2.5)). This strategy provides compelling identification, since I am able to examine "within-matchup" changes in prices in response to plausibly exogenous announcements.

Employing the same-team counterfactual, the primary empirical specification can be written as follows:

$$\ln(Price_{isht}) = \sum_{t=-14}^{14\setminus\{-1\}} \mathbf{D}_t \mathbf{Absence}_{t,ih} + \alpha_{is} + \alpha_h + \epsilon_{isht} \quad (2.6)$$

$\mathbf{Absence}_{t,ih}$ is a vector of binary variables indexed by event-time t. Event-time t is in the half-hours-to-game unit, but is normalized to $t = 0$ based on the half-hours-to-game value when the announcement of a superstar's absence takes place. As is standard in event study estimations, each variable takes a value = 1 if the observation in the data refers to a matchup i where a superstar was absent and the observation of data corresponds to event-time t. $\mathbf{D}_t$ is a vector of estimated coefficients distinguishing the price differential between the treated game and counterfactual games at event-time t compared to an omitted period (which for this analysis will correspond to $t = -1$). As can also be seen in the estimating equation, I restrict the event-time horizon to $t = [-14, 14]$, where the left (right) binned endpoint coefficient represents the average treatment effects for all pre- (post-) periods not included in $t = (-14, 14)$. The dependent variable and fixed-effects remain identical to the simple DID estimating equation.

In addition to estimating an effect for each individual matchup that experienced a superstar absence, I also estimate an aggregate absence effect for each superstar, which requires a slightly more complex method of constructing the same-team counterfactual. Because each "treated" matchup for a specific player has a different announcement time in terms of hours-to-game, one cannot simply assign the same announcement time to all matchups in the counterfactual as was done in the individual matchup case. Rather, announcement times are randomly assigned for all matchups in the counterfactual by sampling from the pool of announcement times observed for the treated matchups. For example, James Harden was absent from six qualifying matchups that were analyzed (1/3/18, 3/11/18, 3/26/18, 4/11/18, 10/25/18, and 2/23/18), and was announced absent for these matchups at 47.5, 22, 26.5, 1.5, 33.5, and 2 hours-to-game, respectively. For each of these 6 treated matchups, I randomly pair a proportional number of counterfactual matchups based on the total set of eligible counterfactual matchups for the Houston Rockets, and assign the announcement time (in hours-to-game) of the treated matchup to each counterfactual matchup with which it was paired. In the case of Harden, there are 148 eligible, untreated matchups in the counterfactual group, so 4 treated matchups receive 25 counterfactual matchups each and the remaining two matchups receives 24 counterfactual matchups. Once the pairings are assigned, the same announcement time is assigned to the group and the announcement time of each grouping is normalized to 0. Finally, all groups are merged into a single table for the given player on which estimation is then performed. The estimating equations remain the same as in the case of the individual matchup analysis with one key difference – for matchups in the counterfactual, $PostAnn_h$ is determined based on the assigned announcement time within each grouping. To ensure robustness of the random counterfactual matchup-pairing algorithm, the aggregate-matchup analysis for each player is performed 3 times, each with a different random counterfactual pairing.

Finally, it is important to note that listed prices are used in this analysis. While a listed price does not necessarily indicate a seller's true willingness-to-sell (i.e. the reservation price of attending the game) since the choice of the listing price is a function of the prices of other listings of comparable seats, changes in listed prices due to superstar absences should reflect the combined effect of sellers' and buyers' lower value of attending the corresponding matchup. Therefore, the effect I estimate is the value loss associated with the absence of a specific superstar for the average NBA game attendee. In addition, I restrict the sample to tickets listings that eventually "sold," since these are listings that at some point reflect a market-clearing equilibrium price between sellers and buyers.

Identification Concerns

With any empirical estimation, there are concerns over identification of a causal estimate. In this estimation, I am inherently assuming that there are no omitted variables correlated with announcements that also affect ticket prices, namely:

$$\mathbb{E}[\epsilon_{isht}|Absence_{t,ih}, X_{isht}] = 0 \tag{2.7}$$

where X_{isht} represents the vector of covariates controlled for. However, because injury announcements are plausibly random (the occurrence of an injury is not predictable), and I only look at price movements 3 days prior to a matchup, there is only concern if a confounding event occurs that adjusts the price trajectory of a treated game differently than counterfactual games during this time horizon. One potential threat to identification is if an absence announcement of a player is correlated with having already made the playoffs and their team's seeding set. This may occur if the propensity to sit a superstar due to injury is higher once a team's playoff seeding is already known. In this case, it would be difficult to disentangle the price effect associated with a team having already made the playoffs and determined their seeding, and the price effect due to the injury of a superstar player.

While it is difficult to imagine important identification issues with respect to injury announcements, announcements about superstars being intentionally rested likely face a different set of concerns. First, decisions to rest superstar players may be dependent on several factors, for example the second night of back-to-back games or third game in four nights may exhibit a higher likelihood of superstars resting (e.g. Joel Embiid all of the 2017-18 season), competitiveness of the opponent, home vs. away games, etc. However, to the extent these characteristics are known prior to the three-days before a matchup, they would be accounted for in the matchup-specific fixed-effect.

2.5 Results

This section presents findings from the panel and quasi-LASSO analyses, as well as the DID estimation and event studies, including important heterogeneous impacts on WTP.

Panel Analysis

Tables 2.5 and 2.6 present the results of two separate estimations of equation (2.1): Table 2.5 using weighted average listed ticket prices (of all tickets that eventually sold) at the matchup-level as the dependent variable, and Table 2.6 using initial TV rating (projected total number of households watching) at the matchup-level as the dependent variable.

In Table 2.5, there are four different specifications presented. The first specification does not cluster the standard errors and does not account for a differential effect on ticket prices associated with a large absolute point spread and the home team being favored. One might think this would be important since the majority of fans attending a game are likely to be

Table 2.5: Impact of Player Popularity, Player Skill, Team Quality, and Parity on Ticket Prices

	Dependent Variable: log(Avg. Listed Price) (Matchup-Level)			
log(Ag. All-Star Votes)	0.2076***	0.2076***	0.2105***	0.1334***
	(0.0104)	(0.0185)	(0.0187)	(0.0194)
log(Ag. PER)	0.0758	0.0758	0.0760	0.1283
	(0.0797)	(0.0926)	(0.0910)	(0.0906)
log(Avg. Current Win PCT)	0.1770***	0.1770**	0.1946**	0.3842***
	(0.0483)	(0.0787)	(0.0855)	(0.0829)
Home Team Favored (HTF)	−0.0017	−0.0017	−0.0444	−0.0134
	(0.0213)	(0.0333)	(0.0276)	(0.0207)
Absolute Pt. Spread (APS)	0.0040*	0.0040	−0.0066	−0.0211**
	(0.0022)	(0.0042)	(0.0105)	(0.0092)
HTF*APS			0.0137	0.0240*
			(0.0129)	(0.0124)
Month FE	Yes	Yes	Yes	Yes
Time of Day	Yes	Yes	Yes	Yes
Day of Week	Yes	Yes	Yes	Yes
Holiday	Yes	Yes	Yes	Yes
Streak FE	Yes	Yes	Yes	Yes
Home Team FE	Yes	Yes	Yes	Yes
Away Team FE	No	No	No	Yes
TV Network FE	Yes	Yes	Yes	Yes
Clustered Robust SEs (Home Team)	No	Yes	Yes	Yes
Observations	2,330	2,330	2,330	2,330
R^2	0.6354	0.6354	0.6365	0.7353
Adjusted R^2	0.6233	0.6233	0.6242	0.7229

Note: *p<0.1; **p<0.05; ***p<0.01

supportive of the home team, and thus may exhibit differentially higher willingness-to-pay in cases when the absolute point spread is high but the home team is favored. Specification 2 clusters standard errors at the “Home Team” level and includes an indicator for whether the home team is favored. Specification 3 is identical to specification 2, except it accounts for differential effects of the absolute point spread on ticket prices, depending on whether or not the home team is favored. In this specification, a 1% increase in cumulative All-Star votes of all players playing in a matchup leads to a 0.2% increase in ticket prices. This effect is almost identical in magnitude to the impact of a 1% increase in the average combined winning percentage of the two teams playing in the matchup.

Specification 4 is identical to specification 3, with the exception that it includes "Away Team" fixed effects in addition to "Home Team" fixed effects, and is the preferred specification. One can see that there are substantial adjustments to several of the estimates, in particular the magnitude of the "Avg. Current Win PCT" variable approximately doubles, the "Absolute Point Spread" now has a statistically significant negative impact on listed prices, and when the home team is favored, absolute point spread does not have an impact on ticket prices (i.e. home fans want to see their team win, regardless of expected competitiveness). Under this specification, a 1% increase in cumulative All-Star votes in a matchup leads to a 0.13% increase in ticket prices. Because the popularity metric used is the cumulative All-Star votes of all players *actually playing* in a given matchup, including Home Team + Away Team fixed effects suggests that the coefficient on the popularity variable relies on *changes* in the lineups of teams across a season to drive residual variation in the popularity metric. This variation is similar to the identifying variation used in the DID and event study estimations, and allows for useful comparisons across the two different estimation approaches.

Table 2.6 presents the impact of each of these factors (omitting the "Home Team Favored" binary variable) on TV ratings for nationally televised games. There are three different specifications: the first includes all nationally-televised games from the 2017-18 and 2018-19 seasons, while specification 2 includes only regular season games and specification 3 only playoff games. Each of these specifications uses an "aggregate team value" continuous control variable to account for the number of people that may be expected to watch independent of other important factors.[14] These team values are calculated each year by Forbes, and are a good indicator of the total size of each team's fanbase (Badenhausen and Ozanian 2019). One can see that aggregate popularity, average current team quality, and aggregate team value are the only statistically significant estimates. The findings suggest that for a 1% increase in the cumulative number of All-Star votes in a matchup, initial rating increases by 0.12%, and similarly for a 1% increase in the average current win percentage of the two competing teams, ratings increase by nearly 0.28%. Additionally, in limiting the sample to regular season games (about 70% of the sample), this estimate increases to 0.16%, suggesting that player popularity may be a more important factor in the regular season than the playoffs. In fact, the estimate of the coefficient on the aggregate popularity variable becomes insignificant when subsetting the set of games to include only playoff matchups. One potential explanation for this is that playoff games have an "elimination" component, and so encompass a different

[14]Since these are national audiences, a home team fixed effect does not make as much sense in these specifications as it does in the context of the ticket price analysis (since there are geographic preferences). Including a dummy for each team present in a matchup leads to insignificant point estimates for all variables, likely because of the insufficient power associated due to the relatively low number of nationally-televised games.

Table 2.6: Impact of Player Popularity, Player Skill, Team Quality, and Parity on Initial TV Ratings

	Dependent Variable: log(Total Proj. Viewers) (Matchup-Level)		
	All Games	Reg. Season Only	Playoffs Only
log(Ag. All-Star Votes)	0.1197***	0.1580***	−0.0393
	(0.0321)	(0.0276)	(0.0742)
log(Ag. PER)	−0.1280	−0.0049	0.0381
	(0.1564)	(0.2018)	(0.2411)
log(Avg. Current Win PCT)	0.2756**	0.0838	0.6932***
	(0.1274)	(0.1497)	(0.2501)
Absolute Pt. Spread (APS)	−0.0007	−0.0020	0.0035
	(0.0049)	(0.0051)	(0.0115)
log(Ag. Team Value)	0.1753**	0.1112	0.3631*
	(0.0721)	(0.0706)	(0.1950)
Month FE	Yes	Yes	Yes
Day-of-Week FE	Yes	Yes	Yes
Time-of-Day FE	Yes	Yes	Yes
Streak FE	Yes	Yes	Yes
TV Network FE	Yes	Yes	Yes
Dbl Header FE	Yes	Yes	Yes
Holiday FE	Yes	Yes	Yes
Playoff Gm FE	Yes	Yes	No
Clustered Robust SEs (Home + Away)	No	Yes	Yes
Observations	477	329	148
R^2	0.7448	0.6494	0.7084
Adjusted R^2	0.7155	0.5922	0.6068

Note: *p<0.1; **p<0.05; ***p<0.01

viewership utility function that downweights player popularity. Note that I reset each team's record for the playoffs, and so the "Avg. Current Win PCT" variable simply reflects that better teams move to future rounds by construction, where each subsequent playoff round experiences higher viewership. All specifications are clustered at the "Home Team + Away Team" level.

Next, I present the results from the quasi-LASSO estimation, as laid out in equations (2.2) – (2.4). This procedure is conducted for three different, yet correlated, independent variables: (i) combined average current win percentage (a measure of team quality), (ii) cumulative All-Star votes of all players who played in a matchup (a measure of popularity), and (iii) the cumulative player-efficiency rating (PER) of all players who played in a matchup

Figure 2.5: Quasi-LASSO Results Figures

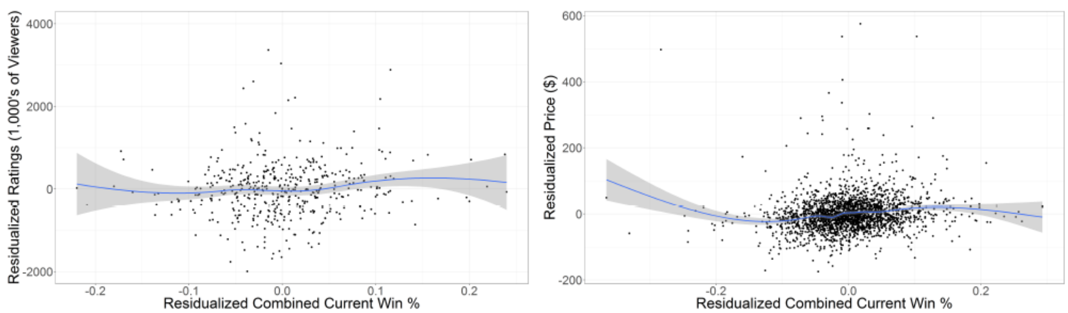

a. Impact of Residualized Combined Current Win % on (left pane) Residualized TV Ratings and (right pane) Residualized Ticket Prices

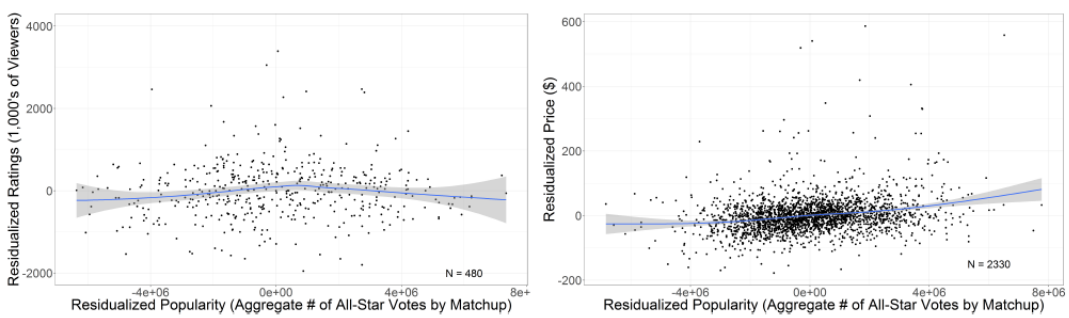

b. Impact of Residualized Cumulative # of All-Star Votes on (left pane) Residualized TV Ratings and (right pane) Residualized Ticket Prices

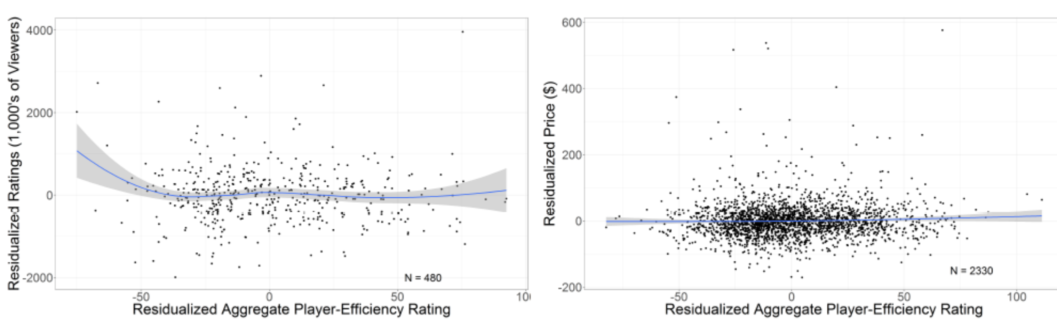

c. Impact of Residualized Aggregate Player-Efficiency Rating (PER) on (left pane) Residualized TV Ratings and (right pane) Residualized Ticket Prices (a measure of ability), and two different dependent variables: (i) weighted average ticket price at the matchup-level, and (ii) initial TV rating (as measured by the projected number of households watching) at the matchup-level. Figures 2.5a, 2.5b, and 2.5c present the results for residualized team quality, popularity, and ability, respectively, with the left pane in each figure corresponding to the impact on residualized TV ratings and the right pane the impact on residualized ticket prices.

One can see that within the primary support of the residualized independent variables, both team quality and player popularity meaningfully affect ticket prices and TV ratings, as seen in Figures 2.5a and 2.5b. In the case of player popularity, the positive relationship on TV ratings is concentrated towards the bottom of the distribution, while in ticket prices there is a clear convex relationship throughout the entire support of the data. The relative magnitudes and significance of the relationships in Figures 2.5a-2.5c are quite similar to those in Tables 2.5 and 2.6.

Difference-in-Differences

Figure 2.6 presents the results of the DID estimation as seen in equation (2.5). Each estimate reflects the average treatment effect on listed ticket prices from the entire sample of analyzed absences for each qualifying superstar. The confidence intervals presented are at the 95% level. Importantly, I only include players where pre-trends in ticket prices between the counterfactual and treated matchups prior to a superstar's injury announcement are parallel, satisfying the identifying assumption that the DID estimate is causal. In examining the results, the reduction in prices due to absence announcements in percentage terms is highest for Dwyane Wade, Kemba Walker, and Dirk Nowitzki, all resulting in 14-16% reductions in prices associated with their absence announcements.

Figure 2.7 exhibits these declines in level price reductions instead of percentage terms. Because the average price of Los Angeles Lakers' and Golden State Warriors' tickets is quite high, absences for LeBron James and Stephen Curry result in the largest magnitude decrease in ticket prices at $42 and $29 per ticket, respectively. There are a number of other players whose absences lead to economically meaningful and statistically significant price reductions, including Dwyane Wade, Dirk Nowitzki, Luka Doncic, Paul George, Kemba Walker, and Kawhi Leonard, each of whom lead to price reductions between $7-$26 per ticket. Somewhat surprisingly, there are no statistically significant price reductions associated with James Harden's or Giannis Antetokounmpo's absences, who are the reigning MVPs from the previous two seasons.

To gain a better understanding of how player popularity factors into the magnitudes of these estimates, Figure 2.8 visualizes the relationship between player absence impact ($ per ticket) on their maximum single-season All-Star vote total over the previous two seasons, and fits a quadratic approximation to showcase the general shape. One can see that there is a convex relationship between each player's impact and their fan votes, again supporting

Figure 2.6: Difference-in-Difference Results for Superstar Absences (Percentage Change in Prices)

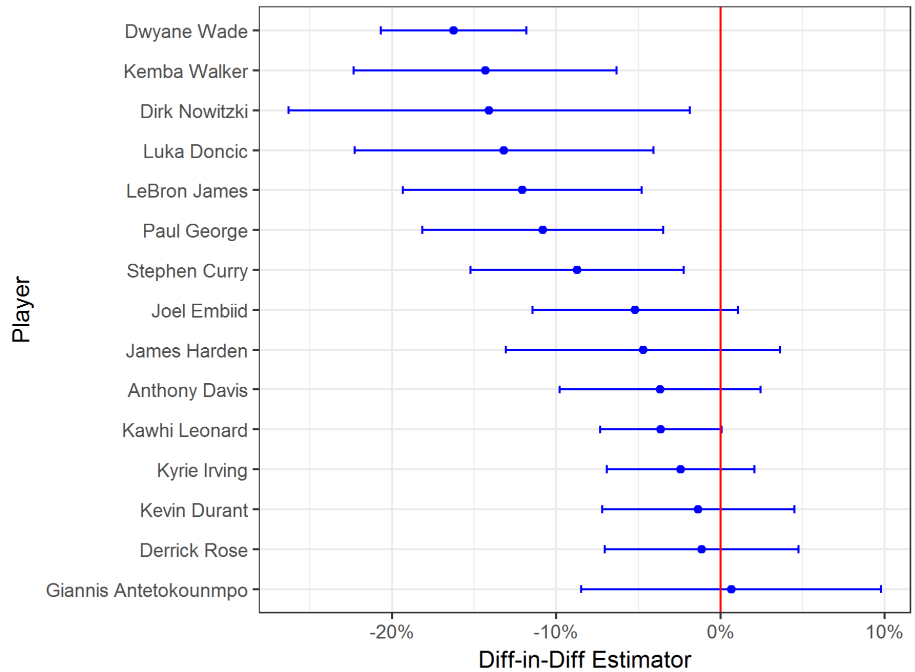

Figure 2.7: Difference-in-Difference Results for Superstar Absences (Level Change in Prices)

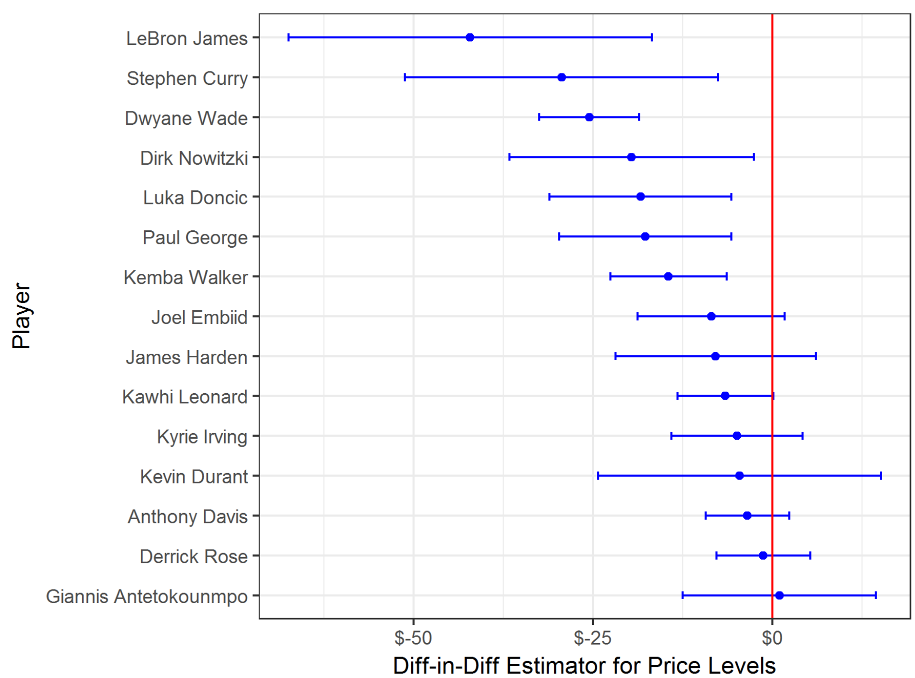

Figure 2.8: Difference-in-Difference Results by All-Star Votes

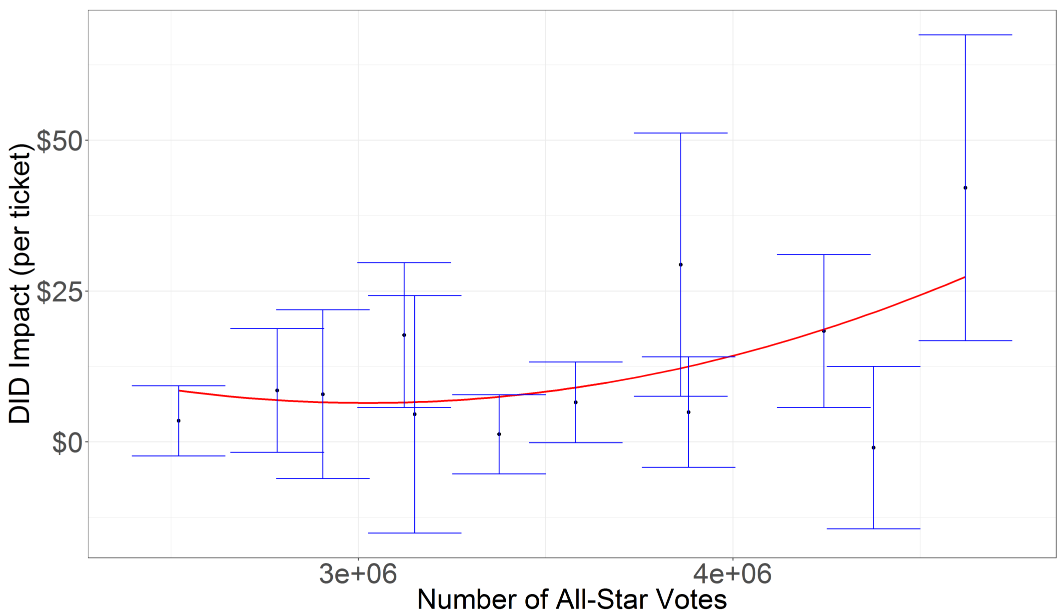

the economic theory of superstars.[15]

To capture different measures of popularity, and in particular "career-long" popularity (which weights legacy players relatively more than recently popular stars), Figure 2.9 plots the player-specific ticket price impacts on total number of career All-Star appearances (left) and championships won (right). While both exhibit a relatively convex relationship, the left pane of Figure 2.9 is slightly more U-shaped, which is largely due to the significant global popularity of Luka Doncic, who played internationally for several seasons as a highly-renowned teenager before coming to the NBA but has not had a chance to accumulate All-Star appearances.

Event Studies

The event study results present coefficients for each of the 30-minute intervals before and after an absence announcement takes place. Figure 2.10 shows the results for the top three impact players with respect to ticket price declines as a result of their absences, again using

[15] This figure omits Dwyane Wade, Dirk Nowitzki, and Kemba Walker. Wade and Nowitzki were "legacy picks" by the NBA commissioner to take part in the All-Star game because of their career achievements, and thus had lower fan vote totals but large impacts on prices when they missed games. Despite Walker's significant impact on prices, his vote total was half the size of the next lowest vote-getter, and only missed two games over the course of these two seasons (both of which were at home), which was the lowest missed game total among all players analyzed here. His somewhat large effect is likely due to a small sample size of missed games and the fact that those games were missed in one of the NBA's smallest markets (Charlotte). The graph including these players is found in the Appendix in Figure B.2.

Figure 2.9: Difference-in-Difference Results by All-Star Appearances (left) and Championships (right)

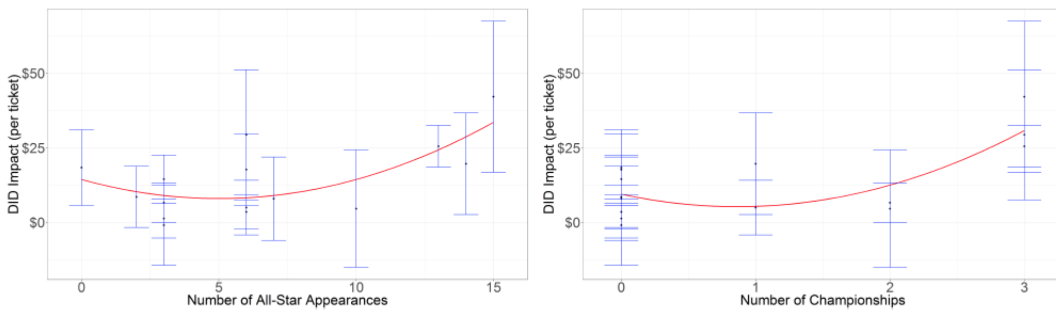

Figure 2.10: Event Study Results for Top Impact Superstars

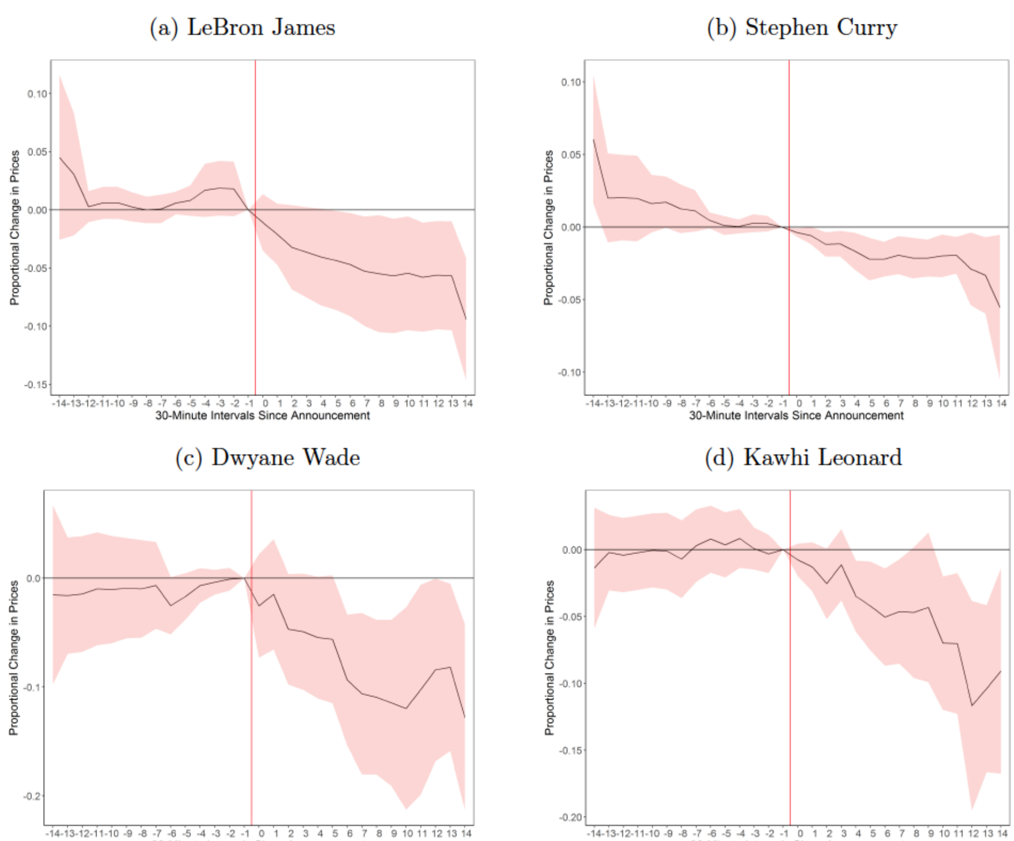

Figure 2.11: Kevin Durant vs. Stephen Curry Absence Impacts

(a) Kevin Durant (b) Stephen Curry

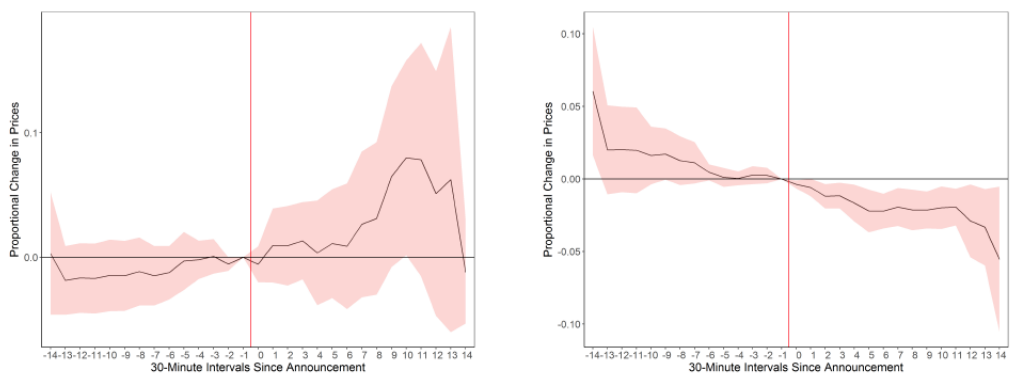

the aggregate estimation, and Kawhi Leonard, who is the reigning NBA Finals MVP.[16] Each point on the graph can be interpreted as the differential effect on listed ticket prices of a superstar absence announcement on the treated group vs. the counterfactual group. Coefficients statistically insignificantly different from zero prior to an absence announcement, which is indicated by the vertical red line, suggest that parallel pre-trends in ticket prices hold in each of these cases. The event study estimates exactly when prices change as a result of an announcement. One can see that there is a slight delay in the full responsiveness of listed ticket prices to the announcement of a superstar's absence – typically the effects are smaller closer to the announcement time and larger further away. This is intuitive, as many sellers and buyers do not have immediate access to announcement information or the ability to immediately change their listing on the secondary marketplace. The endpoints are binned at -7 and $+7$ hours in event-time with respect to when the announcement occurs (at $t = 0$).

In Figures 2.6 and 2.7, one can see that Kevin Durant's absence announcements on average lead to no statistically significant ticket price adjustments. This is particularly interesting given that there is a meaningful reduction for his teammate Stephen Curry's absences. Figure 2.11 presents the event study results for Kevin Durant and Stephen Curry. From a ability standpoint (measured by player efficiency rating or value over replacement player), Kevin Durant and Stephen Curry were nearly identical during the 2017-18 and 2018-19 seasons. However, Curry's popularity with NBA fans as "the best shooter of all-time" and his unique style of play may make him a more desirable player to watch from an entertainment standpoint.

[16]The event study results for the remaining eligible players are presented in the Appendix.

Figure 2.12: Difference-in-Differences Estimator by Home vs. Away Matchup Absence (Level Change in Prices)

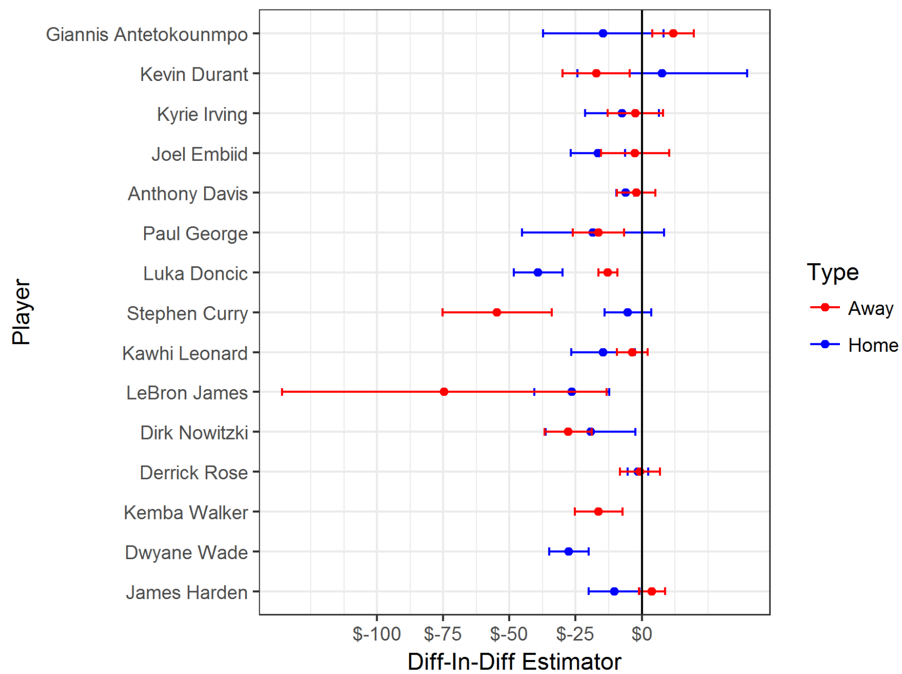

Heterogeneity Tests

The final set of analyses examines heterogeneity with respect to types of games superstars are absent for. Two sets of analyses are presented here – first, the differential impact on ticket prices for home games vs. away games for each qualifying player. Next, I examine the impacts of market size of the home team, matchup competitiveness, and number of other superstars present in a matchup on the absence effect of each qualifying superstar player.

Figure 2.12 presents two distinct DID estimators (exhibiting level price changes) for each player: one for home games missed and another for away games missed.[17] One can see there are some striking differences in effects for certain players. For example, Stephen Curry and LeBron James' absence effects are sizably larger and much more negative for away absences than for home absences. LeBron's average away-game effect is $75/ticket, while Stephen Curry's is $55/ticket. This suggests that the value of these players in away arenas is higher than in their home arena, likely because they only play in opposing arenas at most two times per year, and so there is a geographic scarcity effect of not being able to substitute towards a different game. On the other hand, Luka Doncic and, to a lesser extent, James Harden both exhibit the opposite effect, where their absences are more meaningful for home

[17]Note that Kemba Walker was not absent for any qualifying home games, and Dwyane Wade was not absent for any qualifying away games.

games than for away games. This is also quite intuitive – both of these players are not just entertaining to watch, but without them their teams become much less competitive and much more likely to lose a game. The same argument could be made for LeBron James' impact on the Lakers, who also exhibits a negative effect for home game absences, but his transcendent superstardom leads to an even larger away absence effect. Home fans value the competitiveness of their team, and thus the absence of these stars substantially reduces their team's chances of winning. Figure B.5 in the Appendix exhibits these changes in percentage point terms.

I also conduct heterogeneity tests analyzing the differential effect of absence announcements depending on the competitiveness of the matchup (as measured by the absolute point spread), the total number of other starting-caliber superstars present, and the market size of the home team. This analysis relies on a triple differences estimation, where the DID treatment variable is interacted with the relevant metrics for competitiveness, total star power, and market size, respectively, in separate estimations. The results of these analyses suggest there are no meaningful or statistically significant relationships between ticket prices and these additional differentiators. One potential explanation for this finding is that there are not enough events of different types to estimate a robust statistical relationship. Future work should aim to incorporate additional absences to add to the power of such an estimation.

2.6 Discussion and Player Salaries

The results from the panel and difference-in-differences/event-study approaches yield largely consistent findings at an aggregate level, despite differences in the estimation approach. The panel analysis found that in the presence of a player holding 100% of the average cumulative All-Star votes for a specific matchup, ticket prices and TV ratings were on average 12.0-13.3% higher. This is comparable to the range found in the DID and event-study approaches, where ticket price reductions due to superstar player absences were on the order of 4-16% on average.

For context, LeBron James averaged just over 3.6 million fan votes over the 2017-18 and 2018-19 seasons, which corresponds to approximately 95% of the average aggregate number of All-Star fan votes of all players in a matchup (3.8 million). In other words, LeBron's average fan All-Star vote total is just below the total number of All-Star votes of all players in an average game. Using the results from the panel analysis, the presence of LeBron alone results in a 11.37-12.6% increase in ticket prices and TV ratings. The DID analysis yields a very similar result – the absence of LeBron leads to a 13% average reduction in ticket prices. This implies millions of dollars in welfare lost for each of these matchups, and tens or even hundreds of millions of dollars lost across all superstar absences over the course of a season.

For the remaining superstar players analyzed, Table 2.7 presents their projected season-level impact under the difference-in-differences method (column 3) and panel analysis method (column 4). The impacts in column 3 are based on the results in Figure 2.7, while the impacts in column 4 are using the results from Tables 2.5 and 2.6. One can see that from ticket sales

Table 2.7: Season-Level Difference-in-Differences and Panel Analysis Player Impacts

Player	All-Star Votes	DID Impact (Millions of $)	Panel Impact (Millions of $)
LeBron James	4,620,809	69.06	37.87 - 42.05
Giannis Antetokounmpo	4,375,747	-1.57	35.86 - 39.82
Luka Doncic	4,242,980	30.13	34.77 - 38.61
Kyrie Irving	3,881,766	8.09	31.81 - 35.32
Stephen Curry	3,861,038	48.18	31.64 - 35.14
Kawhi Leonard	3,580,531	10.76	29.34 - 32.58
Derrick Rose	3,376,277	2.09	27.67 - 30.72
Kevin Durant	3,150,648	7.50	25.82 - 28.67
Paul George	3,122,346	29.03	25.59 - 28.41
James Harden	2,905,488	12.98	23.81 - 26.44
Joel Embiid	2,783,833	13.99	22.81 - 25.33
Anthony Davis	2,520,728	5.73	20.66 - 22.94
Dwyane Wade	2,208,598	41.86	18.1 - 20.1
Kemba Walker	1,395,330	23.74	11.43 - 12.7
Dirk Nowitzki	394,622	32.19	3.23 - 3.59

Note: These estimates represent the season-level monetary impacts each player had based on the difference-in-differences (DID) and panel estimations. For the DID estimates, this meant multiplying by 20,000 people on average per arena and 82 games over the course of a season. The panel estimates were estimated via a log-log specification, so the season-level impacts were determined using the player-specific % total of the average cumulative number of All-Star votes present for all players in a matchup and the average listing price across all games.

alone, the DID impacts of the presence of superstars range from millions to tens of millions of dollars over the course of a season, with LeBron James as the leader at $69 million. The panel impacts are slightly larger on average, despite the maximum value (which corresponds to LeBron James) being slightly lower at $42 million. The slightly larger estimates on average may be due to a combination of factors, including potentially high transaction costs of responding on the secondary market in the DID case, or omitted variable bias in the panel analysis.

However, there are a couple of players significantly underestimated by the panel analysis compared to their DID impacts. During the 2017-18 and 2018-19 seasons, Dwyane Wade and Dirk Nowitzki were not playing at an All-Star level, but since the 2018-19 season was known (or in the case of Dirk, almost certainly known) to be their last, they were "hand-picked" by the NBA commisioner to be All-Stars that year. It appears that the DID impact may represent some sort of "legacy" effect representing the aggregate popularity of these players over the entire course of their careers, and the desire to see them play before they officially retired.

Player salaries are an important metric in considering their overall value. In looking at these results, one can see that the salaries of many of the most popular players, including LeBron James ($35.6 million), Stephen Curry ($41 million), Dwyane Wade ($2.4 million), Dirk Nowitzki ($5 million), and Luka Doncic ($7.5 million) appear to be lower than their direct impact on ticket values. This is even more staggering when considering that their impact reaches far beyond tickets, including television ratings, league merchandise, and other league-level sponsorship deals. The NBA generates an estimated $8 billion in annual revenue, which comes from approximately equal shares of ticket sales, television deals, and league sponsorships and merchandise (Sports-Illustrated 2014; Investopedia 2019; Sponsorship.com 2018; Statista 2019). Therefore, when accounting for benefits in each of these domains, it becomes clear that most All-Star players' salaries may reflect an underestimate of their actual value added. For instance, extrapolating LeBron James' season-long impact on ticket value estimated from the DID analysis ($69 million) to the two other categories, James is worth close to $210 million per year to the NBA, assuming his proportional value in ticket sales is similar to television deals and league sponsorships/merchandise. This is likely still an underestimate since it fails to account for both playoff ticket and ratings values, as well as consumer surplus. Between his salary and endorsement deals, LeBron's actual annual income is an estimated $94 million (NBC 2019).

Figure 2.13: Annual Player Salary by Fan All-Star Vote Total

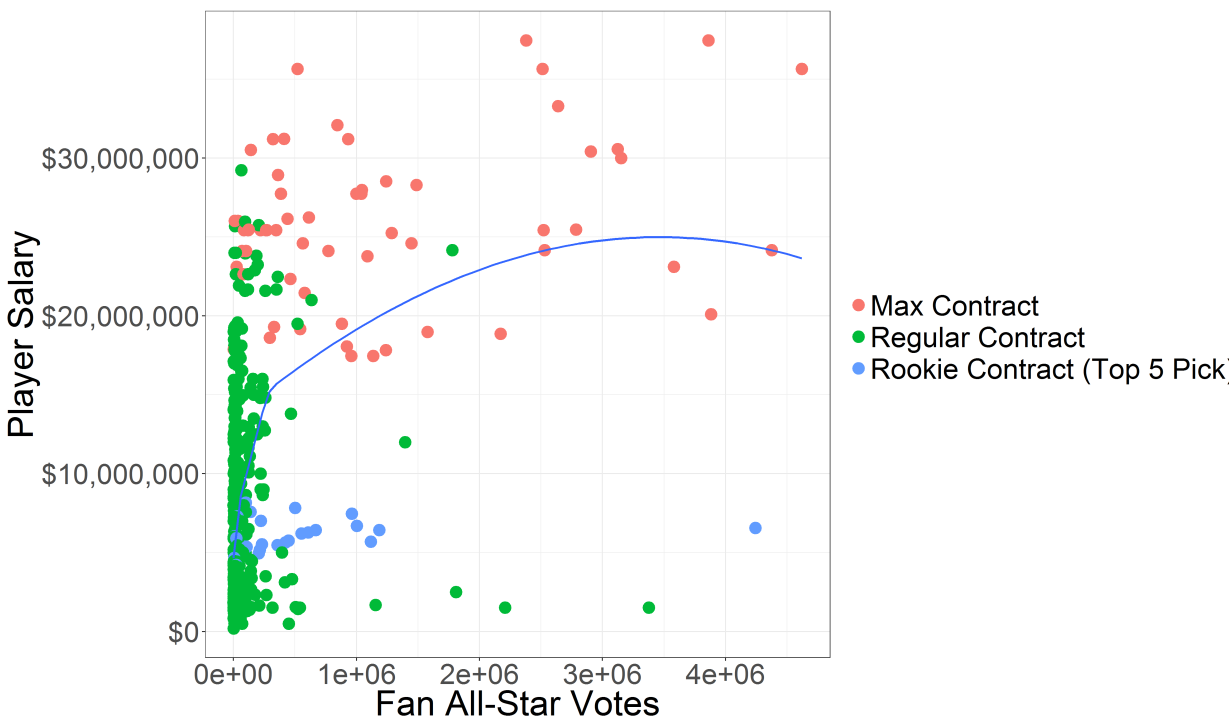

The NBA places restrictions on the maximum amount players can be paid (for a variety of reasons), and so salaries for the league's superstars are "censored," as depicted by Figure 2.13. Figure 2.13 plots the relationship between All-Star fan votes and salaries of each applicable NBA player over the 2017-18 and 2018-19 seasons. One can see that there is a concave relationship between salaries and fan votes. Additionally, the proportion of players under these maximum-level contracts becomes much higher in the high popularity domain of the data.

Table 2.8 reinforces this finding. The data is split into seven different total fan vote bins. Each data point corresponds to a specific player-season combination (e.g. Kevin Durant in 2017-18), and each of these is classified as a "max contract," "rookie contract (top 5 pick)," and "regular contract." I account for rookies selected in the top 5 of the NBA draft still on their initial contract (which lasts as long as four years) since some of these players become extremely popular before they are able to sign their next maximum-level contract. The final column indicates the proportion of players holding max contracts and top 5 picks still on their rookie deals within each popularity grouping. It is clear that as popularity increases, this proportion increases substantially, equaling one in the highest popularity bin.

Table 2.8: Salary Censoring – Number of Player-Season Combinations by Contract Status and All-Star Fan Votes

Fan Votes	# Max	# Rookie (Top 5 Picks)	# Regular	% Max or Rookie (Top 5 Picks)
3.5 Million +	5	1	0	1.00
2.5-3.5 Million	8	0	1	0.93
1.5-2.5 Million	3	0	3	0.81
0.5-1.5 Million	21	8	7	0.81
0.25-0.5 Million	11	3	13	0.71
0.1-0.25 Million	5	7	47	0.50
0-0.1 Million	7	11	431	0.15

Note: An example of a "Player-Season" combination is Stephen Curry in the 2018-19 season. I examine the 2017-18 and 2018-19 NBA seasons. All combinations featuring a player playing less than half of the regular season (41 games) are omitted. The maximum vote tally in the sample was 4,620,809 by LeBron James during the 2018-19 season. Salary information was collected from `ESPN.com`, maximum contract information from `HoopsHype`, and top 5 picks in preceding drafts from `Basketball Reference`.

Finally, Table 2.9 shows an estimation of the impact of number of All-Star fan votes on salary. Specifications 3-5 control for the ability metric of these players (player-efficiency rating), and specification 5 presents the results of a log-log estimation. The findings in specifications 2-4 back-up those presented in Figure 2.13, which depicts a concave relationship between popularity and salary. This paper does not aim to investigate the normative implications of these findings nor discuss how the NBA should treat these types of contract restrictions. It is clear that a maximum salary restriction may be needed to maintain important parity across the league, and not disadvantage certain markets more drastically than others. On the other hand, it is a direct transfer of value from players to owners and other league stakeholders.

Table 2.9: Relationship between Player Salary and Popularity

	Dependent Variable: Player Salary				
Total Fan Votes (100,000)	636,663.50***	1,500,027.00***	1,016,613.00***	985,048.70***	
	(50,718.35)	(138,859.50)	(146,927.20)	(150,643.50)	
Total Fan Votes Squared		−27,157.50***	−18,812.76***	−18,612.53***	
		(4,087.38)	(4,052.53)	(4,058.33)	
log(Total Fan Votes)					0.17***
					(0.02)
Skill (PER) Control	No	No	Yes	Yes	Yes
Skill (PER) Quadratic Control	No	No	No	Yes	No
Log-Log Specification	No	No	No	No	Yes
Observations	592	592	592	592	592
R^2	0.21	0.27	0.33	0.33	0.24
Adjusted R^2	0.21	0.26	0.33	0.33	0.24

Note: *p<0.1; **p<0.05; ***p<0.01

2.7 Concluding Remarks

This paper measures the economic value of superstars using data on spectator WTP. I estimate flexible panel and DID/event study models that rely on plausibly exogenous ticket price changes associated with player absence announcements for NBA games. The results suggest that a 1% increase in the aggregate popularity of a matchup (as measured by the total number of All-Star fan votes of all players playing) increases ticket prices and TV ratings by 0.12-0.13%, while increases in player ability (as measured by PER) have no significant impact. In the DID and event study analyses, I find that absences of several superstars, including some of the most popular like LeBron James, Stephen Curry, and Dwyane Wade, do have statistically significant and economically meaningful impacts, ranging from a 4-16% ($7-$42) reduction in the average ticket price, and that player popularity predicts these impacts in a convex manner. Additionally, I examine the differential in superstar absence effects for home vs. away games, finding that the most popular players, like LeBron James and Stephen Curry, exhibit much larger away game absence effects – prices fall an average of $75/ticket for James' absences and $55/ticket for Curry's absences.

This paper provides a novel methodology and framework to causally estimate the economic value of superstars, particularly in the context of sports and entertainment. The findings explicitly connect the theory of superstars to player marketability and branding, and how the sheer popularity of an individual impacts consumer demand. It is the first study to provide quantifiable, plausibly causal evidence that there are significant reductions in welfare when superstars miss events. It also emphasizes heterogeneity in the financial importance of superstar allure for different types of franchises, especially those experiencing low ticket sell-out rates. In addition, franchises may experience significant financial returns to investing in popular players that do not necessarily have substantial impacts on team performance. Finally, the findings can help inform national TV networks' strategy about which games to televise, and in the case of superstar absences, whether to "flex" to a different matchup.

Future work should aim to expand on the scope of superstar impacts, including examining other leagues, entertainment and arts sectors, and other economically important sectors by using dynamic pricing and other mechanisms that allow the use of high-quality exogenous variation to estimate well-identified impacts. Within sports specifically, there are interesting opportunities to examine different behavioral responses in secondary ticket marketplaces based on the timing of supersar absence announcements and estimated duration of the player's absence. One potential expansion would be to study the impact of long-term injuries on ticket prices. For example, when a player is guaranteed to be out for the remainder of the season, how does this affect the dynamics of ticket price adjustments corresponding to all affected games? This may lead to different results and implications, since sellers and buyers have more time to adjust and process this information. Additionally, do ticket prices for "near-term" games associated with a long-term absence announcement adjust differently than games further in the future? In a similar vein, what is the impact of "uncertainty" associated with some players' timelines in returning from injury or rest on ticket prices for future, potentially impacted games? Each of these scenarios may impact how consumers respond, and provide insight into consumer time preferences and behavior under uncertainty.

Chapter 3

Soda Wars: The Effect of a Soda Tax Election on University Beverage Sales[1]

3.1 Introduction

With the current trend of sugar consumption, exercise, and dietary habits, it is estimated that 40% of Americans born from 2000 to 2011 will get diabetes in their lifetimes, with the percentages for African-American women and Hispanics placed even higher at 50% (Gregg et al. 2014). While researchers and industry participants agree on the health dangers of sugar, and in particular sugar-sweetened beverages (SSB), there is disagreement on how to design laws and policies to change behavior. According to the Center for Disease Control (CDC), SSBs are defined as drinks with added sugar, which includes sweeteners like brown sugar, raw sugar, and corn syrup, among others.[2] Policy proposals to address SSB consumption include bans (James et al. 2004; Fernandes 2008; Huang and Kiesel 2012), taxes (Brownell and Frieden 2009), nutrition education programs (James et al. 2004; Fernandes 2008), and warning labels on SSBs advising the dangers of obesity, diabetes, and tooth decay (Roberto et al. 2016). This raises the empirical questions: how do consumers react to such policies and are there differences between direct regulations and informational campaigns? This paper examines how consumers alter their purchasing behavior due to the campaign attention and election outcome of a local excise tax aimed at curbing SSB consumption.

[1]This work was co-authored with Rebecca Taylor, Sofia B. Villas-Boas, and Kevin Jung. The published version can be found in *Economic Inquiry* at https://248bd2cf-cf74-4b88-a4e6-3076b3a73322.filesusr.com/ugd/42a8fe_de97d13ddd9b4c4f91077b772166f611.pdf. We thank the editor Wesley Wilson and two anonymous referees and acknowledge seminar participants at UC Berkeley, as well as participants at the 2016 *Agricultural and Applied Economics Association* Annual Summer Meeting and the 2017 *Future of Food and Nutrition* Annual Conference, for helpful comments and suggestions. We thank University Dining staff in granting us access to university sales data and for providing institutional knowledge and data support. Special thanks to Tracy Ann Stack, Samantha Lubow, Jaylene Tang, and Jayson Lusk for discussions and feedback.

[2]Added sugar is distinctly different from naturally-occurring sugar, and a beverage with naturally-occurring sugar is not classified as a SSB.

We take advantage of a tax policy change—referred to as Measure D—in the city of Berkeley, California. Measure D imposes a penny-per-fluid-ounce tax to be paid by distributors of SSBs. The aim of the policy is to lower the consumption of SSBs, or if demand is deemed to be unresponsive,[3] to raise tax revenues which could fund nutritional programs and education. On November 4, 2014, Measure D was put to a vote and passed with 75% of voters in favor. An aggressive campaign war preceded this vote, dubbed "Berkeley vs. Big Soda." This campaign cost $3.4 million, with roughly $1 million spent in favor of Measure D and $2.4 million spent against it.[4]

The specific objective of this paper is to examine how consumers reacted to the Measure D media campaign and to the election outcome. There is evidence suggesting that highlighted news coverage can lead to sharp information updates (Huberman & Regev, 2001; Lusk, 2010), and local elections can publicly reveal to individuals previously unknown information about the preferences of their neighbors and peers (Goldstein et al., 2008). Investigating whether campaigns and elections also lead to behavioral changes—whether through information or social norm channels—has important policy implications about how and where SSB policies will be effective in altering SSB consumption. This is especially true if campaigns and elections cause behavioral changes that happen before the policy implementation.

Our study employs detailed data from university retailers in Berkeley, consisting of monthly beverage sales. We use a difference-in-differences (DID) strategy to measure the change in ounces of soda purchased against untreated products (i.e., comparable control beverages) and untreated months (i.e., the pre-campaign period).[5] Additionally, we estimate an event study model to test the identifying assumption of parallel trends in the pre-campaign period. We verify that soda would have evolved in the same trend as other beverage products had there not been a tax campaign or affirmative election outcome.

There are two major advantages of using this empirical setting for our research design. First, products offered, as well as the promotional effort and posted prices, are uniform across campus retail locations. Second, we know exactly when and by how much the SSB tax is passed on to consumers, and do not have to infer the pass-through from the data. Since SSB taxes are often levied on the distributors of SSBs—who have a choice on how much of the tax they will pass on to consumers[6]—there is an empirical literature asking who bears the

[3]There is suggestive evidence that in the first month of the tax, tax revenues increased by $116,000, which is consistent with demand having not responded in an elastic fashion to the one-cent-per-ounce increase in price (*The Daily Californian.* "1st Month of Berkeley 'Soda Tax' Sees $116,000 in Revenue," May 19, 2015. Online, *accessed May 21, 2016*).

[4]*Berkeleyside.* "Around $3.4M spent on Berkeley soda tax campaign," Feb. 5, 2015. Online, *accessed May 21, 2016.*

[5]We focus on soda instead of SSB products more broadly because the media campaign focused on soda.

[6]One reason to tax distributors instead of customers is to make the price change more salient. There is a growing literature providing empirical evidence that consumers have an attenuated response to non-salient costs. With a labeling experiment, Chetty et al. (2009) find that the sales of taxable products at a grocery store are reduced when their tax-inclusive price is displayed in addition to the tax-exclusive price. Thus by taxing distributors of SSBs, if the tax is passed on to consumers, this will affect the displayed price and be more salient than a tax at the point-of-sale.

SSB tax burden.[7] In the university setting, we know exactly when and by how much the campus retailer adjusts prices. In particular, due to the costs of changing prices, campus retailers chose not to pass-through the tax to consumers for a year after receiving the tax invoices. Thus, we are able to look at how soda demand changes on-campus when the prices off-campus react to the tax implementation, yet the prices on-campus remain unchanged.

Our primary findings reveal no significant difference in on-campus retail soda sales compared to control beverage groups during the campaign period before the election (July 2014-October 2014). Conversely, soda sales fell significantly compared to control beverage groups in the period immediately following the election (November 2014-February 2015), decreasing by between 10-20% compared to pre-campaign levels. We also find that on-campus soda sales continued to fall when the tax was implemented in the city but not on campus (March 2015-July 2016)—decreasing by 18-36% compared to the pre-campaign period—and remained at this depressed level after the tax implementation on campus (August 2016-December 2016). Additionally, we find evidence that consumers substituted towards diet beverages as a result of the election outcome.

It is important to note that the university retailers in our analysis are not representative of the average U.S. retail outlet, especially in terms of clientele, and this could have large implications for whether our results will generalize to other locations. For this reason, we supplement the on-campus analysis with an analysis of beverage sales off-campus—at drug stores in the city of Berkeley and eight comparable cities with University of California campuses. Using retail scanner data, we estimate a triple-difference model measuring the change in soda sales in Berkeley during the campaign and election periods relative to untreated beverage products, untreated cities, and the pre-campaign period. The results of this analysis show that the drop in soda sales starting after the election was not unique to campus retailers.

Our paper fits into a growing literature informing policymakers about the potential impacts of SSB taxes.[8] Silver et al. (2017) use weekly panel-level scanner data from two supermarket chains in Berkeley and adjacent cities and find soda consumption fell by 9.6% in Berkeley stores, but rose 6.9% in non-Berkeley stores after the tax. We expand on this study by providing campaign- and election-specific treatment effects in addition to post-tax treatment effects on sales. Additionally, in our drug store analysis, we do not use adjacent cities as control cities, since they could have been affected by the media campaign around the election, confounding the results. Another related study, Falbe et al. (2016), uses survey-based evidence on SSB consumption—comparing the responses of survey participants in Berkeley to survey participants in neighboring Oakland and San Francisco. Falbe et al. (2016) estimate that the quantity of SSBs purchased in Berkeley dropped by 21%. We

[7]Cawley and Frisvold (2015) and Falbe et al. (2015) examine the incidence effects of the Berkeley soda tax and both studies find that roughly half of the tax was passed on to consumers four to five months after the election. Grogger (2017) estimates the incidence of a sugary drink tax in Mexico and finds more than full price pass-through of the tax for sugary drinks.

[8]Please see Cornelsen and Smith (2018) for a recent review of the literature on ex-post soda tax evaluations, and Paarlberg et al. (2017) for a discussion of the spread of local SSB excise taxes in the U.S.

extend their analysis by using actual purchase data, rather than stated consumption levels, which could be biased. Furthermore, Falbe et al. (2016) conduct the surveys in two separate blocks of time—before the campaign and after the tax implementation—and thus they cannot distinguish between the election's effect and the tax's effect on SSB consumption.

A third related study is Debnam (2017), who uses the Nielsen© Homescan Consumer Panel instead of retail scanner data, to analyze the effect of Measure D on soda purchases. An important contribution of Debnam (2017) is the ability to study consumer heterogeneity, and in particular high- and low-SSB consumption households. A drawback is that the location of households is limited to the county level, so the author uses all households in Alameda County as the treated units. However, Berkeley comprises less than 8% of the population of Alameda County, and there is no way to guarantee the sampled households are located in Berkeley, especially since the Nielsen Homescan sample is designed to be nationally representative and not necessarily representative at the county level. Debnam (2017) finds that high-consuming households living in Alameda County *increase* their weekly SSB consumption by 7.41 ounces relative to other U.S. households. Similar to our results, the change occurs after the election, before the tax implementation. Given we find a significant decrease in soda sales at on- and off-campus retailers in Berkeley, our results together with Debnam (2017) and Silver et al. (2017), suggest more work needs to be done to understand border shopping behavior and spillover effects.[9]

The rest of the paper proceeds as follows. Section 3.2 reviews the literature on potential mechanisms behind policy-induced behavioral change. Section 3.3 describes the setting and summarizes the data, while Section 3.4 outlines the empirical design (i.e. the DID and event study strategies). Section 3.5 presents the results from the analysis of the on-campus data while section 3.6 presents the results using off-campus data. Section 3.7 discusses policy implications and future research.

3.2 Literature Review

Mechanisms Behind Behavioral Change

While our results will suggest the election caused a change in purchasing behavior well before the tax led to a price increase (both on- and off-campus), this paper—similar to other natural experiments—cannot distinguish the exact mechanism behind these changes (e.g., media information effects, rational addiction effects, and social norms effects). First, our results are consistent with models where consumers update their beliefs and behaviors based on information provided by the media and by advisory campaigns. Several studies

[9] There have also been several studies examining SSB taxes outside of Berkeley. In the context of the U.S., Fletcher et al. (2010) use variation in soda taxes across states and estimate that a one percentage point increase in the soda tax implies a reduction of 6 soda-calories per day, accounting for 5% of daily caloric intake from soda. Colchero et al. (2016), Colchero et al. (2017), and Aguilar et al. (2017) examine the effects of a countrywide sugary drink tax in Mexico and find a 6–9% reduction in demand for sugary drinks compared to untaxed products.

show that new information about food-related health problems, food-safety, and animal-safety can alter preferences and consumer demand (Chavas 1983; Brown and Schrader 1990; Van Ravenswaay and Hoehn 1991; Yen and Jensen 1996; Schlenker and Villas-Boas 2009; Lusk 2010). The approach of our analysis is particularly close to Lusk (2010), who uses scanner data to examine how consumer demand for eggs changed in the months leading up a statewide election on whether to bar the use of cages in California egg production. The author finds that demand for the types of eggs associated with higher animal welfare standards increased over time in response to articles on the election, whereas demand for other types of eggs fell.

Second, our results are consistent with models of rational addiction (Becker & Murphy, 1988), which model consumption of addictive products as a function of past and future prices, and where permanent price changes can curb addictive behavior for a product. The rational addiction model has been applied to many common consumer goods, including coffee (Olekalns and Bardsley 1996), alcohol (Waters and Sloan 1995), and cigarettes (Chaloupka 1991). Gruber and Köszegi (2001) study consumers addicted to cigarettes using a rational addiction model with time-inconsistent preferences and find that excise tax legislation that has been enacted but not yet implemented can cause immediate behavioral changes, instead of behavioral changes that only result once the tax is actually put into place. With the case of the SSB legislation in Berkeley, the effect we witness may be the result of forward-looking consumers adjusting their behavior as soon as the election outcome was reached, knowing that the SSB tax would go into effect in the near future.

Third, the results in this paper are consistent with the rich literature on peer effects and social norms, and how these effects may lead to changes in consumption behavior (Rosenquist et al. 2010; Christakis and Fowler 2008; Allcott 2011). The Measure D election revealed that 75% of Berkeley voters were in favor of a SSB tax. As many university consumers are not originally from Berkeley, the election may have revealed new information that was not previously known about the social norms of peers and neighbors in Berkeley. Social norms have been found to have the greatest effect when the comparison group is most similar to the treated individual (Goldstein et al., 2008). Thus, we might expect local elections to have stronger social norm effects than state and national elections. In a university setting, peer communication can have a large influence on the way norms spread among groups, especially with respect to "sin" products (Kremer & Levy, 2008; Real & Rimal, 2007). Thus, it may be the case that university students and staff are more susceptible than other populations to social norms about products deemed to be unhealthy.

3.3 Empirical Setting and Data

Background on Measure D

Since 2009, the soda industry has spent more than $117 million to stop soda tax initiatives in the U.S., such as those considered by the U.S. Congress and in states such as Maine, Texas, and New York.[10] For Berkeley's Measure D in particular, the American Beverage Association of California contributed almost $2.5 million to defeat the tax, while supporters of Measure D spent just under $1 million.[11] One of the strongest supporters of Measure D—"Berkeley vs. Big Soda"—gathered industry, individual, and lawmaker support and funded an aggressive advertising campaign promoting "Yes on D" and emphasizing the need to fight "Big Soda." While the SSB tax in Measure D affects all beverages containing added sugar at a rate of $0.01 per ounce,[12] the City of Berkeley and the media paid particular attention to soda, rather than SSB products in general (see Figure C.1 in the Appendix for examples). For instance, we found that a vast majority of popular articles and op-eds written on Measure D in Berkeley refer to a soda tax rather than a sugar-sweetened beverage tax.[13] Thus, we will look at the effects of the campaign war and election on regular soda separately from other SSB products.

Given that time series data on campaign expenditures are not available, we investigate the intensity of the campaign over time by examining web search data. Figure 3.1 depicts Google Trends data for web searches of the terms "soda tax", "sugar tax", and "beverage tax" in the San Francisco-Oakland-San Jose area in the weeks before and after the election.[14] Numbers on the y-axis represent search interest relative to the highest point on the chart for the given region and time. A value of 100 is the peak popularity for a term. A value of 50 means that a term is half as popular. Likewise a score of 0 means a term was less than 1% as popular as the peak. Figure 3.1 shows that the relative search interest for "soda tax" reached 5% in July after the election was first announced. It grew to 7% and 17% in September and October respectively, suggesting the campaign war led to increased awareness of a potential soda tax. Web searches then spiked to 100% in early November, after Measure D was voted on and passed into law. The interested reader can adjust the Google Trends query dates closer to the election in order to see that the web search spike occurred on November 5, the day after the election. Conversely, when analyzing search trends for the terms "sugar tax" and "beverage tax," we find only modest increases in search interest for these terms around the election, evidence that attention was focused on soda rather than SSBs more broadly.[15]

This increased search interest after the election may have several explanations: (1) voters searching for the outcome of the election, (2) prominent national and local news coverage

[10] *The New York Times*. "Berkeley Officials Outspent but Optimistic in Battle Over Soda Tax." Oct. 7, 2014. Online, *accessed May 21, 2016*.

[11] *Berkeleyside*. "Around $3.4M Spent on Berkeley Soda Tax Campaign," Feb. 5, 2015. Online, *accessed May 21, 2016*.

[12] Official legislation regarding which beverages fall under the tax and which beverages are exempt can be found Online, *accessed February 21, 2018*.

[13] Of the 34 articles on Measure D archived at *Berkeleyside*, an independent and reputable news site that reports on Berkeley and the East Bay, 32 articles refer to the tax as a "soda tax" (or Berkeley taking on "Big Soda") while 2 articles refer to it as a "sugar tax."

[14] Source: Google Trends. Online, *accessed February 13, 2018*.

[15] We also examine Google Trends data for web searches of the terms "sugar-sweetened beverage" and "sugar-sweetened beverage tax" in the same region and time period. We find zero interest in these alternate phrases.

Figure 3.1: Google Trends Web Search Interest of "Soda Tax", "Sugar Tax", and "Beverage Tax" in the San Francisco Bay Area Over Time

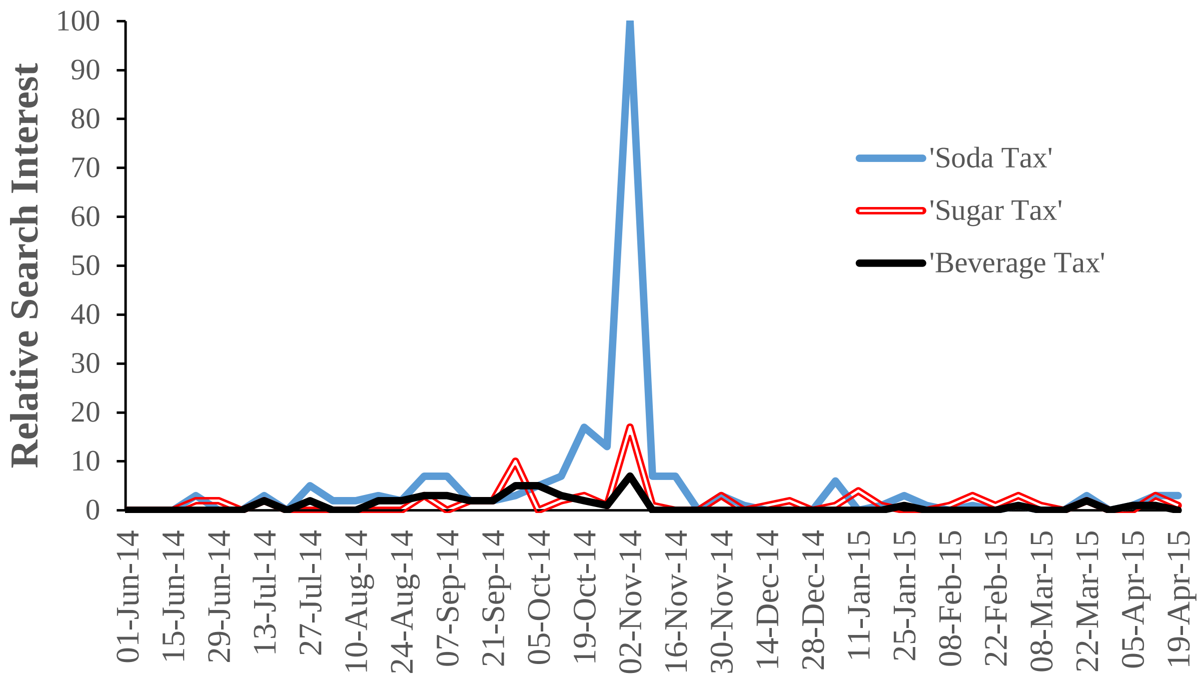

Source: Google Trends. Online, *accessed Aug. 2, 2018.* Note: Numbers on y-axis represent search interest relative to the highest point on the chart for the given region and time. A value of 100 is the peak popularity for the term. A value of 50 means that the term is half as popular. Likewise a score of 0 means the term was less than 1% as popular as the peak.

leading to more searches, as Berkeley historically became the first city in the U.S. to pass a SSB tax, and (3) a delay in exposure to campaign information and searching for more details on the tax. Overall, Figure 3.1 indicates that the campaign did raise some interest in soda taxes in the two months before the election, however, this is dwarfed by interest shown after the election outcome.

University Retail Data

We use a unique data source to estimate the effect of a media campaign and election on consumer purchasing decisions: a retail dataset from dining locations at the University of California, Berkeley. This dataset includes monthly data on the total quantities sold and revenue sales at the product level—i.e., campus sold x ounces of product z in month m, where a product is represented by a unique bar-code. The dataset includes all beverage products for the period January 2013 through December 2016. We categorize products into eight product groups: 1) soda, 2) water, 3) juice, 4) energy drinks, 5) milk, 6) coffee, 7) tea, and 8) diet drinks. We focus on beverage products in order to have a common unit of analysis—fluid ounces.

While the university retailers in our empirical analysis may not be representative of the average U.S. food outlet, there are several advantages of using this empirical setting for our experimental design. First, we have strong institutional knowledge of our setting. Our data come from on-campus retailers, which are open to all people on campus and do not include residential dining halls. Beverages are sold à la carte with individual product prices posted (i.e., drink prices are not hidden in the price of a meal). The products offered, promotional effort, and posted prices are uniform across campus locations, which would otherwise be a concern because we only have aggregate campus retail sales and not sales by individual campus locations. Customers can use cash, credit and debit cards, and university ID cards loaded with "meal points" to make purchases. However, our data does not include information on payment type or customer identifiers, and thus we cannot track customers over time or look at customer heterogeneity. Second, we know exactly when and by how much the soda tax is passed on to consumer, as told to us directly by the campus retail staff. Unlike other studies (i.e., Cawley and Frisvold 2015; Falbe et al. 2015; and Silver et al. 2017), we do not have to infer tax pass-through from observing shelf prices or using scanner data. Third, the university is an important, yet previously unstudied, retailer in the context of Berkeley and Measure D, with the student population more than a third the size of the city population.[16]

We define soda as the treated product category, which we will compare to the seven other beverage groups. It is important to note that some of our beverage categories, aside from soda, may include products that fall under the regulation (in particular, juice, energy drinks, coffee, and tea). Given the wording of Measure D ("The City hereby levies a tax of one cent ($0.01) per fluid ounce on the privilege of distributing sugar-sweetened beverage products in the city"), any drinks with added sweeteners are taxed.[17] So for example, 100% juices are not taxed, but juices with sugar or corn syrup added are taxed. In an attempt to distinguish regulated beverages from unregulated beverages, we group products into the separate diet drinks category if they have "diet", "low-calorie", "zero-calorie", or "unsweetened" in their name. While this works well for categorizing soda and energy drinks into taxed and untaxed, it does a poor job for juices—which tend not to distinguish between natural sugars and added

[16] Source: According the U.S. Census Bureau, the population of Berkeley was 121,240 in 2016. There were 40,173 students enrolled at UC Berkeley in 2016-2017 (Office of Planning and Analysis, UC Berkeley. Online, *accessed Febrary 21, 2018*).

[17] The following beverage products are taxed: regular soda, sport and energy drinks, sweetened tea, and lemonade. Exempted are the following: water (without added sugar), diet drinks (drinks sweetened with zero/low-calorie sweeteners), beverages containing only natural fruit and vegetable juice, beverages in which milk is the primary ingredient, beverages or liquids sold for purposes of weight reduction as a meal replacement, medical beverages (used as oral nutritional therapy or oral rehydration electrolyte solutions for infants and children), and alcoholic beverages, although the last two categories are not sold on campus.

sugars—and for tea and coffee drinks—which do not always have the sugar content in the name, nor do we know if taxed syrups were added.[18] Thus the juice, tea, and coffee categories may contain both regulated and unregulated products.

The substantial amount of advertising and campaigning directed at soda, and not sugar-sweetened-beverages, may have affected consumption of soda differently from other SSBs. For this reason, we focus first on soda sales as the treated product; however, we we also examine the effect of the soda tax election on other beverages, namely, energy drinks (another SSB) and diet drinks (a substitute for SSBs).

Summary Statistics

We use the pre-campaign period data to investigate pre-existing trends in demand for soda versus the control beverage groups. Table 3.1 presents the average monthly ounces sold by beverage group in the academic year before the campaign (August 2013 to July 2014). Figure 3.2(a) unpacks these averages and plots the monthly ounces sold by beverage group over time, both before and after the campaign. The highest selling categories in the pre-campaign period are juice, water, energy drinks, followed by soda and coffee. Milk, tea, and diet drinks experience the lowest levels of sales. While the various products differ in levels, their seasonal patterns are quite similar, with sales peaking in April—the weeks leading up to final exams—and plummeting in June—after the Spring semester ends. While soda has different quantities sold than the other products, to the extent that these differences are constant over time, product-group fixed effects will control for all possible time-invariant determinants of beverage sales, and month fixed effects will control for seasonality in sales. Additionally, Figure 3.2(b) plots the linear trend in average monthly sales by product category, both before and after the start of the campaign. Before the campaign, all categories share similar trends, with the exception of coffee.[19] For this reason, we will examine the sensitivity of our results with and without the inclusion of coffee. After the campaign, juice, water, tea, and milk remain on similar trends as before the campaign while soda and energy drinks experience a decline in trends and diet drinks and coffee experience an increase in trends. This is consistent with consumers substituting coffee and diet drinks for soda and energy drinks after the soda tax campaign and election. While these figures are visually suggestive, we will test this relationship formally in our regression analysis. As a final descriptive statistic, Table 3.1 also presents the average price per ounce by beverage group. Water, soda, tea, and diet drinks are between 8 and 12 cents per ounce while energy drinks, juice, milk, and coffee are between 19 and 27 cents per ounce.

[18] Unfortunately, we do not have information on which juices have natural or added sugar.

[19] As a more rigorous test of parallel trends, we regress quantity sold on a time trend interacted with the eight products. We find that the point estimates of the product trends are not statistically different from each other, again with the exception of coffee. Furthermore, the time series correlation of the sample averages of soda and the other products is high, suggesting that the products share broadly similar time varying patterns in the pre-campaign period.

Table 3.1: Average Monthly Ounces Sold and Price per Ounce by Beverage Group (2013-2014 Academic Year)

	Coffee	Diet	Energy	Juice	Milk	Soda	Tea	Water
	mean/sd	mean/sd	mean/sd	mean/sd	mean/sd	mean/sd	mean/sd	mean/sd
Quantity Sold (Oz)	123846.73	20977.97	167810.88	213262.79	72516.03	114172.37	69005.18	217998.35
	(74436.02)	(11260.22)	(76218.38)	(109572.74)	(31448.30)	(56332.02)	(40961.09)	(108092.20)
Price per Oz ($)	0.27	0.12	0.19	0.23	0.24	0.10	0.11	0.08
	(0.03)	(0.02)	(0.02)	(0.01)	(0.04)	(0.01)	(0.01)	(0.00)
N	12	12	12	12	12	12	12	12

Standard deviations in parentheses. The academic year begins August 2013 and ends July 2014.

Figure 3.2: Monthly Quantities Sold by Product Group (Pre-Campaign Period)

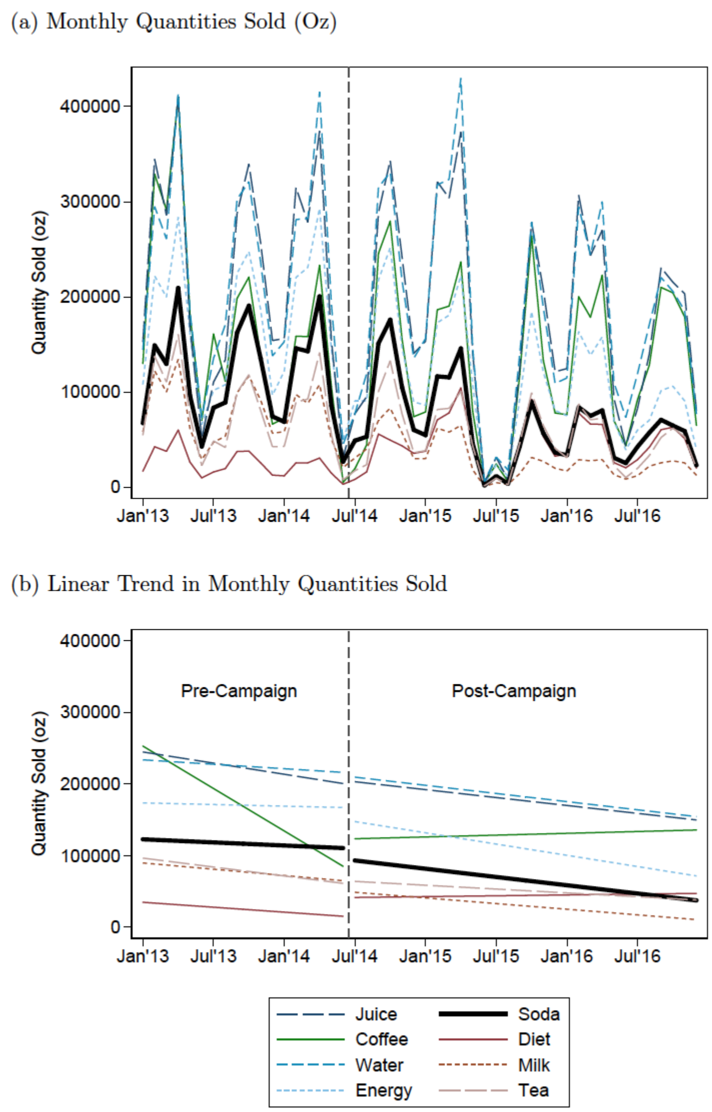

Note: Beverage products are categorized into eight groups: 1) Juice, 2) Coffee, 3) Water, 4) Energy Drinks, 5) Soda, 6) Diet Drinks, 7) Milk, and 8) Tea.

In evaluating the effects of the soda tax campaign, we will compare the pre-campaign period to four separate post-campaign periods: (1) the pre-election campaign period—July 2014-October 2014, (2) the post-election and pre-tax implementation period—November 2014-February 2015, (3) the tax implementation period in City of Berkeley but not on campus—March 2015-July 2016, and (4) the tax implementation period on campus—August 2016-December 2016. It is important to note here that while the City of Berkeley implemented the SSB tax in March 2015, campus retailers did not start receiving the SSB tax on invoices from their vendor until August 2015, and did not pass the tax on to consumers in any form until August 2016.[20] Furthermore, when prices increased on campus, they increased by roughly a penny per ounce for all beverage groups, except water and milk which have no added sugar and are exempted by the SSB tax. Interestingly, diet drinks saw the same price increase as soda (i.e., the price of Diet Pepsi and Pepsi both increased by one cent per ounce). The tax is set up such that it is paid by the distributor, who may or may not pass the cost on to their consumers. Both Falbe et al. (2015) and Cawley and Frisvold (2015)—who examine prices at non-campus retailers in the City of Berkeley—find incomplete pass through of Berkeley's SSB tax on to consumers three months after the policy implementation, with roughly half of the tax passed through. In our setting, campus food and beverage prices are sticky and only change once per year, occurring during the summer months of June, July, or August. Since the tax was not passed through to consumers on campus for almost two years after the campaign, this paper examines how the soda tax campaign, election, and increase in prices off-campus affect the sales of soda on-campus.

3.4 Empirical Strategy

Our approach has two parts. First, we use a difference-in-differences (DID) strategy to measure the change in soda sales due to the soda tax campaign and election. Secondly, we estimate an event study model to test the identifying assumption of the DID model, namely that soda sales would have continued on the same trend as the other products had it not been for the campaign and election.

Difference-in-Differences Model

The DID model compares purchases of soda (i.e., the treated category) with purchases of the seven other beverage groups (i.e., the control categories), in the four policy periods. Using data from January 2013 through December 2016, we compare the pre-campaign period to four subsequent periods: (1) pre-election campaign, (2) post-election and pre-policy implementation, (3) post-policy implementation in the City of Berkeley, and (4) post-policy implementation on campus. For shorthand, we refer to these periods as: pre-campaign, campaign, post-election, post-city, and post-campus. By comparing the soda purchase behavior in the pre-period to each of these policy periods, we attempt to distinguish the effects of the campaign from the effects of the election and the effects of prices increasing off- and on-campus. The DID model specification is as follows:

$$\begin{aligned} Q_{im} = {} & \beta_1(\text{Soda*Campaign})_{im} + \beta_2(\text{Soda*PostElection})_{im} + \beta_3(\text{Soda*PostCity})_{im} \\ & + \beta_4(\text{Soda*PostCampus})_{im} + \alpha_i + \alpha_m + \epsilon_{im}. \end{aligned} \tag{3.1}$$

[20]This was reported to us by campus retail staff and confirmed in the data.

where Q_{im} is the quantity sold (measured in ounces) of beverage group i in month-of-sample m. We estimate equation (3.1) with quantities both in levels and in logs (i.e., Q_{im} and $ln(Q_{im})$). $Soda_i$ is an indicator for beverage group i being in the treated soda group. Four time indicators—$Campaign_m$, $PostElection_m$, $PostCity_m$, and $PostCampus_m$—define the four policy periods. Finally, we include fixed effects for the eight product groups α_i and for the month-of-sample α_m. It should be noted that price is not included in this estimating equation due to endogeneity concerns. Moreover, since prices are adjusted only once every year in either June, July, or August, the month-of-sample fixed effects pick up much of the price variation that may have biased results.

The coefficients of interest are those on the interactions of $Soda_i$ and the policy periods. The coefficient for $Soda * Campaign_{im}$ is the effect of the campaign on soda sales relative to the control product categories, the coefficient on $Soda * PostElection_{im}$ is the effect of the election, the coefficient on $Soda * PostCity_{im}$ is the effect of the SSB tax change in the city of Berkeley, and the coefficient on $Soda * PostCampus_{im}$ is the effect of the SSB tax change on campus.

Event Study Model

The identifying assumption of the DID model is that of parallel trends, where soda sales would have continued on the same trend as the other product groups had it not been for the campaign, election, and tax implementation. To directly test this assumption, we complement the DID model with the following event study model:

$$Q_{imt} = \sum_{t=-5}^{7} \beta_t D_{t,im} + \alpha_i + \alpha_m + \epsilon_{imt} \tag{3.2}$$

where $D_{t,im}$ is a dummy variable equaling one if product group i is in the soda product group and month-of-sample m is t time periods from the election. Once again, we include fixed effects for the eight product groups α_i and for the month-of-sample α_m. The time periods t are four month intervals centered at the election (i.e., t=0 is Nov 2014 to Feb 2015). We omit the period immediately preceding the election (t=-1) to avoid perfect collinearity. Thus, equation (3.2) is the same as equation (3.1), except instead of splitting the sample into 5 periods of unequal length, we instead compare soda sales to the untreated products in every 4-month interval of the sample. The β_t vector contains the coefficients of interest, which we plot over time to trace out the adjustment path from before the soda tax campaign through the election and policy implementations. Importantly, if soda is trending parallel to the control products before the policy periods, there should be no trend in the β_t coefficients in the pre-campaign period and they should be statistically indistinguishable from zero.

Estimation Concerns

An important potential limitation of our analysis is that any beverage could be a substitute for soda. For example, diet drinks sales may increase due to the Measure D election as regular soda sales decrease. Since we are examining soda sales relative to the other beverage groups, an increase in sales in one of the other beverage categories would bias our estimates in the same direction as a drop in soda sales. In other words, our effects would be biased in the correct direction but they would be biased larger in magnitude. To address this concern, we estimate equation (3.1) seven times, excluding one of the other beverage groups each time, in order to evaluate whether substitution towards one of the other products is biasing our results. In this way, we are able to gain some clarity on whether consumers are substituting towards certain beverage products more than others as a result of the election. However, while having a potential substitute as a control may create an upward bias in our treatment effect, it should not bias the timing of when the effect occurs, which is one of our main research objectives.

A second limitation of the university data is that they do not contain a comparison location unaffected by the soda tax campaign and election. To address this limitation, we supplement our campus analysis with an analysis of beverage sales at drug stores in Berkeley and eight comparable cities with University of California campuses. With these data, we measure the change in soda sales in Berkeley during the campaign and election periods relative to untreated beverage products, untreated cities, and the pre-campaign period. This analysis provides evidence that the drop in soda sales starting after the election was not unique to campus retailers.

3.5 Results

Effect of the Soda Tax Campaign on Soda Purchases (Campus Retail Analysis)

We present the results from the reduced form specification of equation (3.1) in Table 3.2, where the dependent variable is the quantity sold (in ounces) of product group i and month-of-sample m, in levels (column 1) and in logs (column 2). The parameters of interest are the four interactions of the soda indicator and the policy period indicators. Standard errors are clustered at the product group by academic year level, to account for the possibility that the errors are correlated within a given product group and academic year, but not across product groups or years.[21]

There are three main takeaways from Table 3.2. First, in both columns the coefficients on the *Campaign* interaction are positive, small in magnitude, and are not statistically different from zero. This suggests the campaign did not alter soda sales, on average, relative to the control beverage groups. We acknowledge that consumer heterogeneity may lead to a statistically insignificant average effect—where decreased consumption by consumers relating more to the "Yes" side of Measure D could cancel increased consumption of consumers relating more to the "No" side. However, without customer level data, we can say little about

[21] We use the academic year instead of the calendar year since campus retail product and pricing decisions are made at the beginning of the academic year.

Table 3.2: Difference-in-Difference: Effect of Soda Tax Campaign and Election on Campus Retail Soda Sales Relative to Other Beverage Products

	(1) Oz Sold	(2) Log Oz Sold
Soda × Campaign	3339.766	0.035
	(12267.539)	(0.112)
Soda × Post-Election	-11172.373	-0.227*
	(10735.197)	(0.126)
Soda × Post-Policy City	-19958.315*	-0.441***
	(10663.730)	(0.145)
Soda × Post-Policy Campus	-23253.053*	-0.443**
	(12644.855)	(0.168)
Mean of Dep. Variable	112376.308	11.179
Num of Obs.	384	384
R squared	0.839	0.928
Product Group FE	X	X
Month-of-Sample FE	X	X

Standard errors in parentheses are clustered at the product group by academic year level. The outcome variable is ounces sold of product group i in month m, in logs (column 1) and in levels (column 2). Asterisks indicate the following: $*p < 0.10$, $**p < 0.05$, $***p < 0.01$

heterogeneous effects. Second, the coefficients on the other three interactions are negative, much larger in magnitude, and statistically different from zero at the 10% significance level, with the exception of the *Post-Election* interaction in column (1). Moreover, the coefficients on the *Post-City* and *Post-Campus* interactions are nearly double the coefficients for the *Post-Election* interaction. Translating the coefficients into percent changes shows that soda sales were 10-20% lower post-election compared to pre-campaign and 18-36% lower post-tax implementation compared to pre-campaign.[22] These results suggest that soda sales began to deviate below the sales of the control beverage groups after the election and this decrease continued through the tax implementation periods. Third, the coefficients on the *Post-City* and *Post-Campus* interactions are nearly equal to one another, suggesting that the price changes that occurred on-campus almost two years after the election did not lead to any

[22] In column 1, the mean of the dependent variable is 112,376 ounces per month, thus a coefficient of -11,172 on $Soda \times PostElection$ translates to a 10% decrease in quantity sold. In column 2, the percent change in the dependent variable can be found using $100 \times (\exp^{\beta_1} - 1)$, thus the coefficient of -0.433 on $Soda \times PostCampus$ translates to 36% decrease in quantity sold.

additional decreases in sales.

To understand whether the use of the other beverage groups as controls for soda is biasing our results away from zero, we estimate equation (3.1) seven times, excluding one of the control beverage groups each time. Comparing column (2) in Table 3.2 to each of the columns of

Table 3.3: Robustness by Control Beverages: Effect of Soda Tax Campaign and Election on Campus Retail Soda Sales Relative to Other Beverage Products

	(1) Excl. Coffee	(2) Excl. Diet	(3) Excl. Energy	(4) Excl. Juice	(5) Excl. Milk	(6) Excl. Tea	(7) Excl. Water
Soda × Campaign	-0.003	0.088	0.054	0.040	0.013	0.013	0.043
	(0.100)	(0.091)	(0.135)	(0.127)	(0.122)	(0.116)	(0.125)
Soda × Post-Election	-0.242*	-0.094	-0.264*	-0.208	-0.305**	-0.250*	-0.223
	(0.124)	(0.100)	(0.145)	(0.147)	(0.128)	(0.141)	(0.148)
Soda × Post-Policy City	-0.424***	-0.266**	-0.463***	-0.439**	-0.597***	-0.470***	-0.425**
	(0.151)	(0.116)	(0.169)	(0.170)	(0.135)	(0.163)	(0.169)
Soda × Post-Policy Campus	-0.382**	-0.277*	-0.515***	-0.423**	-0.594***	-0.477**	-0.433**
	(0.171)	(0.148)	(0.185)	(0.197)	(0.161)	(0.191)	(0.197)
Mean of Dep. Variable	11.126	11.312	11.123	11.076	11.289	11.247	11.066
Num of Obs.	336	336	336	336	336	336	336
R squared	0.933	0.948	0.927	0.918	0.944	0.922	0.918
Product Group FE	X	X	X	X	X	X	X
Month-of-Sample FE	X	X	X	X	X	X	X

Standard errors in parentheses are clustered at the product group by academic year level. This table replicates column (2) of Table 3.2, with one of the seven control beverage groups excluded in each column. Asterisks indicate the following: $*p < 0.10$, $**p < 0.05$, $***p < 0.01$

Table 3.3, the only beverage group when excluded that alters the results is diet drinks (shown in column 2). In particular, while the coefficients estimated excluding diet drinks follow the same sign and pattern as the coefficients including diet drinks, the coefficients estimated without diet drinks are nearly half the size. This is an interesting result in and of itself, suggesting some consumers substituted diet drinks for regular soda after the election. Since these results raise support for the concern that diet drinks may not be a valid control for soda, the specification in column (2) of Table 3.3 is our preferred specification.[23] Translating the coefficients in column (2) into percent changes, soda sales relative to the remaining six beverage groups were 9% lower post-election compared to pre-campaign and 23-24% lower post-tax implementation compared to pre-campaign.

Thus far we have focused on soda, since soda was the target of the election campaign. However, Figure 3.2 suggests other beverage categories were affected by the soda tax election—in particular, energy drinks and diet drinks. In Table 3.4, we estimate the effects of the soda tax campaign on soda, energy drinks, and diet drinks relative to one control category, water. We use water as a comparison because there should be little confusion about whether water is taxed (unlike juice, tea, coffee, and milk) and because trends in water sales are not statistically different from soda, energy drinks, and diet drinks in the pre-campaign period (as shown in Figure 3.2). Columns (1) and (2) of Table 3.4 show that soda and energy drinks experience similar declines in sales (relative to water) after the election but not during the campaign period. Therefore, in column (3) we jointly compare soda and energy drinks to water and find that sales of these SSB drinks fell by 24.9% after the election compared to the pre-campaign period. Sales dropped another 10 percentage points after the tax was implemented in the city and another 10 percentage points after the tax was implemented on campus. Conversely, column (4) shows that diet drink sales increased relative to water in all four treatment periods. For instance, diet drink sales are 77.5% greater in the post-election period than the pre-campaign period. This, once again, suggests some customers substituted diet drinks for SSBs.

In summary, even though the tax was not implemented on campus during the *Post-Election* and *Post-Policy City* periods, we find consumers purchased less soda relative to the other beverage groups. The result in the *Post-Policy City* period is particularly surprising, given soda on-campus would have been relatively cheaper when prices increased off-campus. Our results are consistent with the election causing consumers to update their beliefs and change behavior. In particular, the election revealed a social norm that 75% of people in Berkeley were in favor of the SSB tax. If, instead, sales fell due to rational addiction, two years is a long time for sales to remain depressed without a change in price. Changes in purchasing behavior did not occur during the $3.4 million campaign period; however, we cannot rule out that delays in receiving campaign information led to changes in consumption after the election or that a lack of average effects during the campaign period reflects the consumption changes from the opposing sides of the campaign canceling each other.

[23] We replicate Table 3.2 without diet drinks in Appendix Table C.1.

Table 3.4: Effect of Soda Tax Campaign and Election on SSB and Diet Drink Sales Relative to Water

	(1)	(2)	(3)	(4)
	$Soda\|Water$	$Energy\|Water$	$S\&E\|Water$	$Diet\|Water$
Soda × Campaign	-0.011			
	(0.080)			
Soda × Post-Election	-0.249***			
	(0.003)			
Soda × Post-Policy City	-0.534***			
	(0.036)			
Soda × Post-Policy Campus	-0.501***			
	(0.003)			
Energy × Campaign		0.066		
		(0.139)		
Energy × Post-Election		-0.249***		
		(0.003)		
Energy × Post-Policy City		-0.231***		
		(0.014)		
Energy × Post-Policy Campus		-0.487***		
		(0.003)		
S&E × Campaign			0.028	
			(0.108)	
S&E × Post-Election			-0.249***	
			(0.066)	
S&E × Post-Policy City			-0.382***	
			(0.083)	
S&E × Post-Policy Campus			-0.494***	
			(0.061)	
Diet=1 × Campaign				0.268*
				(0.128)
Diet=1 × Post-Election				0.775***
				(0.141)
Diet=1 × Post-Policy City				0.953***
				(0.141)
Diet=1 × Post-Policy Campus				0.939***
				(0.141)
Num of Obs.	96	96	144	96
Product FE	X	X	X	X
Month-of-Sample FE	X	X	X	X

Standard errors in parentheses are clustered at the product group by academic year level. The outcome variable is ounces sold of product group i in month m. In columns 1-4, soda, energy, and diet drinks are compared to the control category (water). In this table, S&E is an abbreviation for Soda and Energy Drinks. Asterisks indicate the following: $*p < 0.10$, $**p < 0.05$, $***p < 0.01$

Event Study Results

Given the interesting patterns we find in the DID results, we next explore the parallel trends assumption and the dynamics of the treatment effects over time using our event study model. Figure 3.3 plots the estimates we obtain from equation (3.2), excluding diet drinks, with the β_t plotted in black and the 95 percent confidence intervals plotted in gray. Standard errors are once again clustered at the product group by academic year level. Vertical red lines separate the sample into the four treatment periods. The omitted dummy is D_{-1}, which corresponds to the four month interval of the campaign period.

In the periods before the election, we find parallel trends, with each of the β_t not statistically different from zero at the 95% significance level. After the election in November 2014, the β_t estimates begin to decline, indicating that soda sales dropped relative to the control beverage groups. By a year after the election, the β_t are no longer declining, but are at a constant level significantly lower than the pre-campaign period. These event study results suggest that the decline in soda sales on-campus relative to the other beverage categories began after the election, and that sales of soda remained depressed after the tax was implemented off- and on-campus. The event study results also provide evidence against the alternative hypothesis that slowly declining preferences for soda, rather than the campaign and election, are driving the decreases in consumption described. In particular, there are no downward trends in the pre-campaign period to suggest the results are driven by declining preferences for soda. Instead, the downward trend in sales begins only after the election and stabilizes by a year later.

Figure 3.3: Event Study: Effect of Soda Tax Campaign and Election on Campus Retail Soda Sales Relative to Other Beverage Products

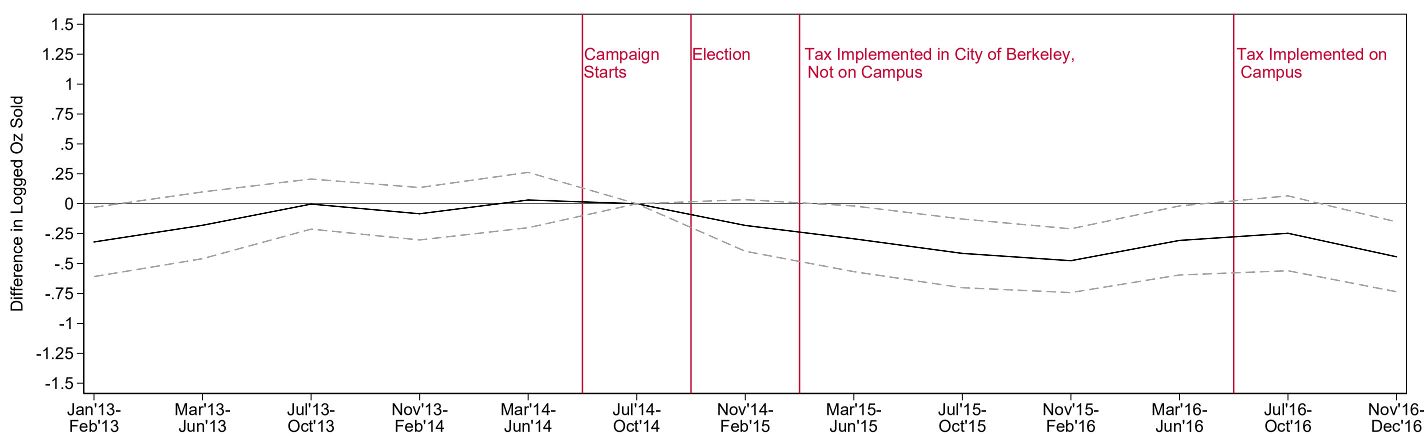

Note: The figure displays the β_t coefficient estimates from event study equation 3.2. The dependent variable is the logged quantity sold (in ounces) of product group i and month-of-sample m. Upper and lower 95% confidence intervals are depicted in gray, estimated using standard errors clustered at the product group by academic year level.

3.6 Supplemental Analysis Using Nielsen Scanner Data

While the university retail data has the benefit of institutional knowledge (i.e., we know exactly when the tax was implemented and the exact pass-through amount), a major drawback of these data is that they do not contain a comparison location unaffected by the soda tax campaign and election. To address this limitation, we supplement our campus analysis with an analysis of beverage sales at drug stores in Berkeley and eight comparable cities with University of California campuses. With these data, we extend the DID model above to a triple-difference model—measuring the change in soda sales in Berkeley during the campaign and election periods relative to untreated beverage products, untreated cities, and the pre-campaign period. This analysis provides evidence that the drop in soda sales starting after the election was not unique to campus retailers.

Drug Store Scanner Data

The drug store scanner data are collected by Nielsen© and made available through the Kilts Center at The University of Chicago Booth School of Business.[24] These data include weekly price and quantity information at the product-by-store level—i.e., store j sold x units of product z in week w, where a product is represented by a universal product code (UPC)—from January 2012 through December 2015.[25]

While the Nielsen database includes several types of retail food outlets selling soda and other beverages (e.g. supermarkets, grocery stores, and mass merchandising stores, among others),[26] we focus on drug stores because these are the only stores in the sample we could uniquely identify as being located in Berkeley. This is because the scanner data do not contain the exact street address of each store in the sample; instead, the county and three-digit zip code of each store is provided. There are two cities in Alameda County and zip code 947 (Albany and Berkeley), and we select the five drug stores we can verify are in Berkeley and not Albany using the retailer codes provided.

Given our goal is to compare the drug store analysis to the campus analysis, we select control drug stores as those in counties and 3-digit zip codes containing one of the nine University of California campuses, other than UC Berkeley. Specifically, we use drug stores in

[24]Researchers own analyses calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researchers and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein."

[25]At the time of this study, the Nielsen data span 2006 to 2015. Since our event of interest took place towards the end of the available data, we chose to subset the data from 2012 onwards for computational ease.

[26]The Nielsen data covers more than 50,000 individual stores in 90 participating retail chains across the entire United States.

the same counties and 3-digit zip codes as UC Davis, Irvine, Los Angeles, Merced, Riverside, San Diego, Santa Barbara, and Santa Cruz.[27] We do not use drug stores near UC San Francisco for three reasons: (1) UCSF does not have an undergraduate program, (2) UCSF persuaded every vendor on campus to stop selling SSBs in 2015, and (3) San Francisco had a SSB tax election at the same time as Berkeley which did not pass. The second two reasons would particularly confound changes in soda sales. We decided to restrict this analysis to locations within California, since they will be identical in terms of state policies that may affect consumption of beverages differently across states.

Finally, we select data for product groups similar to the ones used in the campus analysis—soda, water, coffee, tea, milk, and juice—and aggregate the week-store-upc data to the month-store-product-group level in order to match the level of analysis used with the campus data.[28] Thus in total we have 48 months, 80 stores (5 treated and 75 control), and six product groups (1 treated and 5 control).

Table 3.5 shows the average monthly ounces sold by beverage group per store, averaged across all stores in Berkeley and in the control cities, during 2012–2013.[29] The higher selling categories in this pre-campaign period are milk, water, juice, and soda, while the lower selling categories are tea and coffee. The stores in Berkeley sell more ounces per month across all beverage category than the control stores.

[27]We use the following counties and 3-digit zip codes to select control stores: UCD = Yolo County and 956; UCI = Orange County and 926; UCLA = Los Angeles County and 900; UCM = Merced County and 953; UCR = Riverside County and 925; UCSD = San Diego County and 920; UCSB = Santa Barbara County and 931; UCSC = Santa Cruz County and 950.

[28]While we also have data for diet drinks, we choose not to use them in this analysis given the results in the previous section.

[29]The observations in Berkeley and the control cities reflect the number of stores multiplied by 24 months, since the summary statistics are calculated for 2012–2013 for each product group.

Table 3.5: Average Monthly Ounces Sold per Store by Beverage Group (2012-2013 Nielsen Data)

	Coffee	Juice	Milk	Soda	Tea	Water
Quantity Sold (Oz)	mean/sd	mean/sd	mean/sd	mean/sd	mean/sd	mean/sd
Berkeley [N=120]	6380.86	96030.75	158783.86	76072.18	34546.53	97766.22
	(2718.95)	(41491.24)	(66403.39)	(29953.56)	(20529.05)	(50248.46)
Control Cities [N=1,800]	4362.25	43825.13	59158.68	44824.30	26695.99	70161.95
	(2546.88)	(29852.22)	(60161.21)	(27798.79)	(14836.58)	(66173.17)

Standard deviations in parentheses.

Drug Store Empirical Specification and Results

For the drug store analysis, we extend the DID models in equation (3.1) and (3.2) to include an additional dimension—store s, which is city-specific. We estimate the following triple-difference model:

$$Q_{ims} = \beta_1(Soda \times Berkeley \times Campaign)_{ims} + \beta_2(Soda \times Berkeley \times PostElection)_{ims} + \beta_3(Soda \times Berkeley \times PostCity)_{ims} + \alpha_{im} + \alpha_{ms} + \alpha_{is} + \epsilon_{ims}. \tag{3.3}$$

where Q_{ims} is now the quantity sold (measured in ounces) of beverage group i in month-of-sample m in store s, and the fixed effects are product group by month-of-sample (α_{im}), month-of-sample by store (α_{ms}), and product group by store (α_{is}). The coefficients of interest are the interactions of $Soda_i$, $Berkeley_s$ and the three policy periods: $Campaign_m$, $PostElection_m$, and $PostCity_m$. There are only three policy periods in the drug store analysis because our sample ends in 2015, before the tax was implemented on campus. Similarly, we extend the event study equation (3.2) as follows:

$$Q_{imst} = \sum_{t=-8}^{4} \beta_t D_{t,ims} + \alpha_{im} + \alpha_{ms} + \alpha_{is} + \epsilon_{imst} \tag{3.4}$$

where $D_{t,ims}$ is now a dummy variable equaling one if product group i is in the soda product group, store s is in Berkeley, and month-of-sample m is t time periods from the election.

Table 3.6 and Figure 3.4 present the results of equations (3.3) and (3.4). There are three main takeaways from the drug store analysis. First, the results follow the same pattern as the campus analysis in that there is no statistically significant change in soda sales during the campaign period, there is a significant drop in soda sales in the post-election period, and this drop in sales continues into the post-city period. Second, the magnitude of the effects using the drug store data are similar in size to the Berkeley campus data analysis excluding diet drinks. For instance, the coefficient on the post-election interaction is -0.108 in Table 3.6 and -0.094 in Table 3.3 column (2). This suggests the lack of control cities in the campus analysis is not biasing the results away from zero when diet drinks are excluded. Third, converse to the campus analysis, the drug store results show the decrease in soda sales might have begun in the campaign period; however this decrease is not statistically significant in either Table 3.6 or Figure 3.4.

Table 3.6: Triple-Difference: Effect of Berkeley Soda Tax Election on Drug Store Soda Sales Relative to Other Beverage Products and Other Cities

	(1) Oz Sold	(2) Log Oz Sold
Soda × Campaign × Berkeley	-2313.329	-0.033
	(1884.994)	(0.036)
Soda × Post-Election × Berkeley	-782.082	-0.108***
	(2086.024)	(0.028)
Soda × Post-Policy City × Berkeley	-3610.394*	-0.125***
	(1981.229)	(0.023)
Mean of Dep. Variable	42306.065	6.687
Num of Obs.	23040	23040
R squared	0.976	0.983
Product Group × Store FE	X	X
Store × Month-of-Sample FE	X	X
Month-of-Sample × Product Group FE	X	X

Standard errors in parentheses are clustered at the product group by year by 3-digit zip code level. The outcome variable is ounces sold of product group i in month-of-sample m, store s, and city c, in levels (column 1) and in logs (column 2). Asterisks indicate the following: $*p < 0.10$, $**p < 0.05$, $***p < 0.01$

Figure 3.4: Triple-Difference Event Study: Effect of Berkeley Soda Tax Campaign and Election on Drug Store Soda Sales Relative to Other Beverage Products and Other Cities

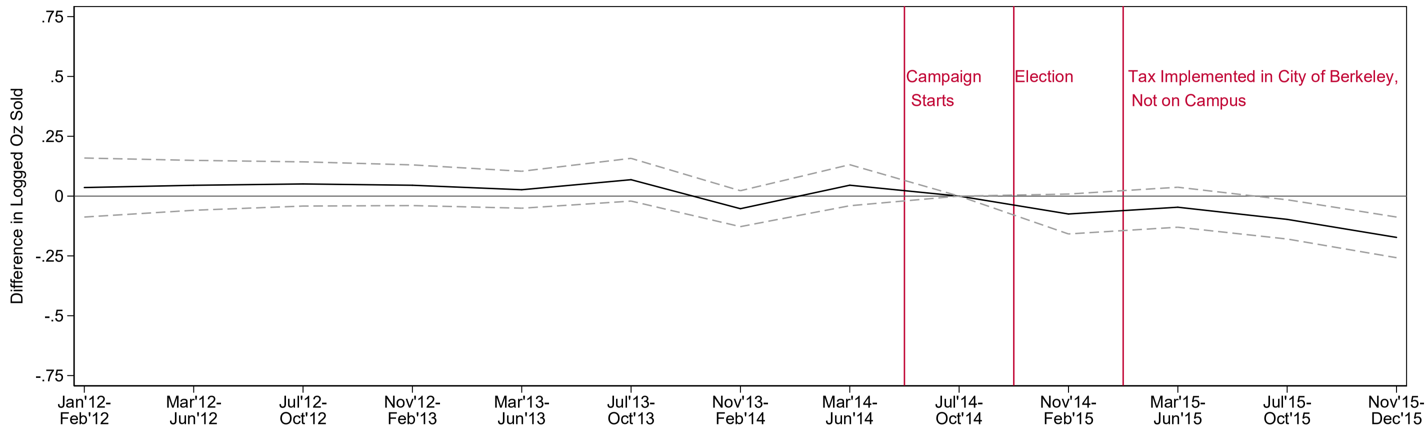

Note: The figure displays the β_t coefficient estimates from event study equation 3.4. The dependent variable is the logged quantity sold (in ounces) of product group i, in month-of-sample m, store s, and city c. Upper and lower 95% confidence intervals are depicted in gray, estimated using standard errors clustered at the product group by year by 3-digit zip code level.

3.7 Discussion

This paper is motivated by the growing adoption of local excise taxes on sugar-sweetened beverages in the U.S. (e.g. Berkeley, San Francisco, Oakland, Boulder, Seattle, and Philadelphia). In particular, Berkeley made history by being the first city to vote and pass a soda tax in a local election. This paper uses a detailed scanner dataset in a university setting to measure the response of the soda tax campaign and election on soda sales. Our results show that soda purchases significantly drop relative to other beverage products immediately after the election, months before the tax is implemented in the city of Berkeley or on campus. Additionally, using scanner data from off-campus retailers in the same city as the campus, we find similar drops in soda sales after the election. Thus, our results are not unique to the university setting. Specifically, we find a 10.8% drop in off-campus sales in the period immediately following the election.

While other studies have examined the effects of the Berkeley SSB tax on beverage sales (Debnam, 2017; Falbe et al., 2016; Silver et al., 2017), this study is unique across multiple dimensions. First, we focus on an understudied yet important setting—university food retailers. This setting is especially important in the context of Berkeley because the student population is more than 1/3 the size the population of the city.poplab A second contribution of this paper is that instead of a simple pre- and post-policy comparison, we explore changes in purchasing behavior during several periods before and after the election (e.g., campaign, post-election, post-implementation in the City of Berkeley, and post-implementation on campus), in an effort to disentangle the mechanisms behind the behavioral changes. Our results show that soda sales fell on-campus after the SSB tax election yet before prices changed due to the tax. This suggests that comparing pre-campaign to post-implementation sales may confound a price elasticity effect with media and social norm effects. Finally, we provide evidence that consumers substituted towards diet drinks. This is particularly interesting given prices on-campus changed for diet soda in the same fashion as regular soda.

An important policy implication of our study is that the effects of election campaigns and outcomes on behaviors may be larger than the effects of the policy itself. In the university setting, we find that sales dropped 10-20% in the post-election period compared to the pre-campaign period. Sales dropped another 8-16% when the policy was implemented off-campus, but did not fall any further when the policy was implemented on-campus. Thus, the change in soda sales post-election but pre-implementation was similar (if not slightly larger) to the change in soda sales post-implementation. These results are consistent with findings from other highly publicized elections. For instance, in the context of standards in egg production, Lusk (2010) finds that the publicity surrounding a vote to pass a proposition pertaining to animal welfare *in itself* had a significant impact on consumer behavior, beyond the effect the policy had once implemented.

The Berkeley tax differs from other soda taxes in several ways, which could have important implications for comparing the results in Berkeley to SSB taxes in other jurisdictions. In particular, it was voted on and passed by the people of Berkeley, and there was an extensive campaign to inform voters about the tax. Conversely, when the Mexican government announced their SSB tax in September 2013, it took the soda industry and the public by surprise, according to media reports.[30] If the Berkeley SSB tax is replicated elsewhere without a proceeding campaign war and affirmative election outcome, its effects on sales may differ substantially. The Berkeley soda tax was also a local election, as opposed to statewide or countrywide election. If the social norm revealed by the election was the driving mechanism behind the consumption change, and not the information revealed by the media campaign, we might expect larger effects when consumers identify more closely with other voters. If instead, rational addiction was the main mechanism behind the reduction in soda consumption after the election, then the size of the jurisdiction, the level of media coverage, and whether the policy is implemented by direct democracy or republic may not matter for external validity.

This study scratches the surface with respect to the mechanisms behind the reduced soda demand. We can reject that SSB taxes only affect beverage demand through current price changes, but there is still much to uncover. We echo the sentiments of Cornelsen and Smith (2018) in that more needs to be done to understand the mechanisms behind the behavioral changes—especially the media, rational addiction, and social norm effects—and how these might vary across heterogeneous consumers. It may be that a combination of these mechanisms is at play for different types of consumers. We suggest three areas of future research. First, to address the limitations of natural experiments such as the one in this paper, laboratory experiments simulating elections could be designed to parse out information effects from social norm effects. Second, rational addiction models could be tested and compared across price changes specifically resulting from elections and price changes resulting from other sources—for instance, an individual may be more likely to believe a tax will be permanent if the vast majority of people voted for it as opposed to the officials currently in power. Lastly, this paper highlights a need for additional research on border shopping behavior and substitution effects from SSB taxes.

[30] *The Guardian.* "How One of the Most Obese Countries on Earth Took on the Soda Giants," November 3, 2015. Online, *accessed February 21, 2018.*

Bibliography

Aaker, J. L. (1997). Dimensions of brand personality. *Journal of marketing research*, *34*(3), 347–356.

Adgate, B. (2018). Why the 2017-18 season was great for the NBA. [link here, (last accessed 2018-11-21)]. *Forbes*. Retrieved November 21, 2018, from https://www.forbes.com/sites/bradadgate/2018/04/25/the-2017-18-season-was-great-for-the-nba

Adler, M. (1985). Stardom and talent. *The American economic review*, *75*(1), 208–212.

Aguiar, M., Hurst, E., & Karabarbounis, L. (2013). Time use during the great recession. *American Economic Review*, *103*(5), 1664–96.

Aguilar, A., Gutierrez, E., & Seira, E. (2017). The Effectiveness of Sin Food Taxes: Evidence from Mexico [Working Paper].

Alavy, K., Gaskell, A., Leach, S., & Szymanski, S. (2010). On the edge of your seat: Demand for football on television and the uncertainty of outcome hypothesis. *International Journal of Sport Finance*, *5*(2), 75.

Allcott, H. (2011). Social norms and energy conservation. *Journal of public Economics*, *95*(9-10), 1082–1095.

Anstine, D. B. (2001). How much will consumers pay? a hedonic analysis of the cable television industry. *Review of Industrial Organization*, *19*(2), 129–147.

Athey, S., & Levin, J. (2001). Information and competition in us forest service timber auctions. *Journal of Political economy*, *109*(2), 375–417.

Badenhausen, K., & Ozanian, M. (2019). NBA valuations [link here, (last accessed 2020-01-30)]. *Forbes*. Retrieved January 30, 2020, from https://www.forbes.com/sites/kurtbadenhausen/2019/02/06/nba-team-values-2019-knicks-on-top-at-4-billion/#4809718be667

Bagwell, K. (2005). *The economic analysis of advertising. columbia university department of economics* (tech. rep.). Discussion Paper Series.

Becker, G. S., & Murphy, K. M. (1988). A theory of rational addiction. *Journal of political Economy*, *96*(4), 675–700.

Becker, G. S., & Murphy, K. M. (1993). A simple theory of advertising as a good or bad. *The Quarterly Journal of Economics*, *108*(4), 941–964.

Berri, D. J., & Schmidt, M. B. (2006). On the road with the National Basketball Association's superstar externality. *Journal of Sports Economics*, *7*(4), 347–358.

Bertrand, M., Karlan, D., Mullainathan, S., Shafir, E., & Zinman, J. (2010). What's advertising content worth? evidence from a consumer credit marketing field experiment. *The Quarterly Journal of Economics*, *125*(1), 263–306.

Bizzozero, P., Flepp, R., & Franck, E. (2016). The importance of suspense and surprise in entertainment demand: Evidence from wimbledon. *Journal of Economic Behavior & Organization*, *130*, 47–63.

Blake, T., Moshary, S., Sweeney, K., & Tadelis, S. (2018). Price salience and product choice. *NBER Working Paper*, (w25186).

Borwein, D., Borwein, J. M., & Marechal, P. (2000). Surprise maximization. *The American Mathematical Monthly*, *107*(6), 517–527.

Brown, D. J., & Schrader, L. F. (1990). Cholesterol Information and Shell Egg Consumption. *American Journal of Agricultural Economics*, *72*, 548–55.

Brown, J. (2011). Quitters never win: The (adverse) incentive effects of competing with superstars. *Journal of Political Economy*, *119*(5), 982–1013.

Brownell, K. D., & Frieden, R., Thomas. (2009). Ounces of Prevention—The Public Policy Case for Taxes on Sugared Beverages. *The New England Journal of Medicine*, *206*(18), 1805–1808.

Bryant, J., Rockwell, S. C., & Owens, J. W. (1994). Buzzer beaters" and "barn burners": The effects on enjoyment of watching the game go "down to the wire. *Journal of Sport and Social Issues*, *18*(4), 326–339.

Bryson, A., Frick, B., & Simmons, R. (2013). The returns to scarce talent: Footedness and player remuneration in european soccer. *Journal of Sports Economics*, *14*(6), 606–628.

Bryson, A., Rossi, G., & Simmons, R. (2014). The migrant wage premium in professional football: A superstar effect? *Kyklos*, *67*(1), 12–28.

Buraimo, B., Forrest, D., McHale, I. G., & Tena, J. (2020). Unscripted drama: Soccer audience response to suspense, surprise, and shock. *Economic Inquiry*, *58*(2), 881–896.

Buraimo, B., & Simmons, R. (2015). Uncertainty of outcome or star quality? television audience demand for english premier league football. *International Journal of the Economics of Business*, *22*(3), 449–469.

Busse, M. R., Knittel, C. R., & Zettelmeyer, F. (2013). Are consumers myopic? Evidence from new and used car purchases. *American Economic Review*, *103*(1), 220–56.

Caplin, A., & Leahy, J. (2001). Psychological expected utility theory and anticipatory feelings. *The Quarterly Journal of Economics*, *116*(1), 55–79.

Cawley, J., & Frisvold, D. (2015). *The Incidence of Taxes on Sugar-Sweetened Beverages: The Case of Berkeley, California* (tech. rep.). National Bureau of Economic Research.

Chaloupka, F. (1991). Rational Addictive Behavior and Cigarette Smoking. *Journal of Political Economy*, *99*(4), 722–742.

Chan, W., Courty, P., & Hao, L. (2009). Suspense: Dynamic incentives in sports contests. *The Economic Journal*, *119*(534), 24–46.

Chavas, J. P. (1983). Structural Change in the Demand for Meat. *American Journal of Agricultural Economics*, *65*, 148–53.

Chay, K. Y., & Greenstone, M. (2005). Does air quality matter? Evidence from the housing market. *Journal of political Economy*, *113*(2), 376–424.

Chetty, R., Looney, A., & Kroft, K. (2009). Salience and Taxation: Theory and Evidence. *American Economic Review*, *99*(4), 1145–1177.

Christakis, N. A., & Fowler, J. H. (2008). The collective dynamics of smoking in a large social network. *New England journal of medicine*, *358*(21), 2249–2258.

Chung, K. Y., Derdenger, T. P., & Srinivasan, K. (2013). Economic value of celebrity endorsements: Tiger woods' impact on sales of nike golf balls. *Marketing Science*, *32*(2), 271–293.

Clarke, K. (2016). The gains from dynamic pricing for ticket resellers. *PhD Dissertation, University of Minnesota Department of Applied Economics.*

Cleveland, W. S. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American statistical association*, *74*(368), 829–836.

Colchero, A., Popkin, B., Rivera, J., & Ng, S. (2016). Beverage Purchases from Stores in Mexico Under the Excise Tax on Sugar Sweetened Beverages: Observational Study. *BJM*, *352*, doi: https://doi.org/10.1136/bmj.h6704.

Colchero, A., Rivera, J., Popkin, B., & Ng, S. (2017). In Mexico, Evidence Of Sustained Consumer Response Two Years After Implementing A Sugar-Sweetened Beverage Tax. *Health Affairs*, *36*(3), https://doi.org/10.1377/hlthaff.2016.1231.

Cornelsen, L., & Smith, R. D. (2018). Viewpoint: Soda Taxes–Four Questions Economists Need to Address. *Food Policy*, *74*, 138–142.

Courty, P. (2003a). Some economics of ticket resale. *Journal of Economic Perspectives*, *17*(2), 85–97.

Courty, P. (2003b). Ticket pricing under demand uncertainty. *The Journal of Law and Economics*, *46*(2), 627–652.

Courty, P., & Davey, L. (2020). The impact of variable pricing, dynamic pricing, and sponsored secondary markets in major league baseball. *Journal of Sports Economics*, *21*(2), 115–138.

Currie, J., & Walker, R. (2011). Traffic congestion and infant health: Evidence from e-zpass. *American Economic Journal: Applied Economics*, *3*(1), 65–90.

Debnam, J. (2017). Selection Effects and Heterogeneous Demand Responses to the Berkeley Soda Tax Vote. *American Journal of Agricultural Economics*, *99*(5), 1172–1187.

Ebstein, R. P., Novick, O., Umansky, R., Priel, B., Osher, Y., Blaine, D., Bennett, E. R., Nemanov, L., Katz, M., & Belmaker, R. H. (1996). Dopamine d4 receptor (d4dr) exon iii polymorphism associated with the human personality trait of novelty seeking. *Nature genetics*, *12*(1), 78–80.

Ely, J., Frankel, A., & Kamenica, E. (2015). Suspense and surprise. *Journal of Political Economy*, *123*(1), 215–260.

Fairhall, A. L., Lewen, G. D., Bialek, W., & van Steveninck, R. R. d. R. (2001). Efficiency and ambiguity in an adaptive neural code. *Nature*, *412*(6849), 787–792.

Falbe, J., Rojas, N., Grummon, A. H., & Madsen, K. A. (2015). Higher Retail Prices of Sugar-Sweetened Beverages 3 Months After Implementation of an Excise Tax in Berkeley, California. *American Journal of Public Health*, *105*(11), 2194–2201.

Falbe, J., Thompson, H. R., Becker, C. M., Rojas, N., McCulloch, C. E., & Madsen, K. A. (2016). Impact of the Berkeley Excise Tax on Sugar-Sweetened Beverage Consumption. *American Journal of Public Health*, *106*(10), 1865–1871.

Fernandes, M. (2008). The Effect of Soft Drink Availability in Elementary Schools on Consumption. *Journal of American Dietetic Association*, *108*(9), 1445–52.

Fletcher, J. M., Frisvold, D. E., & Tefft, N. (2010). The Effects of Soft Drink Taxes on Child and Adolescent Consumption and Weight Outcomes. *Journal of Public Economics*, *94*(11), 967–974.

Forrest, D., Simmons, R., & Buraimo, B. (2005). Outcome uncertainty and the couch potato audience. *Scottish Journal of Political Economy*, *52*(4), 641–661.

Fou, A. (2014). Comparative media costs: Cost to reach 1,000 users or impressions [, (last accessed 2020-09-25)]. *Marketing Science Consulting Group, Inc.* Retrieved September 25, 2020, from https://www.slideshare.net/augustinefou/cross-channel-media-costs-comparison-cpm-basis-by-augustine-fou

Fournier, G. M., & Martin, D. L. (1983). Does government-restricted entry produce market power?: New evidence from the market for television advertising. *The Bell Journal of Economics*, 44–56.

Fowdur, L., Kadiyali, V., & Narayan, V. (2009). The impact of emotional product attributes on consumer demand: An application to the us motion picture industry. *Johnson School Research Paper Series*, (22-09).

Franck, E., & Nüesch, S. (2012). Talent and/or popularity: What does it take to be a superstar? *Economic Inquiry*, *50*(1), 202–216.

Friedman, W. (2017). Nfl tv networks see generally higher cpms [, (last accessed 2020-09-25)]. *Television News Daily*. Retrieved September 25, 2020, from https://www.mediapost.com/publications/article/307174/nfl-tv-networks-see-generally-higher-cpms.html

Geanakoplos, J. et al. (1996). The hangman's paradox and newcomb's paradox as psychological games. *COWLES FOUNDATION DISCUSSION PAPER*.

Geanakoplos, J., Pearce, D., & Stacchetti, E. (1989). Psychological games and sequential rationality. *Games and economic Behavior*, *1*(1), 60–79.

Goldstein, N. J., Cialdini, R. B., & Griskevicius, V. (2008). A Room with a Viewpoint: Using Social Norms to Motivate Environmental Conservation in Hotels. *Journal of Consumer Research*, *35*(3), 472–482.

Gregg, E. W., Zhuo, X., Cheng, Y. J., Albright, A. L., Narayan, K. V., & Thompson, T. J. (2014). Trends in Lifetime Risk and Years of Life Lost Due to Diabetes in the USA, 1985–2011: A Modeling Study. *The Lancet Diabetes & Endocrinology*, *2*(11), 867–874.

Grimshaw, S. D., & Larson, J. S. (2020). Effect of star power on nba all-star game tv audience. *Journal of Sports Economics*, *22*(2), 139–163.

Grogger, J. (2017). Soda Taxes and The Prices of Sodas and Other Drinks: Evidence from Mexico. *American Journal of Agricultural Economics*, *99*(2), 481–498.

Gruber, J., & Köszegi, B. (2001). Is Addiction "Rational"? Theory and Evidence. *The Quarterly Journal of Economics*, *116*(4), 1261–1303.

Hartmann, W. R., & Klapper, D. (2018). Super bowl ads. *Marketing Science*, *37*(1), 78–96.

Haugh, M. B., & Singal, R. (2020). How to play fantasy sports strategically (and win). *Management Science*.

Hausman, J. A., & Leonard, G. K. (1997). Superstars in the National Basketball Association: Economic value and policy. *Journal of Labor Economics*, *15*(4), 586–624.

Heindl, K. (2018). Free agency in the new player narrative driven NBA. [link here, (last accessed 2018-11-21)]. *RealGM*. Retrieved November 21, 2018, from https://basketball.realgm.com/analysis/250427/Free-Agency-In-The-New-Player-Narrative-Driven-NBA

Huang, R., & Kiesel, K. (2012). Does Limiting Access to Soft Drinks in Schools Result in Compensation at Home? *European Review of Agricultural Economics*, *39*(5), 797–820.

Huberman, G., & Regev, T. (2001). Contagious Speculation and a Cure for Cancer: A Non-event that Made Stock Prices Soar. *Journal of Finance*, *56*(1), 387–96.

Humphreys, B. R., & Miceli, T. J. (2019). The peculiar preferences of sports fans: Toward a preference-based motivation for the uncertainty of outcome hypothesis. *Journal of Sports Economics*, *20*(6), 782–796.

Investopedia. (2019). How the NBA makes money [link here, (last accessed 2020-01-30)]. *Investopedia*. Retrieved January 30, 2020, from https://www.investopedia.com/articles/personal-finance/071415/how-nba-makes-money.asp

Itti, L., & Baldi, P. (2009). Bayesian surprise attracts human attention. *Vision research*, *49*(10), 1295–1306.

James, J., Thomas, P., Cavan, D., & Kerr, D. (2004). Preventing Childhood Obesity by Consumption of Carbonated Drinks: Cluster Randomized Controlled Trial. *British Medical Journal*, *328*(7450), 1236.

Jane, W.-J. (2016). The effect of star quality on attendance demand: The case of the National Basketball Association. *Journal of Sports Economics*, *17*(4), 396–417.

Jiaqi Xu, J., Fader, P. S., & Veeraraghavan, S. (2019). Designing and evaluating dynamic pricing policies for major league baseball tickets. *Manufacturing & Service Operations Management*, *21*(1), 121–138.

Johnson, B. K., & Rosenbaum, J. E. (2015). Spoiler alert: Consequences of narrative spoilers for dimensions of enjoyment, appreciation, and transportation. *Communication Research*, *42*(8), 1068–1088.

Kahn, L. M. (2000). The sports business as a labor market laboratory. *Journal of Economic Perspectives*, *14*(3), 75–94.

Kaplan, E. H., & Garstka, S. J. (2001). March madness and the office pool. *Management Science*, *47*(3), 369–382.

Kaplan, S. (2020). The economic value of popularity: Evidence from superstars in the national basketball association. *Available at SSRN*.

Knowles, G., Sherony, K., & Haupert, M. (1992). The demand for major league baseball: A test of the uncertainty of outcome hypothesis. *The American Economist*, *36*(2), 72–80.

Knox, B. (2012). Breaking down why the NBA will always be a star-driven league. [link here, (last accessed 2018-11-21)]. *Bleacher Report*. Retrieved November 21, 2018, from https://bleacherreport.com/articles/1329836-breaking-down-why-the-nba-will-always-be-a-star-driven-league

Kremer, M., & Levy, D. (2008). Peer Effects and Alcohol Use Among College Students. *Journal of Economic Perspectives*, *22*(3), 189–206.

Kreps, D. M., & Porteus, E. L. (1978). Temporal resolution of uncertainty and dynamic choice theory. *Econometrica: journal of the Econometric Society*, 185–200.

Krueger, A. B. (2005). The economics of real superstars: The market for rock concerts in the material world. *Journal of Labor Economics*, *23*(1), 1–30.

Leavitt, J. D., & Christenfeld, N. J. (2011). Story spoilers don't spoil stories. *Psychological science*, *22*(9), 1152–1154.

Lee, D., Hosanagar, K., & Nair, H. S. (2018). Advertising content and consumer engagement on social media: Evidence from facebook. *Management Science*, *64*(11), 5105–5131.

Lehmann, E. E., & Schulze, G. G. (2008). What does it take to be a star?-the role of performance and the media for german soccer players. *Applied Economics Quarterly*, *54*(1), 59.

Levin, Y., McGill, J., & Nediak, M. (2009). Dynamic pricing in the presence of strategic consumers and oligopolistic competition. *Management science*, *55*(1), 32–46.

Levine, W. H., Betzner, M., & Autry, K. S. (2016). The effect of spoilers on the enjoyment of short stories. *Discourse processes*, *53*(7), 513–531.

Liu, X., Shum, M., & Uetake, K. (2020). Attentive and inattentive tv viewing: Evidence from baseball telecasts. *Available at SSRN*.

Livingston, J. A., Ortmeyer, D. L., Scholten, P. A., & Wong, W. (2013). A hedonic approach to testing for indirect network effects in the lcd television market. *Applied Economics Letters*, *20*(1), 76–79.

Lusk, J. (2010). The Effect of Proposition 2 on the Demand for Eggs in California. *Journal of Agricultural & Food Industrial Organization*, *8*(1), 1–18.

Luttik, J. (2000). The value of trees, water and open space as reflected by house prices in the Netherlands. *Landscape and urban planning*, *48*(3-4), 161–167.

Madrigal, R., & Bee, C. (2005). Suspense as an experience of mixed emotions: Feelings of hope and fear while watching suspenseful commercials. *ACR North American Advances*.

Mills, B. M., Salaga, S., & Tainsky, S. (2016). Nba primary market ticket consumers: Ex ante expectations and consumer market origination. *Journal of Sport Management*, *30*(5), 538–552.

Morris, D. (2018). NFL vs. NBA: Which will be America's biggest sport 10 years from now? [link here, (last accessed 2018-11-21)]. *Fortune*. Retrieved November 21, 2018, from http://fortune.com/2018/05/26/nfl-vs-nba-americas-biggest-sport/

NBC. (2019). With long endorsement list, lebron james remains highest earning nba player [link here, (last accessed 2020-01-30)]. *NBC Sports*. Retrieved January 30, 2020, from https://nba.nbcsports.com/2019/10/23/with-long-endorsement-list-lebron-james-remains-highest-earning-nba-player/

Olekalns, N., & Bardsley, P. (1996). Rational Addiction to Caffeine: an Analysis of Coffee Consumption. *Journal of Political Economy*, *104*(5), 1100–1104.

Oskam, J., van der Rest, J.-P., & Telkamp, B. (2018). What's mine is yours—but at what price? dynamic pricing behavior as an indicator of airbnb host professionalization. *Journal of Revenue and Pricing Management*, *17*(5), 311–328.

Paarlberg, R., Mozaffarian, D., & Micha, R. (2017). Viewpoint: Can U.S. local soda taxes continue to spread? *Food Policy*, *71*, 1–7.

PWC. (2019). Perspectives from the global entertainment & media outlook 2019-2023 [link here, (last accessed 2020-06-01)]. *Price Waterhouse Cooper Reports*. Retrieved June 1, 2020, from https://www.pwc.com/gx/en/entertainment-media/outlook-2019/entertainment-and-media-outlook-perspectives-2019-2023.pdf

Ranganath, C., & Rainer, G. (2003). Neural mechanisms for detecting and remembering novel events. *Nature Reviews Neuroscience*, *4*(3), 193–202.

Real, K., & Rimal, R. N. (2007). Friends Talk to Friends about Drinking: Exploring the Role of Peer Communication in the Theory of Normative Social Behavior. *Health Communication*, *22*(2), 169–180.

Resnik, A., & Stern, B. L. (1977). An analysis of information content in television advertising. *Journal of marketing*, *41*(1), 50–53.

Roberto, C., Wong, D., Musicus, A., & Hammond, D. (2016). The Influence of Sugar-Sweetened Beverage Health Warning Labels on Parents' Choices. *Pediatrics*, *137*(2), e20153185.

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of political economy*, *82*(1), 34–55.

Rosen, S. (1981). The economics of superstars. *The American economic review*, *71*(5), 845–858.

Rosen, S., & Sanderson, A. (2001). Labour markets in professional sports. *The economic journal*, *111*(469), F47–F68.

Rosenquist, J. N., Murabito, J., Fowler, J. H., & Christakis, N. A. (2010). The spread of alcohol consumption behavior in a large social network. *Annals of internal medicine*, *152*(7), 426–433.

Rottenberg, S. (1956). The baseball players' labor market. *Journal of political economy*, *64*(3), 242–258.

Ryoo, J. H., Wang, X., & Lu, S. (2020). Do spoilers really spoil? using topic modeling to measure teh effect of spoiler reviews on box office revenue. *Journal of Marketing, Forthcoming*.

Salaga, S., & Tainsky, S. (2015). Betting lines and college football television ratings. *Economics Letters*, *132*, 112–116.

Sallee, J. M., West, S. E., & Fan, W. (2016). Do consumers recognize the value of fuel economy? Evidence from used car prices and gasoline price fluctuations. *Journal of Public Economics*, *135*, 61–73.

Scarfe, R., Singleton, C., & Telemo, P. (2020). Extreme wages, performance and superstars in a market for footballers. *Performance and Superstars in a Market for Footballers (March 31, 2020)*.

Schlenker, W., & Villas-Boas, S. B. (2009). Consumer and Market Responses to Mad Cow Disease. *American Journal of Agricultural Economics*, *91*(4), 1140–52.

Scully, G. W. (1974). Pay and performance in major league baseball. *The American Economic Review*, *64*(6), 915–930.

Silver, L. D., Ng, S. W., Ryan-Ibarra, S., Taillie, L. S., Induni, M., Miles, D. R., Poti, J. M., & Popkin, B. M. (2017). Changes in Prices, Sales, Consumer Spending, and Beverage Consumption One Year After a Tax on Sugar-sweetened Beverages in Berkeley, California, US: A Before-and-After Study. *PLoS Medicine*, *14*(4), e1002283.

Simonov, A., Ursu, R., & Zheng, C. (2020). Do suspense and surprise drive entertainment demand? evidence from twitch. tv. *Evidence from Twitch. tv (October 14, 2020)*.

Sponsorship.com. (2018). Sponsorship spending on the NBA totals 1.12 billion dollars in 2017-2018 season [link here, (last accessed 2020-01-30)]. *Sponsorship.com*. Retrieved January 30, 2020, from http://www.sponsorship.com/Report/2018/05/14/Sponsorship-Spending-On-The-NBA-Totals-$1-12-nbsp;.aspx

Sports-Illustrated. (2014). NBA announces 9-year TV deal with ESPN, Turner Sports [link here, (last accessed 2020-01-30)]. *Sports Illustrated*. Retrieved January 30, 2020, from https://www.si.com/nba/2014/10/06/new-nba-tv-deal-worth-24-billion

Statista. (2014). Length of tv ad time in 120 minutes of selected sporting events in united states in 2013 and 2014 [, (last accessed 2020-09-25)]. *Statista*. Retrieved September 25, 2020, from https://www.statista.com/statistics/316507/length-tv-ad-time-sports-events-usa/

Statista. (2019). Gate receipts as percentage of total revenue in the National Basketball Association [link here, (last accessed 2020-01-30)]. *Statista*. Retrieved January 30, 2020, from https://www.statista.com/statistics/193410/percentage-of-ticketing-revenue-in-the-nba-since-2006/

Stigler, G. J., & Becker, G. S. (1977). De gustibus non est disputandum. *The american economic review*, *67*(2), 76–90.

Su-lin, G., Tuggle, C. A., Mitrook, M. A., Coussement, S. H., & Zillmann, D. (1997). The thrill of a close game: Who enjoys it and who doesn't? *Journal of Sport and Social Issues*, *21*(1), 53–64.

Sweeting, A. (2012). Dynamic pricing behavior in perishable goods markets: Evidence from secondary markets for major league baseball tickets. *Journal of Political Economy*, *120*(6), 1133–1172.

Van Ravenswaay, E. O., & Hoehn, J. P. (1991). The Impact of Health Risk Information on Food Demand: A Case Study of Alar and Apples. In J. Caswell (Ed.), *Economics of food safety* (p. 356). Elsevier Bioscience TP373.5.E26.

Waters, T. M., & Sloan, F. A. (1995). Why Do People Drink? Tests of the Rational Addiction Model. *Applied Economics*, *27*(8), 727–736.

Whitehead, T. (2017). NBA teams are resting players earlier and earlier. [link here, (last accessed 2018-11-21)]. *FiveThirtyEight*. Retrieved November 21, 2018, from https://fivethirtyeight.com/features/nba-teams-are-resting-players-earlier-and-earlier/

Wilbur, K. C. (2008). A two-sided, empirical model of television advertising and viewing markets. *Marketing science*, *27*(3), 356–378.

Williams, K. R. et al. (2017). *Dynamic airline pricing and seat availability* (tech. rep.). Cowles Foundation for Research in Economics, Yale University.

Yang, Y., & Shi, M. (2011). Rise and fall of stars: Investigating the evolution of star status in professional team sports. *International Journal of Research in Marketing*, *28*(4), 352–366.

Yang, Y., Shi, M., & Goldfarb, A. (2009). Estimating the value of brand alliances in professional team sports. *Marketing Science*, *28*(6), 1095–1111.

Yen, S. T., & Jensen, H. H. (1996). Cholesterol Information and Egg Consumption in the US: A Nonnormal and Heteroscedastic Double-Hurdle Model. *European Review of Agricultural Economics*, *23*(3), 343–56.

Appendix A

Entertainment Utility from Skill and Thrill

A.1 Additional Empirical Strategies for Assessing Thrill

Stakes-Dependency

To understand the interplay between suspense and the stakes of an event, I examine viewership responses to suspense in regular season versus playoff games. The stakes are much higher in playoff games, since a single win or loss carries substantially higher consequences than a single win or loss during the regular season.[1] The empirical strategy to analyze stakes and suspense will be an extension of the strategy for studying viewership responses to suspense in general. I use the following estimating equation:

$$V_{it} = \gamma(|D_{it}| * \text{Playoffs}_i) + (|D_{it}| * \mathbf{Q_{it}})\mathbf{\Gamma} + (|D_{it}| * \mathbf{Q_{it}} * \text{Playoffs}_i)\mathbf{\Lambda} + \alpha_i + \eta_t + \epsilon_{it} \quad \text{(A.1)}$$

In equation A.1, $\mathbf{\Lambda}$ represents the vector of time-varying, differential impacts of score differential during the playoffs on viewership. Again, if stakes are important in heightening sensitivity to suspense, estimates in $\mathbf{\Lambda}$ should be negative and increasingly large for later segments of a game.

Underdog Margin

Underdog margin represents the score differential in reference to the "underdog," which is the team not favored to win a game at the onset. Thus, the underdog margin variable can

[1] The regular season schedule in the NBA consists of 82 games. Thus, the marginal contribution of each game to a team's final record and playoff chances is quite low.

be positive (if the underdog has more points than the favored team) or negative. Again, I estimate impacts of the underdog margin on viewership while controlling for absolute score differential at different stages of a game.

$$V_{it} = (\text{UnderdogMargin}_{it} * \mathbf{Q_{it}})\mathbf{\Lambda} + (|D_{it}| * \mathbf{Q_{it}})\mathbf{\Gamma} + \alpha_i + \eta_t + \epsilon_{it} \tag{A.2}$$

Consecutive Points Scored

Another element of surprise, particularly in sporting events or other types of competitions, is the "run effect." This effect takes place when one team performs in a way that during a specific portion of the game, there is relatively large updating in beliefs about an outcome. A useful proxy for the run effect is the total number of consecutive points scored by a single team during a specific portion of the game.

To estimate the impact of this effect on viewership, I subset games to those that had a run of at least Z consecutive points scored by a single team during a single period of the game, testing impacts for different values of Z.[2] Again, I estimate impacts of consecutive points scored on viewership for games with ConsecPoints $> Z$ and during different segments of the game, controlling for absolute score differential.

$$V_{it} = (\text{ConsecPoints}_{it} * \mathbf{Q_{it}})\mathbf{\Lambda} + (|D_{it}| * \mathbf{Q_{it}})\mathbf{\Gamma} + \alpha_i + \eta_t + \epsilon_{it} \tag{A.3}$$

[2]For context, $Z = 15$ makes up approximately 10% of games in the sample while $Z = 13$ approximately 25% of games.

A.2 Additional Estimation Results

Absolute Score Differential

Figure A.1: Individual Viewership Results by Score Differential Bin by Quarter Segment (Level Change)

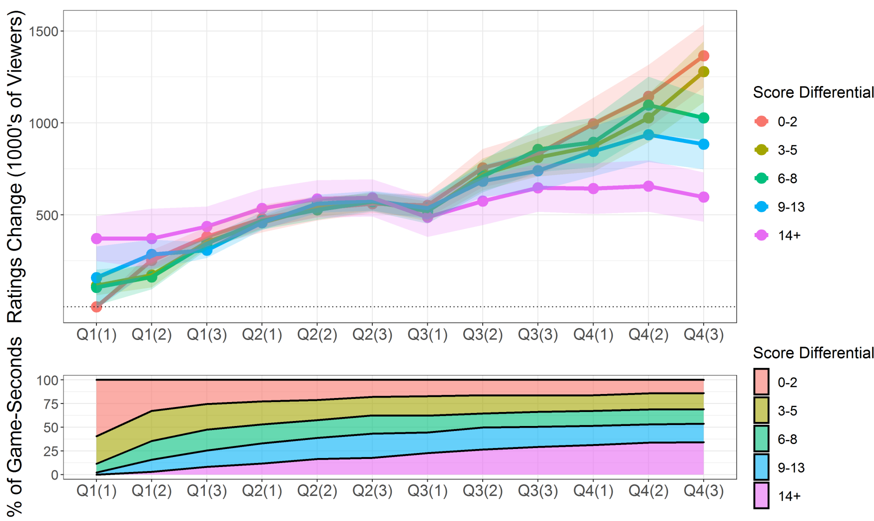

Note: Average 15-minute viewership was 2,694,597 individuals across entire sample.

Figure A.2: Individual Viewership Results by Score Differential Percentile (Within Quarter Segment) by Quarter Segment

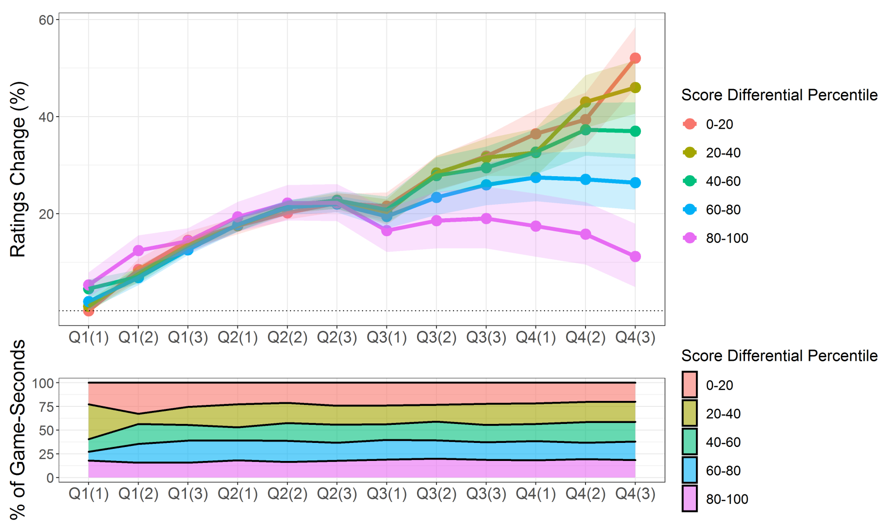

Note: Average 15-minute viewership was 2,694,597 individuals across entire sample.

Figure A.3: Individual Viewership Results by Score Differential Percentile (Within Quarter Segment) by Quarter Segment (Level Change)

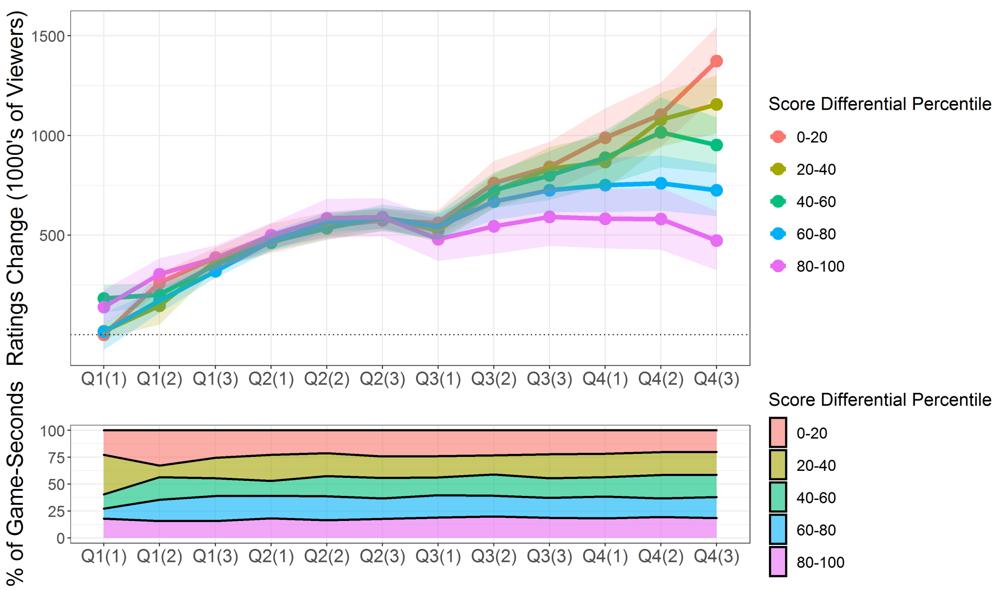

Note: Average 15-minute viewership was 2,694,597 individuals across entire sample.

Figure A.4: Individual Viewership Results by Score Differential Bin by Quarter Segment (Tails, Level Change)

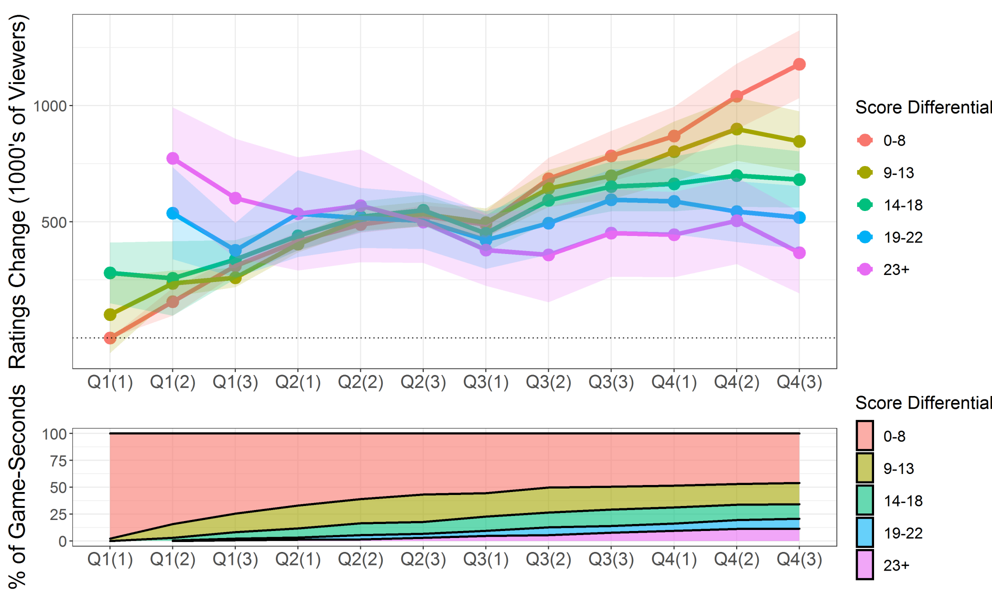

Stakes-Dependency

Here, I examine the stakes associated with an event, hypothesizing that games with different stakes (i.e. playoff games versus regular season games) exhibit different viewership responses to suspense. Table A.1 presents the results examining viewership responses to absolute score differential in playoff versus non-playoff games. Column (1) presents the naive estimation and column (2) the heterogeneous by period results aimed to capture suspense-driven effects. Both specifications include game and quarter segment fixed effects.

Table A.1: Impact of Stakes on TV Ratings

	Dependent Variable: log(Total Proj. Viewers Watching)	
Absolute Score Diff.	-0.0078^{***}	0.0001
	(0.0018)	(0.0021)
Absolute Score Diff. * Q2		-0.0002
		(0.0019)
Absolute Score Diff. * Q3		-0.0049^{*}
		(0.0025)
Absolute Score Diff. * Q4		-0.0105^{***}
		(0.0030)
Absolute Score Diff. * Playoffs	0.0069^{**}	0.0057
	(0.0030)	(0.0043)
Absolute Score Diff. * Q2 * Playoffs		-0.0044
		(0.0036)
Absolute Score Diff. * Q3 * Playoffs		-0.0052
		(0.0046)
Absolute Score Diff. * Q4 * Playoffs		-0.0062
		(0.0048)
Period x Playoff Controls	No	Yes
Game FE	Yes	Yes
Quarter Segment FE	Yes	Yes
Observations	1,381,357	1,381,357
R^2	0.9453	0.9497
Adjusted R^2	0.9453	0.9497
Note:		*p<0.1; **p<0.05; ***p<0.01

In column (1), the term of interest is the interaction between Absolute Score Diff. and Playoffs, where Playoffs is an indicator variable = 1 for playoff games (and = 0 for regular season games). One can see that on average for regular season games, a one-point increase in the absolute score differential decreases viewership by 0.78%, while for playoff games that effect is not statistically different from zero. These findings suggest that viewers are much more responsive to how close a game is if it takes place during the regular season versus the playoffs. One potential mechanism is that for select subsets of viewers, substitution from nationally televised games towards local market games is possible during the regular season, where there are times when nationally-televised games overlap with strictly locally televised games. On the other hand, in the playoffs all games are nationally-televised and there is almost no overlap of games.[3]

In column (2), I examine heterogeneous viewership impacts of absolute score differential across periods in playoff versus regular season games. Thus, if viewers respond to suspense differently when stakes are higher, the effects would be witnessed in the triple interaction terms. One can see that for regular season games, the interaction between absolute score differential and time remaining in the game has the same magnitude and direction of impacts as those in Table 1.3. In addition, the sign of the interaction between Absolute Score Diff. and Playoffs is positive similar to column (1), but no longer statistically significant. Importantly, examining this coefficient along with the coefficients on the triple interaction terms suggests that playoff-level stakes do not statistically significantly enhance the viewership response to suspense, although the signs and ordering of the coefficients are in the expected direction.

Underdog Margin

Table A.2 presents results from estimations of the impact of underdog margin on viewership. Columns (1) and (2) examine the naive estimation of impacts of underdog margin on viewership, while columns (3) and (4) estimate differential effects by quarter and represent the model in equation A.2. Columns (2) and (3) control for the average impact of score differential on viewership, while column (4) controls for the differential impacts by quarter segment. In columns (1) and (2), one can see that the average impact of underdog margin on viewership is positive and statistically significant, whereas column (2) suggests that a one point increase in the underdog score differential margin increases viewership by 0.21%. Thus, the naive model suggests that viewers do respond positively as the score differential margin favors the underdog.

In columns (3) and (4), it is clear that the impact of the underdog margin on viewership depends on the time remaining in a game, where less time remaining increases the impact of the underdog margin on viewership, which is indicative of surprise. The effects are actually quite different across the game and in the expected direction based on the definition of surprise – the impact of a one point increase in the underdog margin in the fourth quarter on viewership is nearly double the impact witnessed in the third quarter. Again, this can be explained by the notion that the impact of a marginal underdog score differential change

[3]The only exception to this is during weekdays when there are 3 games scheduled. Since all games begin after a certain time, there is small overlap between games, where one of the games is typically shown on a less-viewed national network (e.g. NBA TV).

Table A.2: Impact of Underdog Margin on TV Ratings

	Dependent Variable: log(Total Proj. Viewers Watching)			
Underdog Margin	0.0033***	0.0021**	−0.0008	−0.0002
	(0.0009)	(0.0009)	(0.0016)	(0.0016)
Underdog Margin * Q2			0.0005	0.0007
			(0.0011)	(0.0011)
Underdog Margin * Q3			0.0023	0.0018
			(0.0015)	(0.0015)
Underdog Margin * Q4			0.0044**	0.0026
			(0.0017)	(0.0016)
Score Differential Control	No	Yes	Yes	Yes
Score Differential x Quarter Segment Control	No	No	No	Yes
Game FE	Yes	Yes	Yes	Yes
Quarter Segment FE	Yes	Yes	Yes	Yes
Observations	1,381,357	1,381,357	1,381,357	1,381,357
R^2	0.9440	0.9450	0.9455	0.9470
Adjusted R^2	0.9440	0.9449	0.9455	0.9470

Note: *p<0.1; **p<0.05; ***p<0.01

in later stages of a game leads to wider swings in outcome probabilities. When accounting for heterogeneous absolute score differential controls (column 4), the statistical significance disappears, however the signs and magnitudes of the coefficients by quarter support the argument that underdog margin provides meaningful surprise that viewers react to in the expected way. The magnitudes of the estimates are smaller than those seen in characteristics of suspense, albeit still economically meaningful. A one-standard deviation increase in the underdog margin (approximately 10 points) increases viewership 1.6-3.7%, where the effects are monotonically increasing as a game reaches its end.

Consecutive Points Scored

Next I examine the impact of the "run effect," as measured by consecutive points scored by a single team during a specific portion of the game, on viewership. Table A.3 presents results from both the naive (columns 1 and 2) and surprise-focused (columns 3 and 4) estimations. As pointed out in section 1.4, this analysis only includes games with $Z \geq 15$ points, which makes up approximately 10% of all games in this sample.

One can see in the naive estimation, a one-point increase in consecutive points scored during a game leads to an average increase in viewership of 0.35-0.46%. These average

Table A.3: Impact of Consecutive Points Scored on TV Ratings

	Dependent Variable: log(Total Proj. Viewers Watching)			
Consecutive Points	0.0035*	0.0046*	0.0029	−0.0050
	(0.0019)	(0.0022)	(0.0032)	(0.0033)
Consecutive Points * Q2			−0.0009	0.0056
			(0.0036)	(0.0033)
Consecutive Points * Q3			−0.0020	0.0058
			(0.0028)	(0.0032)
Consecutive Points * Q4			0.0075*	0.0154***
			(0.0035)	(0.0043)
Score Differential Control	No	Yes	Yes	Yes
Score Differential x Quarter Segment Controls	No	No	No	Yes
Game FE	Yes	Yes	Yes	Yes
Quarter Segment FE	Yes	Yes	Yes	Yes
Observations	170,959	170,959	170,959	170,959
R^2	0.9410	0.9431	0.9434	0.9483
Adjusted R^2	0.9409	0.9431	0.9434	0.9482

Note: *p<0.1; **p<0.05; ***p<0.01

effects suggest that people enjoy watching teams go on runs, but do not tell a story about runs driving surprise. In the estimations in columns (3) and (4), which draw from the model presented in equation A.3, one can see that consecutive points scored does not have a significant impact on viewership during the first three quarters, but that runs during the fourth quarter are differentially appealing to viewers. One can see that a one-point increase in consecutive points scored during the fourth quarter results in an approximately 0.75-1.5% increase in viewership. The time-dependent nature of the run effect connects directly to the definition of surprise – runs in the fourth quarter are likely to lead to larger swings in outcome probabilities compared to runs in earlier parts of games, and the relationship is almost completely monotonic. So, for a 15-point run in the fourth quarter, viewership increases during that portion of the game by approximately 15% compared to a game in the first quarter without any such run.[4]

[4]The assumed mechanism for these viewership changes is through individuals keeping track of games while not watching (on their phones, for instance) or receiving notifications updating them about a game, and tuning in as a result of this surprise event. In addition, because the viewership data is at the 15-minute interval level, it may be the case that I'm capturing viewers that opt-in at the end of or after a specific run, but still fall into the same 15-minute rating interval. The viewership data provides the average number of viewers during a 15-minute interval, and so these results may be an underestimate of the true viewership impact of a run.

Table A.4 provides robustness by looking at games with $Z \geq 13$ points, making up approximately 25% of all games in the sample. The results are consistent across the two sub-samples.

Table A.4: Impact of Consecutive Points Scored on TV Ratings (Games with Maximum Consecutive Points > 12)

	Dependent Variable: log(Total Proj. Viewers Watching)			
Consecutive Points	0.0023*	0.0033**	0.0025	−0.0039
	(0.0012)	(0.0015)	(0.0020)	(0.0028)
Consecutive Points * Q2			−0.0016	0.0034
			(0.0021)	(0.0028)
Consecutive Points * Q3			−0.0016	0.0043
			(0.0017)	(0.0027)
Consecutive Points * Q4			0.0055**	0.0116***
			(0.0022)	(0.0034)
Score Differential Control	No	Yes	Yes	Yes
Score Differential x Quarter Segment Controls	No	No	No	Yes
Game FE	Yes	Yes	Yes	Yes
Quarter Segment FE	Yes	Yes	Yes	Yes
Observations	367,375	367,375	367,375	367,375
R^2	0.9324	0.9347	0.9349	0.9386
Adjusted R^2	0.9323	0.9347	0.9348	0.9386

Note: *p<0.1; **p<0.05; ***p<0.01

Appendix B

The Economic Value of Popularity: Evidence from Superstars in the National Basketball Association

Figure B.1: Distribution of All Unique Absence Announcement-by-Matchup Pairs (for Starting-Caliber Players) by Hours to Game

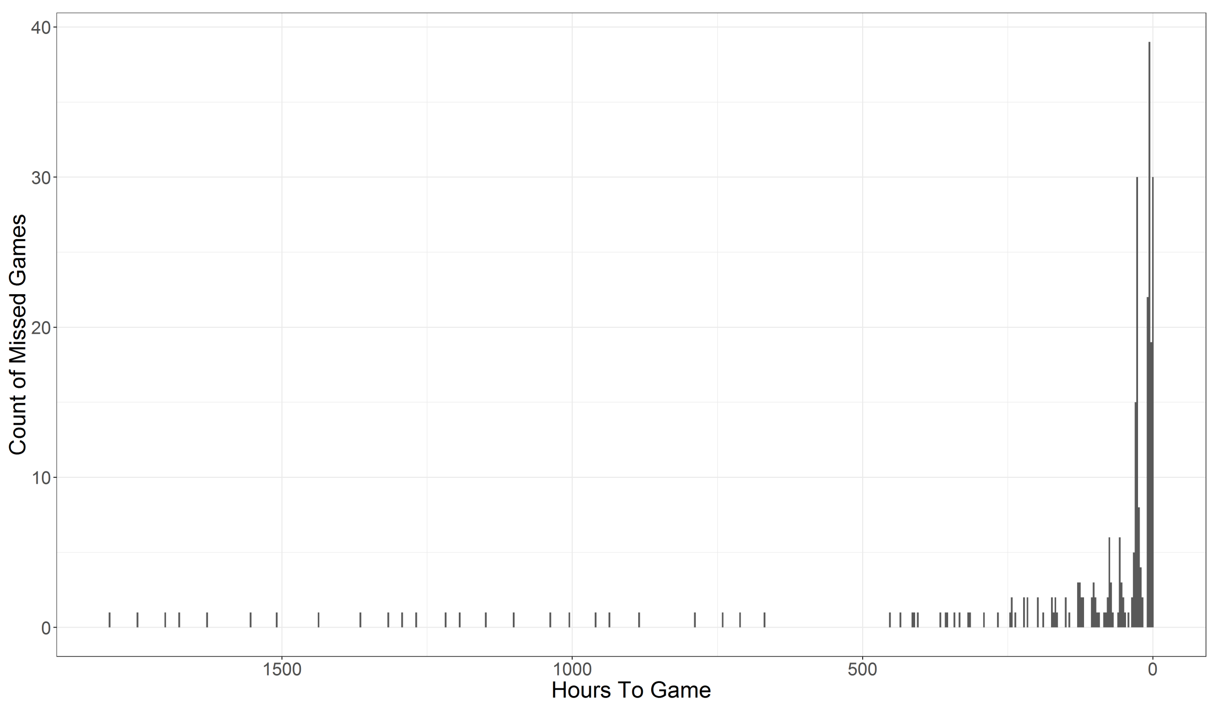

Figure B.2: Difference-in-Differences Results by All-Star Votes

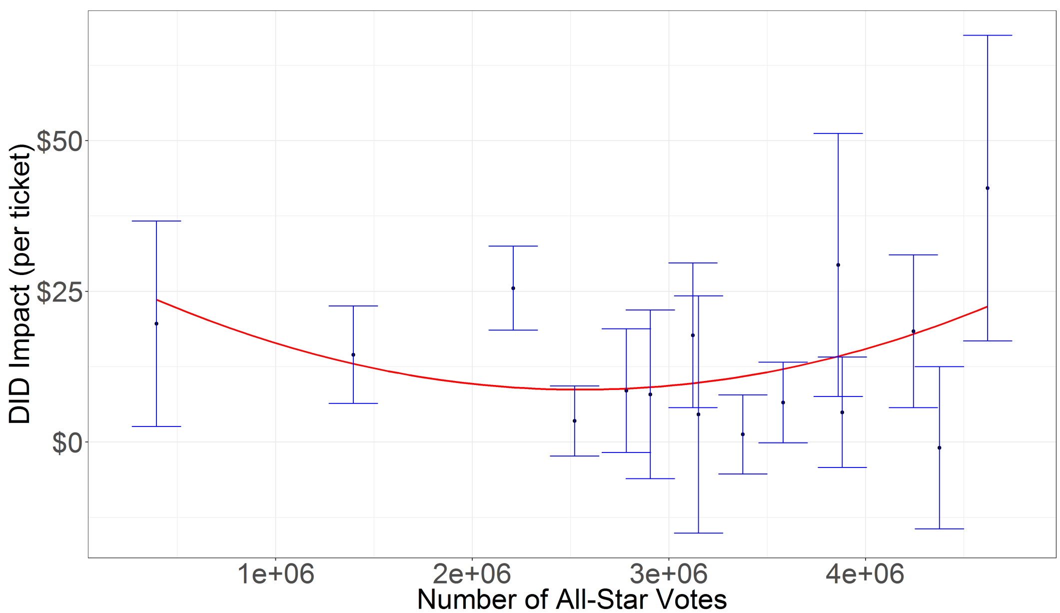

Figure B.3: Event Study Results for Other Superstars

(a) Event Study for Joel Embiid

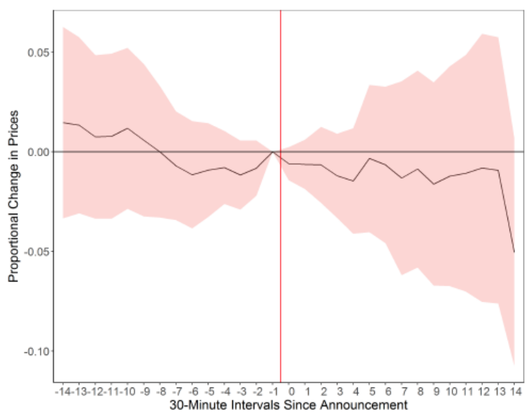

(b) Event Study for James Harden

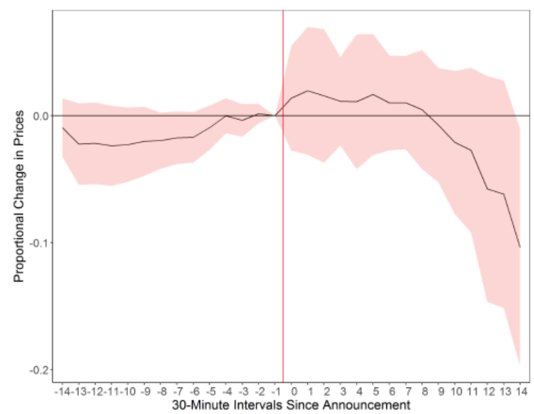

(c) Event Study for Giannis Antetokounmpo

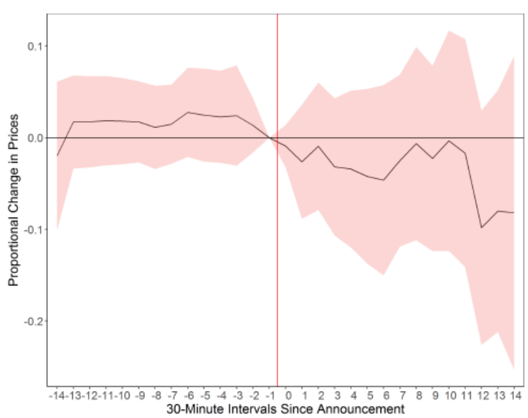

(d) Event Study for Dirk Nowitzki

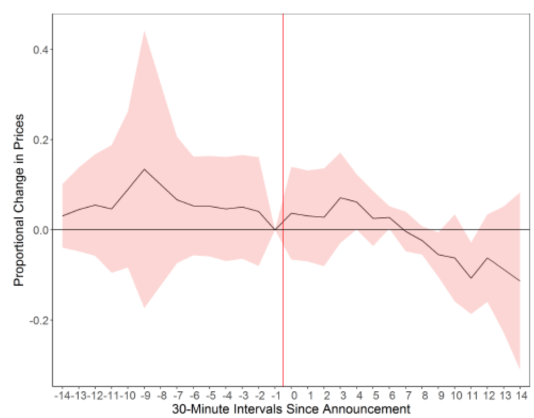

(e) Event Study for Derrick Rose

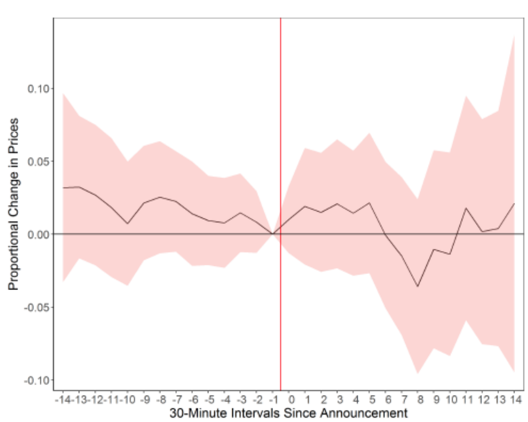

(f) Event Study for Anthony Davis

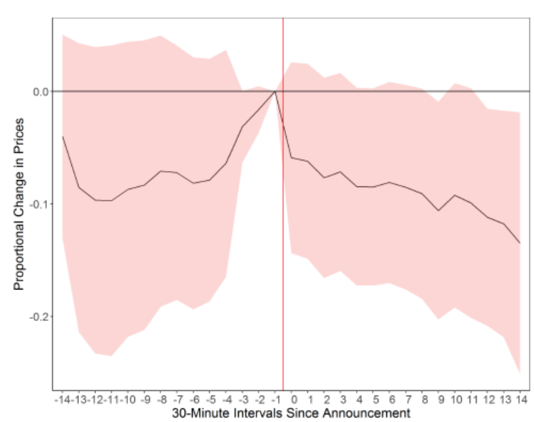

Figure B.4: Event Study Results for Other Superstars

(a) Event Study for Paul George

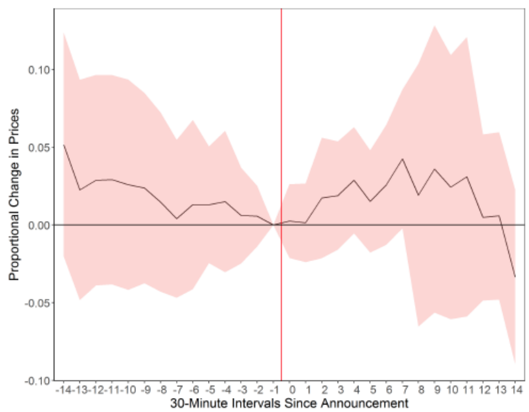

(b) Event Study for Luka Doncic

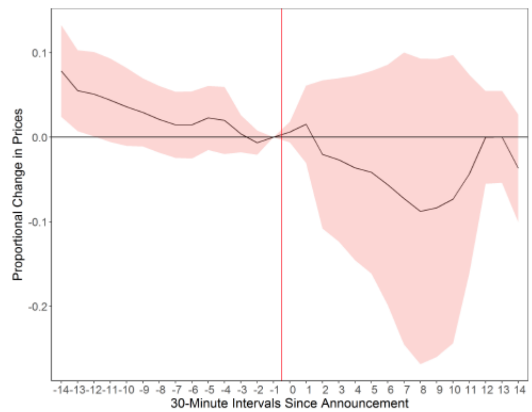

(c) Event Study for Kyrie Irving

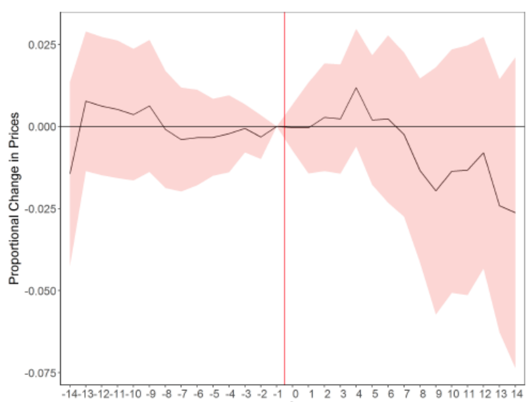

(d) Event Study for Kemba Walker

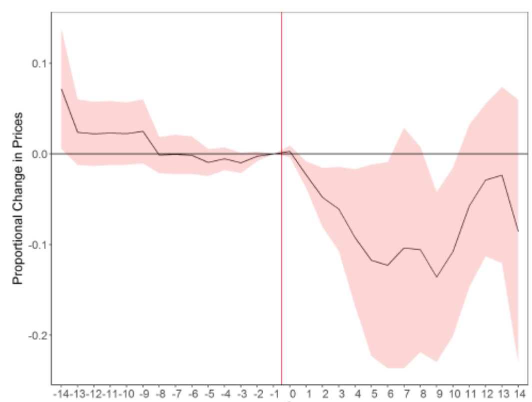

Figure B.5: Difference-in-Differences Estimator by Home vs. Away Matchup Absence (Percentage Change in Prices)

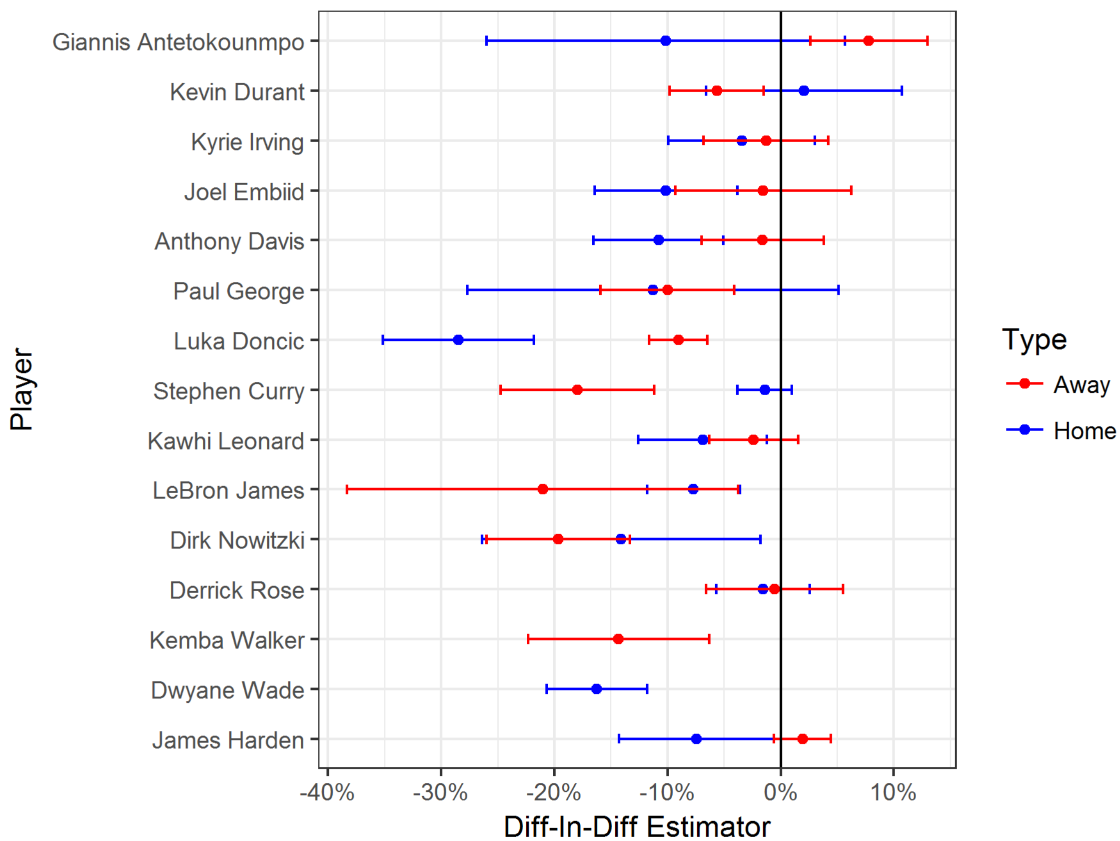

Appendix C

Soda Wars: The Effect of a Soda Tax Election on University Beverage Sales

Table C.1: Replication of Table 3.2, excluding Diet Drinks

	(1) Oz Sold	(2) Log Oz Sold
Soda × Campaign	6663.962	0.088
	(12515.253)	(0.091)
Soda × Post-Election	-3996.684	-0.094
	(10883.869)	(0.100)
Soda × Post-Policy City	-9420.324	-0.266**
	(9778.545)	(0.116)
Soda × Post-Policy Campus	-13156.130	-0.277*
	(12476.123)	(0.148)
Mean of Dep. Variable	123120.144	11.312
Num of Obs.	336	336
R squared	0.865	0.948
Product Group FE	X	X
Month-of-Sample FE	X	X

Standard errors in parentheses are clustered at the product group by academic year level. The outcome variable is ounces sold of product group i in month m, in logs (column 1) and in levels (column 2). Asterisks indicate the following: $*p < 0.10$, $**p < 0.05$, $***p < 0.01$

Figure C.1: Measure D Campaign Advertisements

(a) Yes on D Advertisement

(b) Yes on D Advertisement

(c) No on D Advertisement

(d) No on D Advertisement

Sources: (a) Berkeleyside, Online, (b) Clancey Bateman, MPH, Online. (c) Berkeleyside, Online. (d) MotherJones, Online. *Accessed February 18, 2018.*