



HAL
open science

Three essays on the impact of digitization on economic behavior

Théo Marquis

► **To cite this version:**

Théo Marquis. Three essays on the impact of digitization on economic behavior. Economics and Finance. Université Paris-Saclay, 2022. English. NNT : 2022UPASI006 . tel-03950296

HAL Id: tel-03950296

<https://theses.hal.science/tel-03950296>

Submitted on 21 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Trois Essais sur l'impact de la numérisation sur les comportements économiques

*Three Essays on the Impact of Digitization on Economic
Behaviors*

Thèse de doctorat de l'université Paris-Saclay

École doctorale n°630, Droit Economie Management (DEM)
Spécialité de doctorat: Sciences Économiques
Unité de recherche: Université Paris-Saclay, Réseaux Innovation Territoires et Mondialisation,
92330, Sceaux, France.
Réfèrent: Faculté de droit, économie et gestion
Graduate School: Economie & Management

Thèse préparée dans l'unité de recherche RITM,
sous la direction de **Fabrice Le Guel**, maître de conférences HDR
et la co-direction de **Grazia Cecere**, professeure des universités

Thèse soutenue à Paris-Saclay, le 13 décembre 2022, par

THÉO MARQUIS

Composition du Jury

Membres du jury avec voix délibérative

José DE SOUSA Professeur, Université Paris-Saclay	Président
Emeric HENRY Professeur, Sciences Po	Rapporteur & Examineur
Thierry PENARD Professeur, Université Rennes 1	Rapporteur & Examineur
Laurence DEVILLERS Professeure, Université Paris-Sorbonne	Examinatrice

Titre: Trois Essais sur l'Impact de la Numérisation sur les Comportements Économiques

Mots clés: Économie Numérique, Économie Comportementale, Économie Expérimentale

Résumé: Cette thèse traite des impacts de la numérisation sur les comportements économiques. Le chapitre 1 traite de la réduction des coûts de transport et de recherche et de son impact sur le marché du travail en ligne. Ces dernières années, on a assisté à un passage croissant du travail en personne au travail à distance, ce qui a créé de nouveaux emplois et un nouvel environnement de travail dans lequel les travailleurs peuvent être plus enclins à gérer leurs heures de travail et les compromis entre travail et loisirs. J'utilise l'introduction du jeu mobile Pokémon Go, largement répandu, pour observer son impact sur les comportements des travailleurs sur la plateforme de travail en ligne, Amazon Mechanical Turk. Je montre qu'une augmentation de l'utilisation relative de Pokémon Go entraîne une diminution quotidienne de la proportion de travailleurs américains sur Amazon Mechanical Turk. Le chapitre 2 étudie les résultats d'une expérience de laboratoire sur le terrain menée dans une école primaire française en juillet 2019, où un haut-parleur intelligent, un robot et un

adulte tentaient d'influencer les enfants dans leur choix de partager des billes avec d'autres enfants. Nous avons adapté un jeu de dictateur pour le public d'enfants, puis estimé l'impact de deux stratégies de nudging différentes (proximité sociale et effet de pair) sur le résultat du jeu de dictateur. Au cours de l'interaction avec les enfants, les nudges étaient moins efficaces lorsqu'ils étaient mis en œuvre par l'adulte que par les assistants vocaux, ce qui met en lumière le potentiel de ces dispositifs émergents lorsqu'il s'agit de manipuler leur public vulnérable. Le chapitre 3 étudie les motifs économiques de la manipulation de la popularité sur les médias sociaux par l'acquisition de faux followers par les joueurs de football professionnels. Nous tirons parti de la suppression des comptes suspects de Twitter tenue en juillet 2018 pour approximer les faux followers. Les résultats montrent que les faux followers ont un impact significatif sur la valeur des joueurs, c'est-à-dire sur les frais de transfert, uniquement si le transfert a lieu dans les 1 à 6 mois suivant la création du compte Twitter.

Title: Three Essays on the Impact of Digitization on Economic Behaviors

Keywords: Digital Economics, Behavioral Economics, Experimental Economics

Abstract: This thesis addresses the impacts of digitization on economic behaviors. Chapter 1 deals with the reduction in transportation and search costs and how it impacts the online labor market. In recent years, there has been a growing shift from in-person to remote work, creating new jobs and a new work environment in which workers may be more inclined to manage their work hours and work/leisure trade-offs. I use the introduction of the widely played mobile game Pokémon Go to observe its impact on the behaviors of workers on the online labor platform, Amazon Mechanical Turk. I show that an increase in the relative use of Pokémon Go leads to a daily decrease in the proportion of US workers on Amazon Mechanical Turk. Chapter 2 investigates the results of a lab-in-the-field experiment conducted in a French primary school in July 2019, where a smart speaker, a robot, and an adult were attempting to influence children in their choice of sharing marbles with

other kids. We adapted a dictator game for the children audience and then estimated the impact of two different nudging strategies (Social Proximity and Peer-Effect) on the outcome of the dictator game. During the interaction with the children, the nudges were less effective when implemented by the adult as compared to the voice assistants, shedding light on the potential that these emerging devices have when manipulating their vulnerable audience. Chapter 3 investigates the economic motives to manipulate social media popularity through fake follower acquisition by professional soccer players. We take advantage of Twitter's suspicious account removal held in July 2018 to proxy fake followers. Results show that fake followers impact significantly players' value, i.e., transfer fees, only if the transfer occurs within 1 to 6 months after the Twitter account creation.

Acknowledgements

I am grateful for the financial support of the DATAIA Institute, the University of Paris-Saclay, and the RITM laboratory.

I am deeply indebted to my advisors, Grazia Cecere and Fabrice Le Guel, for their demanding guidance, support, high expectations, and excellent suggestions throughout the years I worked on this thesis. This journey started with an internship during which they started to show me and transmit their passion for research.

I also express my gratitude to my committee members, Emeric Henry and Thierry Pénard, who kindly accepted to be referees, and Laurence Devillers. I want to thank José De Sousa for his participation in the jury, but mainly for the help and advice, he gave me during the thesis. I would also like to thank every assistant professor and professor at RITM, especially Lisa, Serge, Jean, Claire, Miren, and Ugo, for their help, advice, and kindness before and during my Ph.D.

I thank all my colleagues, who have been an endless source of support and joy. I have been lucky to meet inspirational people during my studies and demanding ones who pushed me beyond what I intended and expected. I have no words to describe how much I have benefited from all other past or present Ph.D. students, Kevin and Ismaeel. They gave me support and friendship. A special thanks to Hugo Allouard, Enxhi Leka, Clara Jean, Vincent Lefrere, Alexis Noir-Luhalwe, and Alex Newham for their tremendous help during my Ph.D. and, more importantly, in the end. I'm sure we will keep collaborating in the future.

I send all my love to my family, my mother Carole, my sister Manon and my grandmothers.

Contents

Introduction	1
1 The Work/Leisure Trade-off in Online Labor Markets: The Case of Amazon Mechanical Turk	8
1.1 Introduction	10
1.2 Background	12
1.2.1 <i>Amazon Mechanical Turk</i>	12
1.2.2 <i>Pokémon Go</i>	12
1.3 Literature Review	13
1.4 Data	15
1.5 Empirical Strategy	18
1.5.1 <i>OLS Estimation</i>	18
1.5.2 <i>IV Estimation</i>	19
1.6 Results	20
1.6.1 <i>Impact of Pokémon Go on US workers</i>	20
1.6.2 <i>Robustness Analysis</i>	26
1.7 Conclusion	30
2 Bad Nudge, Kids and Voice Assistants: A Social Preferences lab-in-the-field Experiment	32
2.1 Introduction	34
2.2 Experimental Design	38
2.2.1 The Dictator Game	38
2.2.2 Experiment Design	39
2.2.3 Experimental Procedures (Experiment Setting)	42
2.3 Data	43
2.3.1 Post experimentation verification	45
2.4 Models	46
2.5 Results	47
2.5.1 Mean differences	47
2.5.2 Econometric Results	48

2.5.2.1	Direction of the nudge	48
2.5.2.2	Social Proximity and Peer-effect	50
2.6	Conclusion	52
3	Online Popularity, Fake Followers and Soccer Players' Value	54
3.1	Introduction	56
3.2	Popularity, Value and Online Information Manipulation	59
3.2.1	Superstars' Value: Talent, Popularity and Bargaining Power	59
3.2.2	Cheating Offline and Online: Performance and Online Information Manipulations	63
3.3	Data and Empirical Explorations	66
3.3.1	Player Career and Performance	66
3.3.2	Social Media and Online Popularity	69
3.3.3	Twitter and Fake Followers	70
3.3.4	Alternative Scenario	72
3.4	Transfers, Social Media and Fake Followers	73
3.4.1	The Rise of Instagram	74
3.4.2	Social Media and Transfer	75
3.4.3	Online Popularity Manipulation Strategy	76
3.5	Econometric Model and Descriptive Statistics	78
3.6	Results and Discussion	81
3.6.1	Main Results	81
3.6.2	Robustness Check	85
3.7	Conclusion	88
A	Appendix - The Work/Leisure Trade-off in Online Labor Markets: The Case of Amazon Mechanical Turk	91
A.1	Supplementary Appendix A: Amazon Mechanical Turk operating	91
A.1.1	Amazon Mechanical Turk operation	91
A.2	Supplementary Appendix: Data	93
A.2.1	Amazon Mechanical Turk	93
A.2.2	Pokémon Go	95
A.3	Supplementary Appendix: Robustness Checks & IV test	97
A.3.1	Placebo Test	97
A.3.2	IV Test	98
B	Appendix - Bad Nudge, Kids and Voice Assistants: A Social Preferences lab-in-the-field Experiment	99
B.1	Figure	99
B.1.1	Interlocutor	99

B.1.2	Distribution	100
B.2	Statistics on real participants : by treatment	101
B.3	Possible Selection Bias	101
B.4	Speech of the Interlocutor	102
B.4.1	French	102
B.4.2	English	102
C	Appendix - Online Popularity, Fake Followers and Soccer Players' Value	103
C.1	Sample descriptive statistics	103
C.2	Transfer: years and players' ages	104
C.3	'Real' and fake followers	105
C.4	Account creation date and number of followers on Twitter and Instagram	106
C.5	Twitter account creation and transfers	108
C.6	Results	109

List of Figures

1.1	Proportion of US workers and the Google Trends for Pokémon Go and part-time job searches	17
1.2	Proportion of US workers and the Google Trends for Pokémon Go and part-time job searches	27
1.3	Threshold Effect of Pokémon Go	30
2.1	Proceeding of the experiment by Treatment	42
2.2	Distribution of marbles given at each stage	44
2.3	Distribution of marbles given at the first stage by treatment	45
3.1	Descriptive statistics about Twitter and Instagram	70
3.2	Twitter’s removal of FFs on 12 July 2018	72
3.3	Twitter account creation, FFs and Instagram growth	74
3.4	Number of Twitter accounts created by month (2009-2018)	75
3.5	Distributions of transfers next to Twitter account creation (2009-2019)	77
3.6	Distribution of Transfer fees	85
3.7	Effect of Fake Followers by quantile	86
A.1	Distribution of Mechanical Turk workers in the US by state	93
A.2	Proportion of workers by country	93
A.3	Top iPhone apps in the US in July 2021	95
A.4	Overview of the Pokémon Go user interface	96
B.1	Robot and Smart Speaker	99
B.2	Distribution of marbles given after the first nudge by treatment	100
B.3	Distribution of marbles given after the second nudge by treatment	100
C.1	Transfer types and players’ age	104
C.2	Transfer types by year	104
C.3	Twitter and Instagram account creation dates	106
C.4	Number of followers on Twitter and Instagram (July 2018)	106
C.5	Only Twitter or Instagram account creation dates	107
C.6	Number of followers for only Twitter or Instagram (July, 2018)	107
C.7	Transfers following Twitter account creation (2009-2014)	108

List of Tables

1.1	Variables' description	16
1.2	Summary statistics	17
1.3	Main specifications	21
1.4	Effect on the Time Spent and Amount Earned on AMT	23
1.5	Main specification with IV by demographic group	24
1.6	Main specification by location	25
1.7	Summary statistics	27
1.8	Main specification with IV	28
1.9	Distribution of <i>PokeGo Trend</i> by decile	29
2.1	Timeline of the Experiment	41
2.2	Variables description	43
2.3	Summary statistics	44
2.4	Ttests	45
2.5	Effect of the direction of first nudging	47
2.6	Effect of the direction of second nudging	48
2.7	Control estimation	48
2.8	Main estimation, initial split and after first nudge	49
2.9	Main estimation, after second nudge	50
2.10	Estimation with interactions after first nudge	51
2.11	Estimation with interactions after second nudge	52
3.1	Summary of literature review	63
3.2	Descriptive statistics on transfer types	68
3.3	Descriptive statistics on transfer fees	68
3.4	Correlation between followers and FFs by decile	73
3.5	Twitter account creation and transfer types	76
3.6	Summary statistics	80
3.7	Main specifications	82
3.8	Quantile Regression Specification	87
A.1	Description of tasks available on Amazon Mechanical Turk	92
A.2	Description of categorical variables	94
A.3	Summary statistics: Demographics at an individual level, 96 observations per day	95
A.4	Main specification with placebo explicative variables	98
A.5	IV test	98
B.1	Summary statistics Human	101
B.2	Summary statistics Robot	101
B.3	Summary statistics Smart Speaker	101
B.4	Ttest: Participants vs. Non Participants	101
C.1	Summary statistics of players	103
C.2	Correlation between followers and FFs by ventile	105
C.3	Main specification: control variables results	109

Introduction

Research on digital economics examines whether and how digital technology changes economic activities and behaviors. This thesis through three different empirical works contribute to this research field. As digitization increases the time we spend online, both for work and leisure, its impact on our everyday life must be studied. Understanding the effects of digital technology does not require a fundamentally new economic theory. However, the fact that information is digitized has new influences on its costs and access. The impact of this change is the focus of the digital economics literature. In their literature review, Goldfarb et al. (2019) identified five large cost reductions due to digitization. In search, replication, transportation, tracking, and verification costs. In this dissertation, I address the effect of those reductions in transportation, search, replication, and verification costs through three chapters with three different contexts and research questions. Chapter 1 deals with the reduction in transportation and search costs by studying the trade-off between online work and online leisure performed by online American workers. Chapter 2 addresses the effect of the reduction in search costs by investigating the effect of information delivery through smart speakers and robots to children to change their altruism. Chapter 3 discusses the reduction in replication and verification costs with a study on the effect of the purchase of fake followers on social media on the value of professional soccer players.

The reduction in transportation costs is illustrated, among other things, by the creation of the online labor market (OLM). In this market, and especially on online work platforms, workers perform tasks that so-called requesters post, and their compensation depends on the number and type of tasks they perform. Workers and requesters do not have to be located close to the workplace as it is online. However, even without any cost associated with distance, workers tend to work on platforms closer to them. For example, most workers on Amazon Mechanical Turk (AMT), an American online labor platform, are located in the

United States. Most of the workers on ResponDi, the European AMT equivalent, are European, even though both platforms are accessible in Europe and the US. Same as requesters who prefer local workers when they submit a task. Research on digital economics shows how this transportation cost reduction has resulted in substitution between online and offline products or services in all markets. Duch-Brown et al. (2017) show that this reduction in transportation costs has resulted in substitutions between online and offline sales of personal computers. This substitution exists for news (Gentzkow, 2007), advertising (Seamans et al., 2014), and media consumption (Wallsten, 2013). Blum et al. (2006) show that even for a product with zero shipping costs (visiting websites), people are more likely to visit websites from nearby countries than from faraway countries. Kuhn et al. (2014) show that individuals that used the Internet in job search were more likely to match with an employer. Autor (2001) and Stanton et al. (2016) demonstrate that the reduction in search cost permits a higher screening of workers and that intermediate in the OLM permits a better job allocation. Kroft et al. (2014) show that Craigslist reduced search costs, significantly lowered classified job advertisements in newspapers, caused a significant reduction in the apartment and house rental vacancy rate, and had no effect on the unemployment rate. Chapter 1 deals with the reduction in transportation and search costs and how it impacts the online labor market. In recent years, there has been a growing shift from in-person to remote work, creating new jobs and a new work environment in which workers may be more inclined to manage their work hours and work/leisure trade-offs. I identify a lack in the literature on the impact of this cost reduction on the substitution between online and offline work. There is also a lack of empirical work on the trade-off between work and leisure in the OLM. I address this impact using data collected on Amazon Mechanical Turk. I study how workers work or do not work on the platform depending on the attractiveness of a geo-localized online mobile game and the availability of offline part-time jobs.

The reduction in search cost has been studied in several contexts. Zhang (2018) shows that the reduction in search cost with the rise of Spotify significantly increased long-tailed album sales. This decrease in search costs permits more varied media consumption and leisure activities, according to Boxell et al. (2017a,b). However, this increasing variety does not backfire. Even if the segregation is higher in online news than offline, Gentzkow and Shapiro (2011) show it does not increase over time. Farronato et al. (2018) show that it

permits better matching and increases welfare in some cities or during specific periods (i.e., New Year's Eve) as hosts adapt. The reduction of search costs has also allowed for better access to information. Digitization has made it possible to have access, through an interface and by interacting with it, to online information. Nevertheless, recent advances allow all types of agents, with voice interaction, to request information from machines (smartphones, computers, smart speakers) without going through an interface with written requests. This accessibility can be a blessing, as more people have access to information. However, it can also have drawbacks, as the transmission of information can be done by voice and can be designed to influence the user. Reducing research costs also translates into the increasing use of IA and chatbots to interact with users. They have already been studied and found to be as effective as humans in some contexts and even more effective if they do not reveal their identity. Luo et al. (2019) show chatbots' efficiency in banking services sell. Reduced search costs made devices we interact with by voice and smart speakers more efficient and led to increased sales and usage. Those devices that provide information through only voice interactions return one or a few different pieces of information and can nudge users or influence beliefs. Chapter 2 studies their efficiency in nudging children, compared to humans, in the context of a dictator game. The second chapter relates an experiment with sensible economic agents, children, in which we try to induce them to be more or less altruistic through the intervention of a robot, a smart speaker, and a human.

In economics, signals help economic agents and firms by reducing information asymmetry and allowing them to make more informed decisions. Firms and economic agents who send signals about their high or relative higher quality often allow them to have more bargaining or market power. Ratings, reviews, or popularity measures are indicators of reputation and quality. Waldfogel et al. (2006) show that higher-quality firms may develop more influential brands and, therefore, command higher prices. Tucker et al. (2011) document that popularity information has asymmetrically large effects on niche products. More starkly, Ellison (2011) argues that peer review may be in decline because of low online search costs. In particular, he shows that high-profile researchers do not need to rely on academic journals to disseminate their ideas. They can post online, and people will find their work. In other words, similar to the superstar effect in products, low search costs combined with thousands of research articles benefit the superstar researchers. As feedback from other consumers or

information about products and services is valued, the reduction in verification costs has led to an increase in the use of reviews and other types of evaluation. Chevalier et al. (2006) demonstrate that positive reviews lead to higher sales, and Luca (2016) shows how online restaurant reviews on Yelp affect restaurant demand. Nonetheless, reducing verification costs should induce better quality measurement, but the digitization of ratings and reviews has created opportunities to manipulate online information. Mayzlin et al. (2014) show that if a characteristic is imperfectly observed, there are incentives to manipulate it. They find a positive impact of online information manipulation in the hotel market. The effect of digital information manipulation has been studied in other contexts; Luca and Zervas (2016) found a positive impact of online information manipulation in the restaurant market. Gans et al. (2021) show that Twitter users' influence, through their feedback posted on social media, affects brands, especially when they react to different quality goods. In a digital context, this information can also be manipulated, not only by its means of transmission. In the third chapter of this thesis, we show that the reduction of replication costs due to digitization has created problems with online information manipulation. We use the context of social networks and the purchase of fake followers to show the impact on the popularity and value of "superstar" agents, for whom popularity and talent increase their revenue. I identify in chapter 3 a lack in the literature on the impact of this reduction in verification cost on online information manipulation, especially on online popularity and the possibility of enhancing it unfairly on the value of superstars. I address this impact using data collected on social media and websites that combine data on soccer players' careers. I study how the purchase of fake followers increases the soccer player's transfer value.

In each chapter of this thesis, I use different methodologies for the collection of the data, the implementation of the experimental design, the analysis of the data, and the specification of the econometric models. Each chapter uses a unique and separate database tailored to the specifics of the research question.

For chapter 1, which illustrates the trade-off that online workers can perform between work and leisure, I use a natural experiment design with the release of a geo-localized mobile game. I collected data about the demographics of the online labor market using a python programmed script. These data were collected with a task that ran on Amazon Mechanical

Turk for more than four years (Difallah et al., 2018). I also automated the collection of google trends data to measure the interest in the game. I use a lab-in-the-field experiment design for chapter 2, which studies the smart device's impact on nudging sensible agents. We set up a lab in a school in which we made children play a dictator game. Finally, in chapter 3, which studies the effect of online information manipulation on social media, we use a natural field experiment design, using Twitter's exogenous removal of fake followers. I web-scraped data about international soccer players' performances, careers, and social media metrics using python scripts. In chapters 1 and 3, as these data are unstructured, I use big data tools to clean them up.

Chapter 1 illustrates the trade-off between online work and leisure. This chapter highlights that American online workers, for whom Amazon Mechanical Turk is a secondary source of revenue, are susceptible to exploit the work-leisure trade-off on a daily basis. We use the introduction of the widely played mobile game Pokemon Go to observe its impact on the online labor platform Amazon Mechanical Turk. To estimate whether a mobile game for which players must physically go outside affects the proportion of American workers on the platform, we create an index to measure the intensity of Pokémon Go use. We have two main findings. First, we show that a one standard deviation increase in the relative use of Pokémon Go leads to a daily decrease in the proportion of American workers on Amazon Mechanical Turk of 4.38 percentage points, corresponding to a decrease of up to 1,095 workers. The mechanism identified behind this result is that an increase in the relative use of Pokémon Go decreases the time spent on the platform by under 32 years old workers, without any impact on their earnings. Second, we find that an increase in the search for part-time jobs raises the proportion of American workers on Amazon Mechanical Turk. With the reduction of transportation and search costs due to digitization, we found how the online labor market can be impacted in this chapter. Further research may arise from this result if the non-reduction in productivity can be generalized when workers have more power over managing their working time and how it can affect the general labor supply. It may affect and concentrate the labor supply on some time slot, unbalancing the labor market.

In chapter 2, I study the impact of connected devices using voice recognition as a form of input (Google Home, Alexa, Siri) which are increasingly popular. This mode of interaction

introduces new possibilities to influence the user. We question the extent to which these devices can manipulate their audience, particularly the sensitive population of children. This chapter investigates the results of a lab-in-the-field experiment conducted in a French primary school in July 2019, where a smart speaker, a robot, and an adult were attempting to influence children in their choice of sharing marbles with other kids. We adjusted a dictator game for the children’s audience. We then estimated the impact of two different nudging strategies (Social Proximity and Peer-Effect) on the outcome of the dictator game. Results confirm our first hypothesis that humans should be more persuasive than smart speakers and robots using the Social Proximity strategy (Kosse et al., 2020). Our second hypothesis, which states that smart speakers and robots are more persuasive using the Peer-Effect strategy (Shirado et al., 2017; Vollmer et al., 2018), is also confirmed. During the interaction with the children, the nudges were less effective when they were implemented by the adult as compared to the voice assistants, shedding light on the potential that these emerging devices have when it comes to manipulating their vulnerable audience.

In chapter 3, I study the manipulation of online information in the context of the soccer player’s market. Rosen (1981) and Adler (1985) show that celebrities’ talent and popularity are valuable assets. While talent improvement involves efforts, popularity, and especially online popularity, it can be easily manipulated due to firms selling fake social media followers. Mayzlin et al. (2014) show that if a valuable characteristic is imperfectly measurable, there are incentives to manipulate it. This chapter investigates the economic motives to manipulate social media popularity through fake follower acquisition by professional athletes. Focusing on soccer, we create a unique dataset including 1,075 international players and take advantage of Twitter’s suspicious account removal held in July 2018 to proxy fake followers. Empirical explorations provide evidence of consistent patterns of soccer players associated with fake followers. In their cases, Twitter account creation is quickly followed by a transfer into a new club and predominantly occurred before 2015, while Twitter was more popular than Instagram. Results show that fake followers significantly impact players’ value, i.e., transfer fees, only if the transfer occurs within 1 to 6 months after the Twitter account creation. These fake followers are associated with an average rise of 6% in transfer fees (€650,000). This result is consistent with the acquisition of fake followers to boost the online popularity of a recently created Twitter account and monetize it during transfer bar-

gaining. This online information manipulation can create mistrust, and as verification costs decrease over time with digitization, it should become easier to protect users from this kind of manipulation. However, as replication costs also decrease, there is a need for regulation from platforms.

The results of this thesis have major policy implications regarding how digitization impacts economic behaviors. The first chapter indicates that an online worker can manage his or her time between work and leisure or even reduce working time without suffering a loss of income and, thus, a reduction in productivity. It should lead to a discussion about letting workers manage their working hours and study more specifically with experiments (lab or field) and how long-term productivity and output are impacted. The second chapter finds that smart speakers and, more generally, devices with voice recognition can nudge sensible agents as much as humans. The final chapter shows that online popularity is valued and increases the income and value of superstars. Then there are possibilities for manipulating this online information and having a positive return on investment. This manipulation of online information can lead to mistrust. As the costs of verification decrease over time with digitization, it should become easier to protect against this type of manipulation. Nevertheless, as replication costs also decrease, regulation from platforms and public power is necessary. Otherwise, trust in online information may continue to decline for a portion of the population.

Chapter 1

The Work/Leisure Trade-off in Online Labor Markets: The Case of Amazon Mechanical Turk

Abstract

This article illustrates the trade-off between online work and leisure. We use the introduction of the widely played mobile game Pokémon Go to observe its impact on the behaviors of workers on the online labor platform, Amazon Mechanical Turk. To measure whether a mobile game, for which players must physically go outside, affects the proportion of US workers on the platform, we create an index to measure the intensity of Pokémon Go use. This article highlights that American online workers, for whom Amazon Mechanical Turk is a secondary source of revenue, are susceptible to exploit the work-leisure trade-off on a daily basis. We have two main findings. First, we show that a one standard deviation increase in the relative use of Pokémon Go leads to a daily decrease in the proportion of US workers on Amazon Mechanical Turk of 4.38 percentage points, corresponding to a decrease of up to 1,095 workers. Second, we find that an increase in the search for part-time jobs raises the proportion of US workers on Amazon Mechanical Turk.

1.1 Introduction

In a period where remote work has become widespread, either forced by the health crisis or due to workers' preferences, it is interesting to know how individuals organize their working time. In this article, we link leisure time use with online work. To do so, we use data from the micro-work platform Amazon Mechanical Turk (AMT). Online platforms, such as AMT, allow users to work remotely on tasks or projects posted by employers. The analysis in this article focuses on AMT, one of the biggest online micro-job platforms and the most used in the United States (US). This online platform represents a particular case of remote working; workers have no schedules, only deadlines, and their working hours are therefore more flexible. Each of these online platforms constitutes what we will refer to in this article as an Online Labor Market (OLM). Unlike traditional offline marketplaces, which historically have been highly localized, OLMs allow for long-distance transactions and offer individuals on both sides of the market access to a larger, potentially global, pool of participants for individuals on both sides of the market. OLMs are experiencing considerable growth as they offer new ways to connect employers and employees.

To observe how online workers react to economic changes or exogenous events, we conduct a study of this online market in the US. This market allows us to monitor workers on a daily basis, and then measure the trade-off between work and leisure. This article exploits a natural field experiment, the release of Pokémon Go, a video game that uses real-time player location. We measure how the release and diffusion of this game, for which players have to go physically outside, impact potential AMT workers. This game, which is widespread worldwide, requires workers to leave their work environment, which automatically reduces their presence on the OLM.

This experimental design makes it possible to calculate how workers reach a trade-off between work and leisure, and to observe the impact of this trade-off on the supply of labor. The identification strategy for measuring usage of a leisure activity is inspired by Baker et al. (2017). They used Google Trends to create a Google Job Search Index. Using a similar technique, we created an index to measure the daily usage of Pokémon Go. We also use Google Trends to induce daily part-time job searches.

This article uses daily data about the demographics of AMT workers to measure the variations in the platform's composition. This data on AMT was collected by a daily survey

which was conducted by Difallah et al. (2018). The data covers a three-and-a-half-year period from March 2015 to November 2019. Each observation is constructed from a five-question survey asking workers their age, gender and location. During the three-and-a-half-year period, we measure the daily usage of both Pokémon Go and the search for part-time jobs on Google.

We study the daily variation of the proportion of US workers present on AMT depending on the part-time job and the platform search on Google and on the Pokémon Go index, which measures the daily variation of the importance of the game. For this purpose, we use the OLS and IV 2SLS estimates.

We show that the search for a part-time job on Google increases participation in OLMs. Our findings confirm that workers perform a trade-off between work and leisure. The results show that an increase of one standard deviation (s.d.) in Pokémon Go use leads to a daily decrease in the proportion of US workers on AMT of 4.38 percentage points, corresponding to a decrease of up to 1,095 workers. We exploit the gender and generational differences. The effect is larger on young workers, with no significant difference between young women and young men. We also separate our analysis by area of residence, and we find a larger effect in the most populated areas, where the game is also the most attractive. The mechanism that explains these results is the time spent by workers on the platform, which varies according to the Pokémon Go use and part-time job search. We observe that when the usage of Pokémon Go increases, their time spent on the platform decreases, without affecting the revenue they generate from the platform. To strengthen our results, we conduct several robustness checks. We exclude the release period of the game, which saw a larger number of users, and still find that Pokémon Go has significant effects on online work.

The article is organized as follows. Section 1.2 provides background information on AMT and Pokémon Go. In Section 1.3, we review the literature on online labor markets and on media impact. Section 1.4 presents the data and the main descriptive elements. We describe the econometric modelling in Section 1.5. Section 1.6 presents the results and the robustness checks. Finally, Section 1.7 concludes the article.

1.2 Background

1.2.1 *Amazon Mechanical Turk*

Due to the growing number of workers and employers participating, OLMs are experiencing considerable growth. Indeed, over 10% of the US workforce participates in the online labor market each year.¹ There are several online job platforms, each with its own specifications; some require highly qualified workers, for instance, Freelancer and UpWork, while others, such as Amazon Mechanical Turk (AMT) and Microworkers, do not. The analysis presented in this article focuses on AMT, one of the largest online micro-job platforms, launched by Amazon in November 2005 and the most widely used in the US. An AMT user can be either a “worker” (contractor) or a “requester” (employer). A worker performs short and simple tasks called HITs (Human Intelligent Task) proposed by a requester.² Although Amazon does not publish the number of users registered and active on AMT, Difallah et al. (2018) show that the platform has between 100,000 and 200,000 unique workers, of which 2,000 to 5,000 workers are active on AMT at any given time, which is equivalent to having 10,000 to 25,000 full-time employees.

1.2.2 *Pokémon Go*

Pokémon Go is an augmented reality mobile game in which the player is geo-located and must go outside to participate in the game, and therefore, potentially leave their workplace. As players move through their real-world environment, their avatars move across the game map. It was released in the US on 6 July 2016 and has been downloaded more than a billion times worldwide.³ Pokémon Go generated revenue of \$1.23 billion in 2020, higher than its peak usage in 2016. In addition, over 150 million people play Pokémon Go, higher than in previous years but lower than the 233 million in 2016.⁴ Currently, it still ranks as one of the most grossing applications.⁵ The literature shows that Pokémon Go has reached a large spectrum of the population and potential online workers (Althoff et al., 2016). Figure A.4 in the Appendix shows the game interface with three screenshots taken when the game is running. The second screenshot shows an actual map of the surroundings and the activities

¹GigEconomyData.

²A more detailed explanation of the platform operation is provided in the Appendix (A.1) with Table A.1.

³Niantic via Gamestop.

⁴Niantic via Sensor Tower.

⁵Figure A.3 in the Appendix.

available to the player outdoors. They can capture Pokémon, battle in a gym with their Pokémon, collect virtual items (e.g., Pokéballs), participate in raids with other players and interact with them (trading, sending gifts). Most activities in the games are cooperative, and virtual items and Pokémon are non-rival goods.

1.3 Literature Review

This work contributes to two literature streams. Firstly, we contribute to the OLM literature by analysing the effect of part-time jobs on the online presence and how online workers choose between a leisure activity and being present on the platform. Secondly, to the literature on the media impact by computing the effect of a video game such as Pokémon Go on the online labor participation.

Most studies on online platforms and the gig economy discuss their impact on economic outcomes. This body of work allows us to move from a broad perspective related to the impact of platforms on the economy to a more focused view of platform operations and the workers involved. A study on Uber (Burtch et al., 2018) shows that when the platform is set up in a city, it reduces the quality of entrepreneurial activity. This study shows how online platforms can affect economic outcomes in the form of work activities and allow a better allocation of labor capital. Gig economy platforms affect both the economy and also, in some cases, the labor market and can create new OLMs. Cook et al. (2021) demonstrate the existence of a gender-earnings gap among drivers on the Uber platform but explain this in terms of the different behaviours of men and women drivers. As on Uber, a gender wage gap exists on the OLMs. Gomez-Herrera et al. (2019) found a significant 4% gender wage gap among workers on the largest OLM in the EU, PeoplePerHour. We also document that US women workers on AMT report earning less than US men workers, without reporting fewer working hours.

One of the main results concerning the influence of labor markets on OLM can be found in Borchert et al. (2018). They make a key contribution by showing that local unemployment has an impact on OLMs, and, in particular, on the number of people who register on micro-job platforms. They find that as local unemployment increases, the number of registrations on micro-job platforms increases. They highlight that OLMs can substitute the offline labor market. Dube et al. (2020) provide evidence that the low labor supply elasticity of around

0.1% on OLMs is linked to monopsony power in the marketplace. We observe that an increase in the search for part-time jobs reduces the number of US workers on AMT. Thus, slight variations in the prices on these platforms for executing tasks are negligible.

Some studies focus on AMT, but most examine platforms more generally, e.g., Paolacci et al. (2010) and Difallah et al. (2018) discuss how the platforms function and the demographics of the workers. As mentioned previously, they find that AMT is a secondary revenue source for US workers and that the two largest populations represented on AMT are residents of the US and India. Some previous studies discuss gig workers. Graham et al. (2017) find that these platforms can improve workers' livelihoods by providing jobs for the long-term unemployed population or through the anonymity that the OLMs allow.⁶ There is a recent but growing strand of work on the impact of media, such as video games or television shows, on economics and social outcomes. In this paper, we construct indices to measure interest in a leisure activity to estimate how this affects presence on AMT. To obtain the most accurate indices possible, we draw on Baker et al. (2017), who use Google Trends to construct a Google Job Search Index. Kearney et al. (2015) also use Google Trends data in their studies that show how the MTV reality-TV show "16 and Pregnant" reduced teen births by up to 4.3%. La Ferrara et al. (2012) and Chong et al. (2009) also demonstrate an effect on social outcomes, showing in the first paper how "telenovelas" impact fertility and in the second how this affects the number of divorces in Brazil. Other papers, such as Gentzkow (2006) and Gentzkow and Shapiro (2008), show how television has replaced other media and changed voter behaviour. The impact of video games, and especially Pokémon Go, has been studied in the literature. Faccio et al. (2020) demonstrate that Pokémon Go has induced a significant increase in driving accidents due to people playing Pokémon Go while driving; the costs of driving accidents increased fivefold in the 148 days after the game was introduced. Pokémon Go also increased physical activity for all American players regardless of gender or age (Althoff et al., 2016). Sari (2019) also investigates how spatial disparities affect access to particular leisure facilities using Pokémon Go. Ward (2011) demonstrates how video games are associated with significant declines in crime and death rates and that other youth-related leisure activities (such as sports and movie viewing) generate smaller or no effects. Aguiar et al. (2021) calculate that innovations in recreational computing increased the marginal value of time for younger men, thereby reducing their market hours by about 2%. We document

⁶This provides an ambiguous benefit by hiding potential discrimination.

that the use of Pokémon Go influences the presence on OLM of US workers. When the usage in the game increase, fewer US workers are present on the platform because they spend less time on it, this effect is due to younger workers and is balanced between man and woman. This article argues that the Pokémon Go index, created using daily Google Trends, is a proxy for interest in the game or the number of people playing the game.

1.4 Data

To estimate the effect of the usage of Pokémon Go on the online workers, we use data on AMT workers, the usage of the mobile game and the search for part-time jobs. We construct a data set that includes information on workers on AMT at a daily level over a three-and-a-half-year period, a measure of the intensity of internet research on Google on the subject of “Pokémon Go”, also at a daily level, and data on the search for part-time jobs and AMT jobs.

First, the data provide information on the demographics of workers present on the platform each day between March 2015 and November 2019. The data were obtained via a survey conducted on AMT by Panos Ipeirotis⁷ (Difallah et al., 2018). The survey was run as a task available on AMT and comprised of questions to determine an online worker’s gender, age, household income and size, location (city and country), and marital status with a reward of 5 cents per question. This yielded 96 observations per day, with one observation characterising one AMT worker. Similar to other papers, such as Schilbach (2019), we use self-reported labor data, which permits us to have information on demographics and workers’ habits. The task was well paid compared to the average wage on the platform (Hara et al. (2018) found a median wage of 2USD per hour) and thus was attractive to all workers on the platform. We consider this demographic data to be reliable due to the quality of the work provided by AMT workers and because the literature provides evidence that this is consistent (Buhrmester et al., 2011; Peer, Vosgerau, et al., 2014). Ghose et al. (2012) used similar crowdsourcing techniques to design rankings for search results. The paper of Benson et al. (2020) highlights that employees also select employers through third-party reputation systems such as Turkopticon and that a good reputation attracts higher-quality workers, resulting in no loss in quality and faster completion of tasks. From the data, we observe that

⁷Source: Mturk-Tracker (last retrieved in July 2021).

around 73% of respondents live in the US and 18% in India.⁸ Previous results show that AMT is a primary source of revenue in India, while in the US, it is a secondary income source (Difallah et al., 2018). Therefore, AMT workers who are American or are located in the US should be more likely to engage in leisure activities. These leading presences are especially interesting because of the labor market and mobile network penetration differences.⁹ The respondents are, on average, 35 years old (median at 32) and are more evenly split between men and women (about 46% of men).

Based on these data, we construct the variables that give the daily proportion of workers according to location, gender, and age for 1,655 days. Table 1.1 describes all the variables used in our specifications. This data allows us to create more specific categories, such as the proportion of US women workers under 32, for example.

Table 1.1: Variables' description

Variable	Description
Dependent variables	
PropWorker	Proportion of US workers on AMT
PropWorker Man	Proportion of US man workers on AMT
PropWorker Woman	Proportion of US woman workers on AMT
PropWorker Under 32	Proportion of US workers aged 32 and under on AMT
PropWorker Over 32	Proportion of US workers older than 32 years old on AMT
Variables of interest	
PokeGo Trend	Index of search of "Pokémon Go" on Google in the US
PartTimeJob Trend	Index of search for "part-time job" on Google in the US
AMT Trend	Index of search for "Amazon Mechanical Turk" on Google in the US
Instrumental variables	
RarePokemon Trend	Index of search of "Mewtwo" on Google in the US
RarePokemon Trend NZ	Index of search of "Mewtwo" on Google in New Zealand

Table 1.2 presents summary statistics for the variables constructed based on the collected data.¹⁰ The observation level is per day. The variable *PropWorker* shows that, on average, there are almost 73% of US workers on AMT each day, the other 27% being located in different countries but mainly in India. US workers are more likely to be women (39.26%). The index for Pokémon Go, which measures usage in the US, ranges from 0 to 100, with an average of 2.53.

Figure 1.1 depicts the evolution of the proportion of US workers on the platform, the *PokeGo Trend* index, the index of search for a part-time job on Google and the index of

⁸Unlike other international workers, American and Indian workers can transfer their earnings to bank accounts from their Amazon Payments account. Others can only transfer their earnings to Amazon gift cards.

⁹81% in the US vs 49% in India at the end of 2015 - Source: OpenSignal.

¹⁰Table A.3 presents data from the survey for which the observation level is by workers.

search for AMT search on Google.¹¹

Table 1.2: Summary statistics

Variable	Mean	Std. Dev.	Min.	Max.
Dependent variables				
PropWorker	72.91	8.34	43.299	100
PropWorker Man	33.65	5.83	0	62.5
PropWorker Woman	39.26	7.25	6.25	87.5
PropWorker Under 32	35.249	7.062	0	67.010
PropWorker Over 32	37.659	6.68	12.5	100
Variables of interest				
PokeGo Trend	2.525	6.76	0	100
PartTimeJob Trend	59.62	13.89	20.09	100
AMT Trend	49.24	12.34	18.24	100
Instrumental variables				
RarePokemon Trend	8.23	8.29	1.8	100
RarePokemon Trend NZ	5.48	8.48	0	100
Observations	1,655			

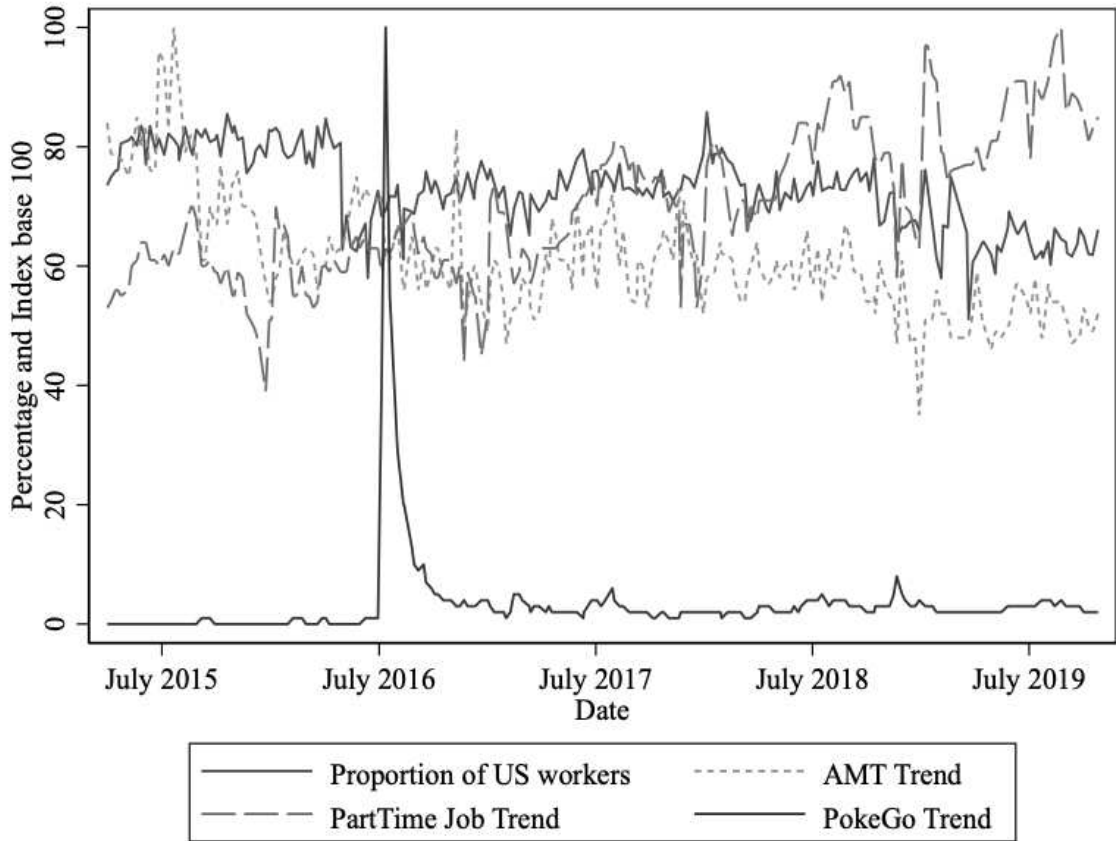


Figure 1.1: Proportion of US workers and the Google Trends for Pokémon Go and part-time job searches

Notes: This figure depicts the variation in the proportion of US workers in relation to the leisure variable *PokeGo Trend* and the job search variables *PartTimeJob Trend* and *AMT Trend*.

We chose *PokeGo Trend* because players and potential AMT workers have to perform a

¹¹Figure A.2 in the Appendix displays the evolution of the proportion of workers per country of residence.

trade-off between staying at home and working or going outside to play the game.

The use of Google Trends¹² in the literature is increasing; Choi et al. (2012) show how to use search engine data to forecast near-term values of economic indicators. Ma et al. (2021) and Baker et al. (2017) use it to construct their variables of interest. We construct our variable *PokeGo Trend* and our job search variables with Google Trends. This variable is motivated by the paper of Comino et al. (2019). They show that updating an application allows it to stay at the top of the charts by remaining among the top search results. To follow up or research an event or mobile application demonstrate the level of interest. Google Trends provides information on the weekly intensity of searches on the keyword “Pokémon Go”, and Google groups the searches that lead to a Pokémon Go result.¹³ To obtain the daily variability, we weighted this average weekly search intensity by day using a Python program. *PokeGo Trend* is the variable that measures the leisure activity, and ranges from 0 to 100. Table 2.3 and Figure 1.1 provide an overview of the leisure activity indices – one descriptive and one graphical.

1.5 Empirical Strategy

1.5.1 OLS Estimation

The analysis aims to determine how the proportion of workers, depending on their country of residence, age or gender, varies with the use of Pokémon Go. We control for the search for part-time jobs and the search for the AMT online labor platform. Since we want to measure the daily work-leisure trade-off, we use OLS regression:

$$\begin{aligned} PropWorker_t = & \beta_0 + \beta_1 PokenGo_t + \beta_2 AMT_t + \beta_3 PartTimeJob_t \\ & + \alpha_t + \epsilon_t \end{aligned} \tag{1.1}$$

We estimate the dependent variable *PropWorker*, which measures the daily proportion of workers on AMT according to their characteristics (gender and age) at time t . The independent variable of interest, *PokeGo*, is a leisure activity variable that varies at time t and measures Pokémon Go usage. We include Google searches for part-time jobs and for the

¹²Google Trends is a Google website that analyses the popularity of the top search queries in Google Search across various regions and languages.

¹³This includes spelling mistakes and related searches.

platform of micro-job AMT with variables *PartTimeJob Trend* and *AMT Trend*, which vary at time t . Finally, we control for time with α_t , a time fixed effect (by day, week and year), which controls for unobserved temporal trends or shocks (e.g., seasonality). ϵ_t denote the error term. Observations are at a daily level. As *PokeGo* should induce workers to leave the working platform, we expect β_1 to be negative. We expect the search for a part-time job to be a complement to AMT, so we expect β_2 to be positive in Equation (1.1).

1.5.2 IV Estimation

However, there might still be concerns about unobserved factors that may be correlated with both the popularity of the game Pokémon Go and the participation rate on the OLM. Perhaps the game is particularly appealing to individuals living in locations where participating in the OLM is much more difficult. If this is the case, the estimated relationship of interest in an OLS model, even with fixed effects, would have a bias towards finding a negative effect of the game on participating in the OLM.

We estimate a 2SLS IV estimation; we instrument the usage of Pokémon Go using the search for a “legendary” Pokémon on Google.¹⁴ We use this variable because the release of this Pokémon is quite an important event for the players’ community.¹⁵ This Pokémon is also linked to other events that do not depend on Pokémon Go, such as other games, TV shows or movies. This instrument is then correlated with our explanatory variable *PokeGo*, but not with our dependent variable *PropWorker*. We use the search for this rare Pokémon in the US and also in New Zealand (NZ), where the game was released on the same date, although events do not take place at the same time all over the world but at the same local time. For example, a rare Pokémon will be available for capture at 11 am local time in each country.

The IV Equation (1.2) takes the form of the OLS equation as represented in Equation (1.1) above, the predicted usage replaces the variable of interest:

¹⁴This legendary Pokémon is Mewtwo, a very rare Pokémon available only during events.

¹⁵In 2017 there was an exclusive event in Chicago, and the 20,000 places for the event, priced at 25 USD, were sold out in 15 minutes. This event was located in Chicago Grant Park, and the Pokémon was later released worldwide.

$$\begin{aligned}
PropWorker_t = \beta_0 + \beta_1 \widehat{PokeGo}_t + \beta_2 AMT_t + \beta_3 PartTimeJob_t \\
+ \alpha_t + \epsilon_t
\end{aligned}
\tag{1.2}$$

Where \widehat{PokeGo} is generated by the first stage regression represented by Equation (1.3) in the IV framework:

$$\begin{aligned}
PokeGo_t = \beta_0 + \beta_1 RarePokemon_t + \beta_2 AMT_t + \beta_3 PartTimeJob_t \\
+ \alpha_t + \epsilon_t
\end{aligned}
\tag{1.3}$$

Tests show that the instruments are correlated with the endogenous regressors and are not weak instruments.¹⁶ All the specifications do not use IV, as some do not require instruments.

1.6 Results

As mentioned in section 1.5, we run OLS and IV regressions. We regress the proportion of US workers depending on their characteristics on the leisure activity variable (*PokeGo*), controlling by time with a year, week and day fixed effects. We control for part-time jobs search and the search for AMT on Google.

1.6.1 Impact of Pokémon Go on US workers

Table 1.3 shows the regressions for the proportion of US workers regressed on the *PokeGoTrend* variable and controlling for search for part-time jobs. Column (1) of Table 1.3 reports the OLS specification. Columns (3) and (4) of Table 1.3 report the IV specification using 2SLS or GMM. Column (5) presents the estimates of the reduced form specification.

¹⁶Table A.5 in the Appendix.

Table 1.3: Main specifications

	OLS (1) PropWorker	First Stage (2) PokeGo Trend	2SLS (3) PropWorker	GMM (4) PropWorker	Reduced Form (5) PropWorker
PokeGo Trend	-0.082*** (0.022)		-0.134*** (0.032)	-0.130*** (0.032)	
AMT Trend	-0.009 (0.018)	0.006 (0.008)	-0.011 (0.018)	-0.011 (0.017)	-0.012 (0.018)
PartTimeJob Trend	0.153*** (0.034)	-0.025* (0.013)	0.147*** (0.030)	0.149*** (0.033)	0.149*** (0.033)
RarePokemon Trend NZ		0.030 (0.040)			-0.034 (0.032)
RarePokemon Trend		0.687*** (0.068)			-0.070*** (0.027)
Constant	86.550*** (2.120)	-6.091*** (1.013)	86.818*** (1.927)	86.733*** (2.067)	87.548*** (2.107)
Day of the week FE	✓	✓	✓	✓	✓
Week FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓
Observations	1,655	1,655	1,655	1,655	1,655
R-squared	0.432	0.804	0.431	0.431	0.435
Durbin			0.002		
Wu-Hausman			0.002		
C-stat				0.001	
First Stage			2,122.221	85.565	
Sargan			0.316		
Basman			0.326		
Hansen's J statistic				0.355	

Notes: OLS, 2SLS and GMM IV with time fixed effects (FE) estimations. *PropWorker* is the dependent variable in all columns, except column (2) which shows the output of the first-stage IV specification. Robust standard errors reported in parentheses.

Significance levels: * $p < .10$, ** $p < .05$, *** $p < .01$

Our main interest is the impact of Pokémon Go on the proportion of workers on AMT. The variable *PokeGo Trend* has a significant and negative impact on the proportion of US workers on the platform, as shown in column (3) of Table 1.3. When the usage of Pokémon Go increases by one, the proportion of US workers on AMT decreases by 0.134 (significant at 1% level). A one s.d increase of the *PokeGo Trend* variable decreases the proportion of US workers on AMT by up to 0.9 percentage point. As Difallah et al. (2018) state, there are between 2,000 and 5,000 workers on the platform at any given time, which is equivalent to having 10,000 to 25,000 full-time employees. This decrease of US workers by 0.9 percentage point due to the use of Pokémon Go can be rewritten as a decrease of 90 to 225 US workers on AMT. Then it exists as a trade-off between work and leisure on this online working platform.

Next, we investigate whether the results are consistent with the previous findings in the OLM literature. We focus on the effect of the job search variables on the proportion of US workers on AMT. One of the main findings is that searching for a part-time job increases the proportion of US workers on AMT, and this increase applies particularly to young US workers. In the period of observation, the relative search for part-time jobs has a standard

deviation of between 13.89 (Table 2.3). Column (3) in Table 1.3 shows the positive coefficient of 0.147 indicates that an increase in the search for a part-time job would raise the proportion of US workers on AMT by about 2 percentage points. This can be rewritten as an increase of between 200 and 500 workers on the platform.

Mechanism: Impact on Duration of Work

As the results show, an increase in the use of Pokémon Go leads to a decrease in the presence of US workers on the platform. A possible mechanism behind this result is that online workers spend less time on AMT. If so, we want to see if this impacts their earnings. In order to test these hypotheses, we use data from the same survey. Indeed, in June 2017, questions were added; one on the time spent on AMT per week, another on the amount earned on the platform per week. These are categorical variables we describe in Table A.2 in the Appendix. The level of observation for these models is at the worker level. We estimate here two ordered logit models (Table 1.4) on US-located online workers:

$$\begin{aligned}
 TimeSpent_{it} = & \beta_0 + \beta_1 PokeGo_t + \beta_2 AMT_t + \beta_3 PartTimeJob_t \\
 & + X_{it} + \alpha_t + \epsilon_{it}
 \end{aligned}
 \tag{1.4}$$

$$\begin{aligned}
 AmountEarned_{it} = & \beta_0 + \beta_1 PokeGo_t + \beta_2 AMT_t + \beta_3 PartTimeJob_t \\
 & + X_{it} + \alpha_t + \epsilon_{it}
 \end{aligned}
 \tag{1.5}$$

The dependent variables *TimeSpent* and *AmountEarned* vary per time t and individual i . We include a vector of variables X representing the characteristics of workers (age, gender). The variable *TimeSpent* varies from 1 to 7, corresponding to “less than 1 hour per week” to “more than 40 hours per week”. *AmountEarned* varies from 1 to 9, corresponding to “less than 1USD per week” to “more than 500USD per week”.¹⁷ We control for time with α_t time fixed effects (by day of the week, week and year) which control for unobserved temporal trends or shocks (e.g., seasonality). ϵ_{it} denote the error term. The point of this model is to estimate what makes an individual increase or decrease the time they spend on AMT. The variable is categorized from less to more time on the platform. As in the previous

¹⁷*TimeSpent* and *AmountEarned* are detailed in Table A.2 in the Appendix.

model, we expect a negative β_1 , as an increase in the usage of Pokémon Go should decrease the time spent on the platform as it reduces the participation in the OLM.

Table 1.4: Effect on the Time Spent and Amount Earned on AMT

	Time Spent			Amount Earned		
	(1) All	(2) Over 32	(3) Under 32	(4) All	(5) Over 32	(6) Under 32
PokeGo Trend	-0.024* (0.012)	-0.021 (0.017)	-0.029* (0.017)	-0.011 (0.012)	-0.016 (0.017)	-0.010 (0.017)
AMT Trend	-0.001 (0.001)	0.000 (0.001)	-0.001 (0.001)	0.000 (0.001)	0.001 (0.001)	-0.000 (0.001)
PartTimeJob Trend	-0.000 (0.001)	-0.000 (0.002)	-0.001 (0.002)	0.001 (0.001)	0.001 (0.002)	0.001 (0.002)
Man	0.019 (0.030)	-0.116** (0.045)	0.158*** (0.038)	0.218*** (0.031)	0.074 (0.048)	0.358*** (0.039)
Day of the Week FE	✓	✓	✓	✓	✓	✓
Week FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Observations	55751	29137	26614	55751	29137	26614
Pseudo R-squared	0.007	0.004	0.007	0.007	0.004	0.012

Notes: Ordered logit with time fixed effects (FE) estimations. *TimeSpent* is the dependent variable in columns (1) to (3). *AmountEarned* is the dependent variable in columns (4) to (6). Standard errors clustered by worker id are reported in parentheses. Significance levels: * $p < .10$, ** $p < .05$, *** $p < .01$

The results are reported in Table 1.4. This set of regressions demonstrates that an increase in the practice of Pokémon Go decreases the time spent on AMT, even if the period of observation of this variable is after the release period of the mobile game. Indeed, column (3) of Table 1.4 shows that for an increase of one in *PokeGo Trend*, we expect a 0.029 increase in the log odds of being in a lower level of *TimeSpent*. The results confirm the hypothesis on the mechanisms behind the decrease in the proportion of US workers on AMT.

If workers then spend less time on AMT when the usage of Pokémon Go increases, we can expect workers to earn less on the platform. The dependent variable *AmountEarned* allows us to measure this. We do not see an impact of *PokeGo Trend* on the amount of money they earned on AMT per week. This suggests that workers increase their productivity on AMT when the usage of Pokémon Go increases as they spend less time on the online working platform but still earn as much.

Age and Gender Differences

Table 1.5 presents the regression for the proportion of US workers according to their characteristics (age and gender), regressed on the usage of Pokémon Go, and controlling for the search for part-time jobs in the US. Columns (1) to (4) in Table 1.5 report the specifications for the sub-sample of US workers (man, woman, older than 32 years old, younger than 32

years old, and so on).

Table 1.5: Main specification with IV by demographic group

	2SLS (1)	2SLS (2)	2SLS (3)	2SLS (4)
PropWorker	Over 32	Under 32	Man Under 32	Woman Under 32
PokeGo Trend	0.001 (0.029)	-0.135*** (0.029)	-0.074*** (0.023)	-0.061*** (0.020)
AMT Trend	-0.002 (0.017)	-0.009 (0.016)	-0.015 (0.013)	0.006 (0.012)
PartTimeJob Trend	0.013 (0.028)	0.134*** (0.027)	0.085*** (0.022)	0.049** (0.020)
Constant	48.355*** (1.765)	38.463*** (1.742)	19.315*** (1.378)	19.149*** (1.241)
Day of the Week FE	✓	✓	✓	✓
Week FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Observations	1,655	1,655	1,655	1,655
R-squared	0.255	0.351	0.141	0.310
Durbin	0.534	0.004	0.038	0.087
Wu-Hausman	0.543	0.005	0.042	0.093
FirstStage	2122.221	2122.221	2122.221	2122.221
Sargan	0.064	0.003	0.000	0.903
Basmann	0.070	0.003	0.000	0.905

Notes: 2SLS IV with time fixed effects (FE) estimations. *PropWorker Over 32* is the dependent variable in column (1), *PropWorker Under 32* in column (2), *PropWorker Man Under 32* in column (3), and *PropWorker Woman Under 32* is the dependent variable in column (4). Robust standard errors are reported in parentheses. Significance levels: * $p < .10$, ** $p < .05$, *** $p < .01$

This decrease in the proportion of US workers on AMT seems to be driven by a generational effect. The negative and significant signs of the coefficients in columns (2) to (4) in Table 1.5 illustrate the generational effect driven by both men and women. It seems that the AMT population affected by Pokémon Go use are the workers aged 32 and younger. The effect of an increase in the usage of Pokémon Go results in a decrease in the proportion of young US workers by 0.135 percentage points. When considering splitting the population to obtain a more specific effect, we find that only workers younger than 32, regardless of their gender, are affected. Indeed, the increase in *PokeGo Trend* (and so the usage of Pokémon Go) reduces the proportions of men and women younger than 32 by almost the same level (-0.074 vs -0.061).

Table 1.5 shows a significant and positive effect of the search for part-time jobs on the proportion of young US workers. Both young men and women are affected, but the effect is slightly larger and more significant for young US men workers (0.085 vs 0.049). We can interpret these results as follows: the number of part-time jobs is limited, and therefore their supply is lower than the demand, while the search for this type of job is high, and the surplus is directed towards AMT, which is a substitute (even though it less well paid, it is

more flexible) for the part-time job market. This is consistent with the fact that changes in the part-time job market more impact young adults.

Urban vs Rural

Pokémon Go’s impact on online working activity by worker location is another area of interest because of the geographical heterogeneity of the game. In large cities, the game offers more activities to players and is more attractive, so the impact of Pokémon Go should be larger and more significant in those areas. Workers are more sensitive or not to in-game events if they live in a more urban or rural area, a more populated city or not, and so we started to study their locations. We observe that the distribution of AMT workers follows the distribution of the population in the US.¹⁸

To see if the impact of Pokémon Go changes based on the population of a city, we ran another set of regressions on two other sub-samples of the database. We split the population into larger and smaller city areas. We define large cities as American cities with a number of AMT workers above the median of our database and small cities as those with a number of workers below the median of our database. The results are reported in Table 1.6.

Table 1.6: Main specification by location

	OLS (1)	2SLS (2)	OLS (3)	2SLS (4)
PropWorker	Large City	Large City	Small City	Small City
PokeGo Trend	-0.065*** (0.021)	-0.101*** (0.028)	-0.016 (0.020)	-0.033 (0.026)
AMT Trend	0.003 (0.017)	0.002 (0.016)	-0.012 (0.016)	-0.012 (0.015)
PartTimeJob Trend	0.104*** (0.027)	0.100*** (0.027)	0.049* (0.026)	0.047* (0.025)
Constant	41.960*** (1.895)	42.143*** (1.728)	44.590*** (1.677)	44.675*** (1.611)
Day of the Week FE	✓	✓	✓	✓
Week FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Observations	1,655	1,655	1,655	1,655
R-squared	0.272	0.271	0.198	0.198
Durbin		0.016		0.230
Wu-Hausman		0.018		0.240
FirstStage		2122.221		2122.221
Sargan		0.751		0.124
Basman		0.756		0.131

Notes: OLS and 2SLS IV with time fixed effects (FE) estimations. *PropWorker Large City* is the dependent variable in columns (1) and (2), and *PropWorker Small City* is the dependent variable in columns (3) and (4). Robust standard errors are reported in parentheses.

Significance levels: * $p < .10$, ** $p < .05$, *** $p < .01$

Columns (2) and (4) in Table 1.6 report the 2SLS estimation results that show the differ-

¹⁸Figure A.1 in the Appendix shows graphical evidence of the distribution of AMT workers in the US.

ences in the impact of Pokémon Go between large and small cities. As we expect, Pokémon Go has a larger and significant effect (-0.101 significant at 1% level vs -0.033) because of the greater number of players and activities in larger cities, and because the game also needs more players in the surroundings to be more interesting. Workers in smaller cities then have fewer incentives to play the game and proceed to a trade-off between staying at home and working on the platform or going outdoors and playing Pokémon Go. The other difference is in the impact of the search for part-time jobs. This increases the proportion of workers located in denser areas much more than in small towns. This can be explained by the fact that there is a more saturated part-time labor market in large cities than in small ones.

1.6.2 Robustness Analysis

Released Period Excluded

In order to strengthen our results, we performed several robustness analyses.¹⁹ To check whether our result depended on the release period of Pokémon Go, which is when interest in the game is highest, we excluded it from the data. Thus, we reconstructed our variables starting, in this section, on 1 January 2017. The 100 based index created with Google Trends is then recalculated as the spike for Pokémon Go is excluded from this sub-sample. We see in Figure 1.2 and Table 1.7 that the average value for *PokeGo* is higher than in the analysis with the entire sample as the peak due to the release of the game is excluded, as it compresses the subsequent usage less.

¹⁹We also performed a placebo test that is described in the Appendix A.3.1

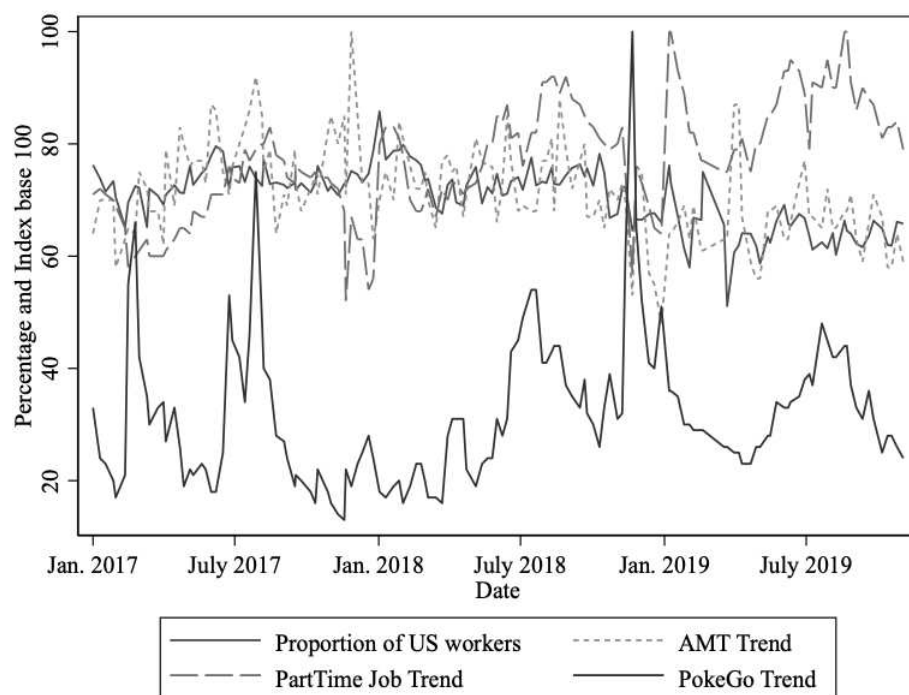


Figure 1.2: Proportion of US workers and the Google Trends for Pokémon Go and part-time job searches

Notes: This figure depicts the variations in the proportion of US workers in relation to the leisure variable *PokeGo Trend* and the job search variables *PartTimeJob Trend* and *AMT Trend*.

Table 1.7: Summary statistics

Variable	Mean	Std. Dev.	Min.	Max.
Dependent variables				
PropWorker	70.597	7.639	43.299	100
PropWorker Man	33.434	5.343	18.75	47.917
PropWorker Woman	37.163	6.96	6.25	75
PropWorker Under 32	33.106	5.833	0	52.632
PropWorker Over 32	37.492	7.071	12.5	100
Variables of interest				
PokeGo Trend	23.931	10.723	6.600	100
PartTimeJob Trend	66.195	12.814	22.4	100
AMT Trend	57.284	12.57	22.55	100
Instruments				
RarePokemon Trend	18.259	9.503	5.5	100
RarePokemon Trend NZ	8.854	11.484	0	100
Observations	1,011			

Table 1.8: Main specification with IV

	OLS (1) PropWorker	First Stage (2) PokeGo Trend	2SLS (3) PropWorker	GMM (4) PropWorker	Reduced Form (5) PropWorker
PokeGo Trend	-0.104*** (0.032)		-0.409*** (0.107)	-0.410*** (0.104)	
AMT Trend	0.018 (0.018)	-0.129*** (0.025)	-0.028 (0.026)	-0.028 (0.024)	0.024 (0.018)
PartTimeJob Trend	0.105*** (0.040)	0.027 (0.058)	0.111*** (0.037)	0.116*** (0.041)	0.099** (0.040)
RarePokemon Trend NZ		-0.037 (0.024)			0.001 (0.019)
RarePokemon Trend		0.266*** (0.043)			-0.105*** (0.028)
Constant	81.461*** (2.672)	24.849*** (3.993)	91.330*** (4.238)	91.131*** (4.136)	81.255*** (2.722)
Day of the week FE	✓	✓	✓	✓	✓
Week FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓
Observations	1,011	1,011	1,011	1,011	1,011
R-squared	0.443	0.495	0.346	0.345	0.442
Durbin			0.001		
Wu-Hausman			0.002		
C-stat				0.001	
First Stage			27.071	19.271	
Sargan			0.477		
Basman			0.492		
Hansen's J statistic				0.469	

Notes: OLS, 2SLS and GMM IV with time fixed effects (FE) estimations. *PropWorker* is the dependent variable in all columns, except (2), which shows the output of the first-stage IV specification. Robust standard errors are reported in parentheses.

Significance levels: * $p < .10$, ** $p < .05$, *** $p < .01$

Table 1.8 reports the results of our model, excluding the release period and the end of the year 2016. In this period, the *PokeGo Trend* variable has a standard deviation of 10.72. A one s.d. increase in the usage of Pokémon Go leads to a decrease in the proportion of American workers on the platform of 4.38 percentage points. It corresponds to a decrease of up to 1,095 American workers on AMT.

Threshold Effect

We use the subsample with the release period excluded. We constructed a *PokeGo Decile* variable that split our observations into ten deciles. The observations in the 1st decile, “Decile 1”, are those with the lowest Pokémon Go usage. Conversely, those in the 10th decile are those with the highest usage. The distribution of the *PokeGo Trend* by decile is detailed in Table 1.9. We estimate an OLS estimation, similar to our main specification, but we replace our *PokeGo Trend* variable with the *PokeGo Decile*:

$$\begin{aligned}
PropWorker_t = & \beta_0 + \beta_1 PokeGo_Decile_t + \beta_2 AMT_t + \beta_3 PartTimeJob_t \\
& + \alpha_t + \epsilon_t
\end{aligned}
\tag{1.6}$$

The dependent variable is *PropWorker*. As in the previous specification, the level of observation is at the day level. We control for seasonality with α_t time fixed effects and for part-time job and AMT searches on Google with variables *PartTimeJob Trend* and *AMT Trend*, which vary at time t . ϵ_t denote the error term. We expect a negative and increasing β_1 as the decile increases. We plot the coefficients given by Equation 1.6. Figure 1.3 displays the results; “Decile 1” is the baseline.

Table 1.9: Distribution of *PokeGo Trend* by decile

PokeGo Decile	Obs.	Mean	Median	Min.	Max.
Decile 1	102	10	11	6.6	13
Decile 2	101	14	14	13	15
Decile 3	101	16	16	15	17
Decile 4	101	18	18	17	20
Decile 5	103	21	21	20	22
Decile 6	100	24	23	22	25
Decile 7	100	26	26	25	28
Decile 8	102	30	30	28	31
Decile 9	102	34	34	31	37
Decile 10	99	47	43	37	100
Total	1,011	24	22	6.6	100




Figure 1.3: Threshold Effect of Pokémon Go

Notes: This figure depicts the variations of the effect of Pokémon Go, with the release period excluded. 95% confidence intervals are displayed.

The higher Pokémon Go usage, the stronger its effect. Indeed, the effect becomes statistically different from 0 at the 5% level from the ninth decile. When *PokeGo Trend* becomes greater than 31, the proportion of US workers on AMT decreases by 1.985 percentage points. It became significant at the 1% level in the tenth decile, when the relative interest in the game exceeds at least 37, with the proportion of US workers on AMT decreasing by 3.891 percentage points. This result highlights that the game affects online workers when its use is at its maximum.

1.7 Conclusion

This article explores whether and how online workers perform a trade-off between work and leisure. We analyse the impact of Pokémon Go and the part-time job search on work on AMT within an OLS and an IV model framework. With a set of ordered logit estimations, we explain the mechanism behind our results. It should be noted that we employ a new measurement for leisure activity usage with the help of Google Trends data. We also create an index that measures the search for a part-time job using Google's tool.

According to the evidence presented in the article, the usage of Pokémon Go has a negative

and significant impact on the presence of workers on AMT. On the contrary, the search for a part-time job increases the proportion of US workers on the platform. This conclusion is robust for various specifications, including IV estimations and a wide range of robustness tests. The mechanism analysis indicates that workers tend to reduce their time on the platform, but without any impact on their earnings, when the use of Pokémon Go increases.

Because of the COVID-19 pandemic, on-site work has become less evident for many jobs. The traditional labor market has had to adapt to remote working, and the role of OLMs has become increasingly important. The fact that workers can divide their time between work and leisure in a more flexible way without suffering a drop in income should concern policymakers. Policies that allow and facilitate this flexibility should be tested.

This article demonstrates how a leisure activity affects the working behaviour of online workers in two dimensions; the fact of being on the platform and the fact that the time spent on the platform is longer or shorter. In the study, in the conditions where workers can engage in this trade-off, the literature shows that welfare increases if they take this option. Further research is needed to investigate other OLMs with higher-skilled jobs, such as Freelancer, to see whether wealthier and more qualified online workers would respond to leisure incentives similarly to AMT workers.

Chapter 2

Bad Nudge, Kids and Voice Assistants: A Social Preferences lab-in-the-field Experiment

WITH F. LE GUEL AND S. PAJAK

Abstract

Connected devices using voice recognition as a form of input (Google Home, Amazon Echo, Apple Homepod) are increasingly popular. This mode of interaction introduces new possibilities to influence the user. How feasible is it for these devices to manipulate their audience, particularly children? This paper investigates the results of a lab-in-the-field experiment conducted in a French primary school in July 2019, where a smart speaker, a robot, and an adult were attempting to influence children in their choice of sharing marbles with other kids. We adapted a dictator game for the children audience and then estimated the impact of two different nudging strategies (Social Proximity and Peer-Effect) on the outcome of the dictator game. During the interaction with the children, the nudges were less effective when they were implemented by the adult as compared to the voice assistants, shedding light on the potential that these emerging devices have when it comes to manipulating their vulnerable audience.

2.1 Introduction

Connected devices with voice recognition (Google Home, Amazon Echo, Apple Homepod) allow for interaction with their users through voice recognition, natural language processing, and Artificial Intelligence (AI)-generated response. Over 36% of US households have a smart speaker in 2021. More than 52% have a smart home device, and over 58% have one or the other.¹ Voice interaction can be seen as more convenient and more natural than other types of input like a keyboard and can be done without interactions with an interface. However, this mode of interaction introduces new possibilities for the device to influence the user.

As the popularity of these devices grows, so does our exposure to voice interaction and familiarity with such interaction. Through smartphones, smart speakers, smart home devices, and cars, we have experienced a sharp increase in this type of interaction in the last few years. In this context, an increasing part of the population can and is exposed to these devices and this type of interaction. These new interactions can be voluntary or involuntary as these devices constantly listen to interact with users. More vulnerable agents can then be exposed, even unintentionally. How feasible is it for these devices to manipulate their audience? More precisely, are these devices more or less effective at manipulating children than an adult experimenter?

In order to bring answers to these questions, we conducted an experiment in July 2019 in a French primary school with children aged from 5 to 11 years old. In this experiment, we use a dictator game (DG) to test whether a manipulative strategy (a nudge) has different effects when a connected device (a robot or a smart speaker) conducts it than when an adult does. Each child is randomly assigned to one of three groups. In the three groups, the only difference is the interlocutor's identity, that is, an adult, a robot, or a smart speaker. The interlocutor will aim to make the children deviate the most from her initial choice, either to make her more or less altruistic. To do so, we repeat the game two times. As both outcomes can be valued in this context, making a child less altruistic will increase her earnings in this game, and more will increase others' earnings. The decision to make them more or less altruistic depends on their previous choice. The interlocutor will be consecutively using two different nudging strategies: a Social Proximity nudge and a Peer-Effect nudge. Both of these strategies are based on previous results found in the literature about the ef-

¹<https://voicebot.ai/2022/06/20/over-half-of-u-s-adults-have-smart-home-devices-nearly-30-use-voice-assistants-with-them-new-report/> Last retrieved in July 2022

fect of social proximity on altruism and norm implementation by Artificial Intelligence (AI) and robots and will exploit behavioral biases. We find that children can be influenced by robots and smart speakers to the same extent as by an adult, if not more. It does not matter which way the nudge goes. The smart speaker and the robot are still as effective as the adult.

Our contribution is twofold. First, we contribute to the literature about experiments in economics through experiments with sensible agents and using nudges in experiments. Second, we contribute to the literature about the use of AI and robot in experiments.

We designed our experiment to evaluate the manipulation potential of smart devices towards sensible agents. Then we chose to focus on the children population as their behaviors in games, especially in dictator games, have been widely studied. In their paper, Murnighan et al. (1998) show the differences between the behavior of adults, older children, and younger children in an ultimatum game, of which the dictator game is a derivative. They find that young children tend to make larger offers and accept smaller offers than older participants. Harbaugh, with co-authors, have primarily contributed to this literature and brought us essential information on how child behave in games in general and about their rationality and their difference in behavior with adult and the elderly who are sensible agents as children. Harbaugh and Krause (2000) study how altruism differs between child and adult and how repetition affects it. They find that altruism is the same, but repetition increases altruism for younger children and decreases for adults and older children. Harbaugh, Krause, and Berry (2001) test children's rationality; older children tend to be more rational, but younger children are more rational than expected. They show that older children are as rational as adults for simple economic decisions. Harbaugh, Krause, and Vesterlund (2002) show that children's choices are consistent with the under-weighting of low-probability events and the over-weighting of high-probability ones. This tendency diminishes with age; on average, adults appear to use objective probability when evaluating risky prospects. In their paper Sutter (2007) made children aged 4 to 9 play an ultimatum game. They find that children and teens propose similar allocations of outcomes than university students but tend to refuse much more unequal allocation, even if no better solution is available to the proposer. Benenson et al. (2007) demonstrated that older children and children from higher SES environments behaved more altruistically, although most children displayed altruistic

behavior even at the youngest age level. Bauer et al. (2014) show that the children of parents with low education are less altruistic, more selfish, and more likely to be weakly spiteful. Gummerum, Hanoch, et al. (2010) and Gummerum, Keller, et al. (2008) show that females tend to be more altruistic due to a higher preference for equal splits and that moral judgment influences altruism. Malti et al. (2009) study how sympathy, depending on children's moral motivation level, influences their altruism. Prosocial behavior increased with increasing sympathy, especially if children displayed low moral motivation. Kosse et al. (2020) show that a prosocial environment and training of the children increase their altruism towards others. In their literature review, Sutter et al. (2019) relate the results we observe in many occidental countries, that females and older children are more altruistic than males and younger children. In their paper, List et al. (2021) thoroughly reviewed the literature about games with children in experimental economics. Our setting follows the ground rule of the classical experiment with children. As in previous studies we also found that older children tend to be more altruistic, but we do not find any difference depending gender. List et al. (ibid.) recall most of the results found in the literature, but they identify a lack in the literature on how the design of games can induce change in children's behaviors.

The second stream of literature we contribute is the one on nudging, especially using AI and machines. One of the main contributors to the nudge literature is Nobel Prize Richard Thaler (Thaler and Sunstein, 2009). His work shows how design decision that still leaves individuals the choice of their decision can influence their economic decision. Thaler and Benartzi (2004) show that policymakers can induce them to make better saving decisions by using people's biases (inertia and status quo). However, Richard Thaler's work on nudging anticipates the possibility of misusing cognitive bias to alter choices in a purely manipulative way (Thaler, 2018a,b). For instance, excessive paperwork or the so-called 'dark patterns' in online interfaces. In their paper, Camerer et al. (1995) considered modifications to the design of the dictator and ultimatum games to change the outcome by explaining non-rational behavior. Sunstein (2020) shows that nudge can quickly become sludge as the design of a process can have high friction to disincentive consumers to adopt the behavior that will bring them the most benefit. In our experiment, we vary the social proximity of the child with the interlocutor in a dictator game. We inspired ourselves with a model of social distance

developed by Akerlof (1997), which helps understand social decisions. We use the fact that people are biased towards people closer to them socially. Charness et al. (2008) show that when the last name of their counterparts was known, dictators allocated a significantly larger portion of the pie. Bohnet et al. (2008) show that people are less willing to take a risk when another person rather than nature determines the outcome, less trust in humans than nature. Krupka et al. (2013) study the importance of social norm compliance. They found that subjects have a generally stable willingness to sacrifice money to take socially appropriate behaviors.

Hummel et al. (2019) recall some results found by the emerging literature on the use of digital nudge. Our work also contributes to the literature on nudging in a digital context or using robots and AI. In their paper, Peer, Egelman, et al. (2020) tested several nudges in the context of online passwords setting. They tailored the nudge based on the user and their decision-making characteristics to have a stronger effect and help the subject improve their password quality. Yeung (2017) shows that the implementation of designs that help consumers to have more control over their data is requested. Before testing how robots and AI can influence behavior, we must understand how humans and robots interact. Bruce et al. (2002) analyze human-robot interactions and how to improve them through gestures and expressions used by the robots. Roth et al. (2019) perform a Stackelberg Security Game using a robot. They vary humanoid robot expressive language (in the form of “encouraging” or “discouraging” verbal commentary). Their results show that a robot opponent that makes discouraging comments causes a human to play a game less rationally and to perceive the robot more negatively.

Machines powered by AI are increasingly present in our society. Rahwan et al. (2019) insist on the importance of studying machine behavior as machines powered by artificial intelligence increasingly mediate our social, cultural, economic, and political interactions. Understanding the behavior of artificial intelligence systems is essential to our ability to control their actions, reap their benefits and minimize their harms. If machines can collectively interact with one another and potentially lead to economic issues (such as online pricing collusion (Calvano et al., 2020)), constant progresses in the field of AI lead these machines to learn from human behavior, mimic them and even make autonomous decisions capable of influencing human decision making. Therefore, it becomes necessary to under-

stand the behavior of these 'AI agents'. Human-machine interactions can take several forms depending on whether the AI agent acts as an advisor, a partner, or a delegate to humans (Köbis et al., 2021). Kleinberg et al. (2018) show that machine learning assistant for judges can lead to better decision. It lead to crime reductions up to 24.7% with no change in jailing rates, or jailing rate reductions up to 41.9% with no increase in crime rates. In this context, where economic agents can rely on AIs to make or influence their decision, how AIs and machines influence humans must be studied. Luo et al. (2019) show that undisclosed chatbots are as effective as proficient workers and four times more effective than inexperienced workers in engendering customer purchases. However, revealing that the chatbot is not a human decrease the success of the chatbot. Unless it is disclosed late, the reduction is lower. Schanke et al. (2021) considers the impact of anthropomorphism on consumers' reaction to a chatbot conversation in the retail clothing industry. In their experiment, they implement varying degrees of social clues for the chatbot, namely humor, communication delay and social presence (greetings). They provide evidence that anthropomorphism is beneficial for transaction outcomes, but that it also leads to significant increases in price elasticity. In our paper we show that when a human, a robot, and an intelligent speaker adopt the same behavior to interact with a child and try to manipulate it, they succeed with similar efficiency. However, our results suggest that smart speakers and robots are more efficient than humans when they use a nudge based on the enunciation of a norm (Peer-Effect).

The remainder of this article is organized as follows. In Section 2.2, we present the experiment design. In Section 2.3, we describe the data collected and in Section 2.4, the models used in the empirical analysis. The results are presented in Section 2.5 and Section 2.6 concludes the article.

2.2 Experimental Design

2.2.1 The Dictator Game

The experiment was designed to measure the effect of nudging on children, depending on the agent which implements the nudging strategy. In order to be understandable to all children, we have chosen a dictator game (DG).

The strengths of the DG rest on its uniform procedure, simplicity, and wide application,

which permit systematic comparisons of altruistic behavior across diverse individuals and contexts. In the DG, one individual (the proposer) dictates how much the other individual (the recipient) will gain (for detailed descriptions, see Kagel et al. (2020)). Unlike other games, the recipient does not have the power to reject the offer. The key components of the DG consist of a one-shot allocation of a valued resource from one proposer to one recipient, both of whom are anonymous to each other and to others. For convenience, resources generally are divided into discrete parts, such as 10 units, which then can be divided between the two players. DGs vary in the degree to which the researcher is aware of the proposer’s allocation, the recipient, and whether a rationale is provided for sharing (e.g., Hill et al. (2004)). Most DGs have been played with university students who allocate, on average, 20–30% of their resources to the recipient (e.g., Forsythe et al. (1994)). Of all the economic games developed to date, only the DG is simple enough to represent a valid measure of altruistic behavior in young children.

2.2.2 Experiment Design

The experiment was conducted in a private Catholic elementary school near Paris in July 2019. The school has kindergarten and primary school classes, with children aged from 3 to 11 years old. As we presented the dictator game in each treatment only verbally, we needed to have children old enough to understand the process without the need to use gestures. We had children from all classes whose parents had consented to their participation in the experiment. However, we selected only children from the last kindergarten class and all primary school classes. The ninety-one children were therefore aged between 5 and 11 years. The five classrooms range from 1st grade, with children age 5, to 5th grade (age 11). Children’s families were asked before the experiment to consent to their child participating in the experiment and being filmed. Each child participating in the experiment received 10 marbles no matter their decision during the experiment in order to not create an injustice feeling. The experiment was approved by the research ethics committee (CER) of the University Paris-Saclay.

Because the study is built on children from primary school, the population of participants was restricted to 120 potential candidates. More than half of the parents’ children accepted the immediate participation upon receiving an invitation, but we had to exclude some in-

dividuals, leaving a total of 91 participants (31 participated in the Human treatment, 29 in the Robot treatment and 31 in the Smart Speaker treatment).

To causally identify the effect of interacting with an adult as opposed to a voice assistant, we implement three independently randomized treatments by varying the nature of the agent with whom the children interact: an adult, a smart speaker (Google Home), or a robot (Pepper)². The school made three rooms available to us, and we set up each room with one of the treatments (Adult, Smart Speaker, Robot). According to our randomization scheme, the teacher sends each child to either one of the three rooms. The child performs an interaction of about 10 minutes and returns to her classroom, and the next child is sent to one of the rooms based on the randomization we provided to the teacher.

Compared to the standard Dictator game, our implementation is adjusted to be understood and performed by children. The child is presented with three bowls, two are empty, and one contains a set of ten marbles. The child is told to split the marbles between those she wants to keep for herself (to be put into bowl 1) and those she is willing to share with other kids from the school (to be put into bowl 2). She is asked to move each marble, one at a time, into one of the two empty bowls.

After the child has made her choice, the split is said aloud. Immediately after, she is given the possibility to change her split choice, and either two strategies of nudging is implemented in random order:

Peer-effect: “You know, other kids choose $(2/8)$ marbles. If you want, you can modify your choice.”

Social Proximity: “You know, I would choose $(2/8)$ marbles. If you want, you can modify your choice.”

In the first instance, we aimed to divert the child as far as possible from her initial choice. So we set up this decision rule, if the child gave 3 marbles or less, we told her to give 8. Conversely, if she gave 4 marbles or more, we told her to be less altruistic and to give only 2. This decision rule was based on the average allocation of resources done by children to the recipient, around 30% of the initial endowment (List et al., 2021). Then we expected to have the same proportion of more and less altruistic nudges.

²Figure B.1 in Appendix shows the appearance of our Robot Pepper that children played with in the Robot treatment and the Smart Speaker.

We chose these two strategies (Peer-effect and Social Proximity) because they rely on different biases, and subjects’ perceived humanity of the interlocutor can have opposite effects (H2.1 & H2.2). Based on the literature, we formulated the following hypotheses:

- H1: A child playing a dictator game with a smart speaker or a robot shares as many marbles with the class as when playing with an adult.
- H2.1: Humans are more persuasive than smart speakers or robots in using the Social Proximity nudging strategy (Kosse et al., 2020).
- H2.2: Robots and smart speakers are more persuasive with the Peer-Effect nudging strategy (Shirado et al., 2017; Vollmer et al., 2018).

As machines and IA are often used and showed their efficiency to help make decisions (Kleinberg et al., 2018; Rahwan et al., 2019), they should be more convincing to state a norm.

A summary presentation of the course of the experiment is given by Table 2.1.

Table 2.1: Timeline of the Experiment

Activity	Description
1. Attribution to a treatment	The child is randomly assigned to a group.
2. Presentation of the game	In each group, the interlocutor presents the game following the same speech.
3. Comprehension task	The interlocutor asks the child to point to the bowl in which, if she puts marbles in, they will be for her.
4. 1st Split of the endowment	The child splits the marbles in her bowl and the other bowl.
5. Presentation of the outcome and First nudge attempt	The interlocutor describes how many marbles were given, nudges the child with the nudge randomly selected and asks the children if she wants to change her allocation of marbles.
6. 2nd Split of the endowment	The child decides to change or not her allocation of marbles between her bowl and the other bowl.
7. Presentation of the second outcome and second nudge attempt	The interlocutor describes how many marbles were given, nudges the child with the nudge randomly selected and asks the children if she wants to change her allocation of marbles.
8. 3rd Split of the endowment	The child decides to change or not her allocation of marbles between her bowl and the other bowl.
9. End of the game	The interlocutor thanks the child and she go back to her classroom.
10. Debrief	Once all children went through the experiment, we present them the robot in a workshop, and a bag of marbles was given to every child.

The payoff for the game is marble in this context because it is not possible and not ethical to pay children with money. Also, to make our experiment more ethical, every child who participated was given the same amount of marbles, no matter their decision in the game. They were all given their gain after every subject went through the experimentation.

2.2.3 Experimental Procedures (Experiment Setting)

The school provided us with three classrooms. Each one was similar to the others and was in a different part of the school. Children were not able to see each other. The experiment was carried out for one and a half days. We proceeded to the experiment during class hours, so the children couldn't communicate about the game they had just played and with which interlocutor. We managed to get all the participants of a class through before any breaks so that the children could not give each other any information before doing the game. Figure 2.1 shows some examples of how the experiment works in each treatment. In each room, two adults were present. One interacted with the children when she arrived in the room, made her seat, and then set the recording. The other adult was either in charge of the interlocutor (Robot and Smart Speaker) or the interlocutor (Human). In the treatment Robot and Smart Speaker, the other adult that welcomed the child also monitored the child's allocation of marbles to communicate them to the adult in charge of the Robot or the Smart Speaker. The detailed speech of the interlocutor is available in Appendix B.4.



(a) Adult Treatment



(b) Smart Speaker Treatment



(c) Robot Treatment

Figure 2.1: Proceeding of the experiment by Treatment

2.3 Data

In this section, we perform the first data analysis and exploration. Table 2.2 presents the variables that were created using the data collected during the experiment. We have basic information about the subjects, their age, gender, in which treatment group they were allocated to and in which order they faced the nudging strategies.

Table 2.2: Variables description

Variable	Description
Age	Age of the subject.
Male	Dummy variable 0= female, 1= male.
Treatment	Categorical variable that takes 3 different values, 0= Human, 1 = Robot, 2= Smart Speaker.
Nudge Order	Dummy variable 0= Peer-effect then Social Proximity, 1= Social Proximity then Peer-effect.
More Altruistic 1	Dummy variable equals 1 if we tried to make the child more altruistic at 1st nudge.
More Altruistic 2	Dummy variable equals 1 if we tried to make the child more altruistic at 2nd nudge.
Given	Number of marbles the subject gave at the first stage.
Given after first nudge	Number of marbles the subject gave after the first nudging strategy.
Given after second nudge	Number of marbles the subject gave after the second nudging strategy.
No Change 1	Dummy variable equals 1 the child didn't change her allocation after the 1st nudge.
No Change 2	Dummy variable equals 1 the child didn't change her allocation after the 2nd nudge.

The socio-demographics data (age and gender) permit us to randomize our subjects into three homogeneous groups. Each group was assigned a treatment (Human, Robot, Smart Speaker), and each subject faced two nudging strategies in a random order (peer-effect, then social proximity, or the contrary). Then at each stage of the experiment, we measure the number of marbles the child gives.

The final sample is composed of 91 children. All participate in the first stage. In the Smart Speaker treatment, we lose one observation at the second stage and another one at the third stage. One because of misunderstanding of the instruction, and another due to sickness. Summary statistics that include how altruistic the children were at each stage of the game are detailed in Table 2.3. The same summary statistics displayed by treatment are available in appendix B.2. Figures 2.2 and 2.3 display graphical evidence of the distribution of marbles operated by children for each stage and by treatment for the first stage.

Table 2.3: Summary statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
Age	8.67	1.687	6	11	91
Male	0.407	-	0	1	91
Treatment	1	0.830	0	2	91
Nudge Order	0.433	-	0	1	90
More Altruist 1	0.211	-	0	1	90
More Altruist 2	0.326	-	0	1	89
Given	4.582	2.186	0	10	91
Given after first nudge	4.7	2.329	0	10	90
Given after second nudge	4.427	2.388	0	10	89
No Change 1	0.495	-	0	1	90
No Change 2	0.516	-	0	1	89

Figure 2.2 shows that the equal split was the children’s preferred allocation of marbles at the first stage. Even after the first and the second nudge, it is still the most common strategy children adopt. As we try to nudge them to be less altruistic, we observe an increasing number of children giving two marbles and an increasing number of children choosing to give eight marbles as we try to make them more altruistic.

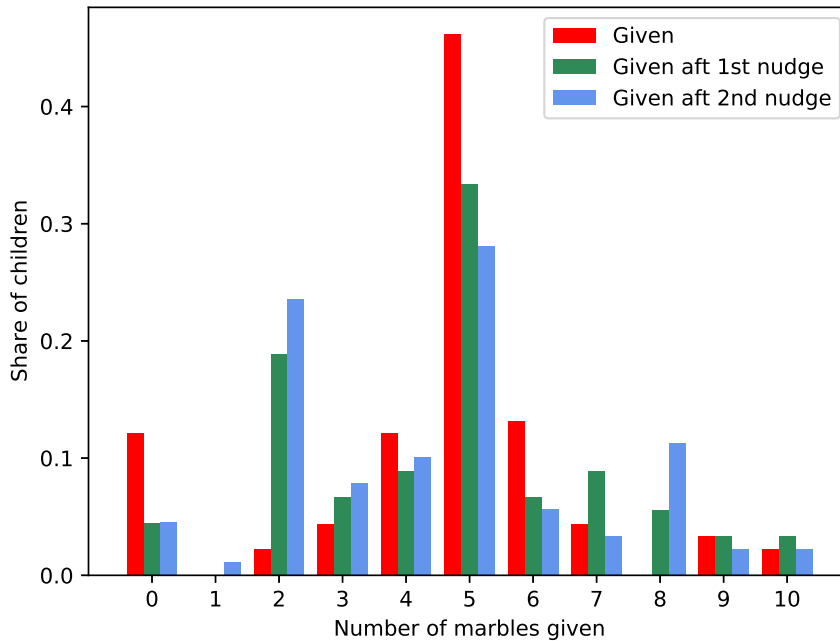


Figure 2.2: Distribution of marbles given at each stage

Figure 2.3 displays the difference in the first split of marbles done by children depending on the treatment. Figure B.2 and B.3 in Appendix display the distribution by treatment of marbles after the first and second nudge. When they face the adult, children tend to perform the equal split more often. However, as we test if the difference was significant using the

ANOVA test, we do not find statistically significant differences.

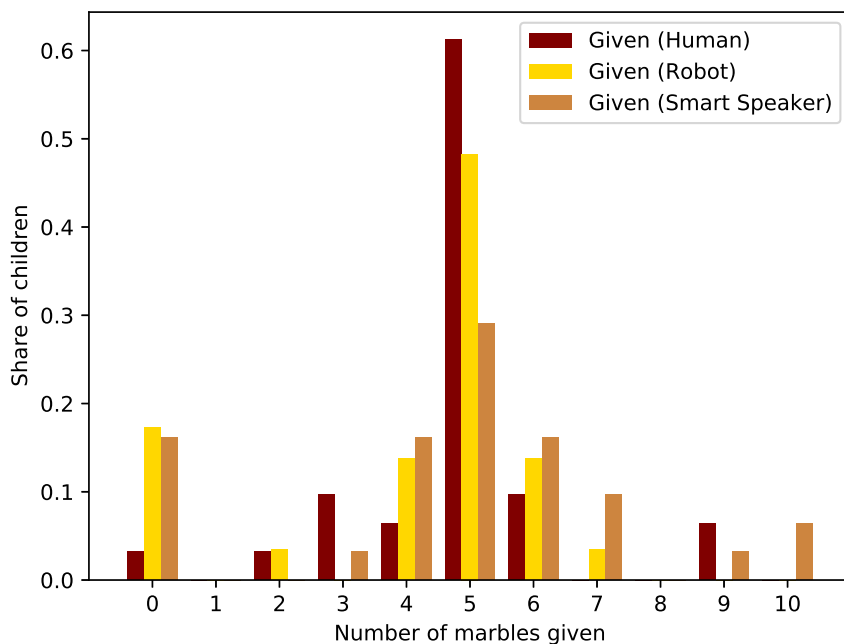


Figure 2.3: Distribution of marbles given at the first stage by treatment

2.3.1 Post experimentation verification

Table 2.4: Ttests

	Non-Robot vs. Robot		Diff.
	Mean non-robots	Mean robot	
Age	8.64	8.72	-0.08
Male	0.4	0.41	-0.01
Nudge Order	0.37	0.55	-0.18
	Non-Human vs. Human		Diff.
	Mean non-human	Mean human	
Age	8.63	8.74	-0.11
Male	0.43	0.35	0.08
Nudge Order	0.424	0.452	-0.028
	Non-Smart Speaker vs. Smart Speaker		Diff.
	Mean non-Smart Speaker	Mean Smart Speaker	
Age	8.73	8.54	0.19
Male	0.38	0.45	-0.07
Nudge Order	0.50	0.30	0.2*

Significance levels: * $p < .10$, ** $p < .05$, *** $p < .01$

With the socio-demographic variables we had prior to the experiment (gender, age), we pre-randomized our sample to have 3 homogeneous sub-sampled in terms of gender, age, and nudge order. After the experiment, we observed that our groups were not significantly dif-

ferent except for the Smart Speaker group that faced more frequently the nudging strategies in the Peer-effect than Social Proximity order.

2.4 Models

We want to estimate the altruism of the child at the three stages of the game, initial decision, after the first attempt of nudging, and after the second. In all specifications, we use OLS estimations.

$$Given_i = \beta_0 + \beta_1 Treatment_i + X_i + \epsilon_i \quad (2.1)$$

Firstly, we regress the number of marbles the child gives (*Given*) on the treatment if the interlocutor is a Human, a Robot, or a Smart Speaker (*Treatment*). We include a vector of variables X representing the characteristics of workers (age, gender). Each variable vary per individual i . ϵ_i denote the error term. *Treatment* takes values 0 to 2 for Human, Robot, and Smart Speaker treatment. This model permits us to test hypothesis H1, which states that the interlocutor should not affect how the children split her first endowment as the game is presented using the same speech in all three treatments. Then in equation 2.1, β_1 should not be significant as we made the hypothesis that in the first stage of the game, the interlocutor who presents to the child the game should not influence the number of marbles given by the child.

$$\begin{aligned} Given_aft_nudge_{is} = & \beta_0 + \beta_1 Given_{is} + \beta_2 Treatment_{is} + \beta_3 MoreAltruistic_{is} \\ & + \beta_4 Treatment_{is} * MoreAltruistic_{is} + \epsilon_{is} \end{aligned} \quad (2.2)$$

Secondly, we regress the number of marbles the child gives (*Given_aft_nudge*) on the treatment if the interlocutor is a Human, a Robot, or a Smart Speaker (*Treatment*) and on the direction of the nudge with *MoreAltruistic* that takes values 0 or 1. An interaction term between *Treatment* and *MoreAltruistic* will measure how the effect of the direction of nudge is different depending on the interlocutor. Each variable vary per individual i and stage s (2,3). ϵ_{is} denote the error term. Equation 2.2 represents the estimations done at

each stage s , i.e., after the first and the second nudging strategy.

$$\begin{aligned}
Given_aft_nudge_{is} = & \beta_0 + \beta_1 Given_{is} + \beta_2 Treatment_{is} + \beta_3 MoreAltruistic_{is} \\
& + \beta_4 NudgeOrder_{is} + \beta_5 Treatment_{is} * MoreAltruistic_{is} * NudgeOrder_{is} \\
& + \epsilon_{is}
\end{aligned}
\tag{2.3}$$

Finally, we regress the number of marbles the child gives (*Given_aft_nudge*) on the treatment if the interlocutor is a Human, a Robot, or a Smart Speaker (*Treatment*), on the direction of the nudge (*MoreAltruistic*) and the type of nudge with *NudgeOrder* that takes values 0 or 1 (0: *Peer-effect* first, 1: *Social Proximity* first). Each variable vary per individual i and stage s (2,3). Equation 2.3 represents the estimations done at each stage s , i.e., after the first and the second nudging strategy. Equation 2.3 implementation of an interaction term between *Treatment*, *MoreAltruistic* and *NudgeOrder* measure how the effect of the direction of nudge associated with the type of nudge is different depending on the interlocutor. These estimations will test hypotheses H2.1 and H2.2.

2.5 Results

2.5.1 Mean differences

In order to evaluate if we succeeded in making the children change their allocation of marbles, we first evaluated the significance level of the difference in the means. Tables 2.5 & 2.6 display the results of the effect of the direction of the nudge at the two stages of the game. No matter the treatment, we succeed in making the child more or less altruistic depending on the nudge. The results of this naive estimation suggest that we made them change their choice more efficiently when we tried to make them more altruistic.

Table 2.5: Effect of the direction of first nudging

	Given	Given aft 1st nudge	Diff.	N
Less Altruistic	5.28	4.71	-0.56***	71
More Altruistic	2.21	4.63	2.42***	19

Significance levels: * $p < .10$, ** $p < .05$, *** $p < .01$

Table 2.6: Effect of the direction of second nudging

	Given aft 1st nudge	Given aft 2nd nudge	Diff.	N
Less Altruistic	5.5	4.23	-1.27***	60
More Altruistic	3.14	4.83	1.69***	29

Significance levels: * $p < .10$, ** $p < .05$, *** $p < .01$

2.5.2 Econometric Results

2.5.2.1 Direction of the nudge

We estimate the altruism of the children at each stage of the game. We first estimate the baseline number of marbles depending on the faced interviewer. The results detailed in Table 2.7 are the estimates associated to the equation 2.1. We observe in column (1) that neither treatment (Robot or Smart Speaker) significantly increased or decreased the number of marbles given by the subject. This finding is consistent with our first hypothesis (H1) that the interlocutor should not impact the first stage decision. The low R^2 statistic shows that the treatment and the interlocutor do not explain the first outcome. We observe a positive effect for Age which is consistent with previous findings in the literature. Older children tend to be more altruistic. For the second and third stages, the interlocutor does not influence the children's altruism. Only the previous allocation affects how many marbles the child gives, which illustrates the path dependence of choices.

Table 2.7: Control estimation

	(1) Given	(2) Given after first nudge	(3) Given after second nudge
Robot	-0.710 (0.470)	-0.313 (0.506)	-0.783 (0.549)
Smart Speaker	0.018 (0.580)	0.510 (0.512)	-0.585 (0.538)
Male	-0.352 (0.481)	-0.012 (0.436)	-0.560 (0.467)
Age	0.249* (0.140)	0.333** (0.128)	0.036 (0.144)
Given		0.434*** (0.151)	0.140 (0.163)
Given after first nudge			0.349** (0.140)
Constant	2.791** (1.313)	-0.267 (1.061)	2.489* (1.445)
Observations	91	90	89
R-squared	0.067	0.290	0.229

Notes: OLS estimations. Robust standard errors are reported in parentheses. Significance levels: * $p < .10$, ** $p < .05$, *** $p < .01$

Table 2.8 displays the results for the first and second stages of the game, the initial share of marbles, and the one after the first attempt of nudging. Column (1) shows that the interlocutor does not have any effect on the children’s altruism at the first stage, as in Table 2.7. We test the nudge’s effect, depending on whether it was used to increase or decrease the children’s altruism. In column (2), we observe that we succeeded when we tried to make the children more altruistic, no matter the treatment. It increased the number of marbles the children gave by 2.155 on average. In columns (3) and (4), we interacted with the treatment with the variable *MoreAltruist*. Then the only difference between these columns is the reference group. In column (3), the reference group is the child that was with the Human, and we try to make them less altruistic. Then the only significant difference is with the children who played the game with the Smart Speaker and were nudged to be more altruistic. In Column (4), as we control that we ask the child to be more altruistic, we find no difference between the groups, whether we ask them to be more or less altruistic.

Table 2.8: Main estimation, initial split and after first nudge

	(1) Given	(2) Given after first nudge	(3) Given after first nudge	(4) Given after first nudge
Robot	-0.735 (0.489)	-0.195 (0.518)		
Smart Speaker	-0.065 (0.566)	0.195 (0.498)		
Given		0.738*** (0.112)	0.743*** (0.119)	0.743*** (0.119)
More Altruistic		2.155*** (0.587)		1.453 (1.276)
Human × More Altruistic			1.453 (1.276)	
Robot × Less Altruistic			-0.388 (0.475)	-0.388 (0.475)
Robot × More Altruistic			2.173 (1.658)	0.720 (1.936)
Smart Speaker × Less Altruistic			0.023 (0.506)	0.023 (0.506)
Smart Speaker × More Altruistic			2.394*** (0.778)	0.941 (1.463)
Constant	4.839*** (0.297)	0.826 (0.679)	0.912 (0.758)	0.912 (0.758)
Observations	91	90	90	90
R-squared	0.023	0.324	0.330	0.330

Notes: OLS estimations. Robust standard errors are reported in parentheses. Significance levels: * $p < .10$, ** $p < .05$, *** $p < .01$

Table 2.9 displays the results of the estimations of the number of given marbles at the third stage of the game, the share of marbles after the second nudging attempt. Column (1)

tests the effect of the direction of the nudge, controlling for the treatment and the previous allocation of marbles. As in the previous game stage, we find a positive and significant effect on altruism. After the first nudge, the effect of the treatment is insignificant. Only the direction of the nudge is. When we wanted to make them more altruistic, children increased, on average, their allocation of marbles to another child by 2.195. In columns (2) and (3), we interact the treatment with the direction of the nudge. The only difference between those two estimations is the reference group and how we read the results.

Table 2.9: Main estimation, after second nudge

	(1) Given after second nudge	(2) Given after second nudge	(3) Given after second nudge
Given	0.200 (0.151)	0.224 (0.145)	0.224 (0.145)
Given after first nudge	0.559*** (0.140)	0.530*** (0.145)	0.530*** (0.145)
Robot	-0.687 (0.522)		
Smart Speaker	-0.747 (0.492)		
More Altruistic	2.195*** (0.493)		1.098 (0.876)
Human \times More Altruistic		1.098 (0.876)	
Robot \times Less Altruistic		-1.259*** (0.477)	-1.259*** (0.477)
Robot \times More Altruistic		1.567 (1.059)	0.469 (1.235)
Smart Speaker \times Less Altruistic		-1.218** (0.548)	-1.218** (0.548)
Smart Speaker \times More Altruistic		1.404** (0.538)	0.306 (0.968)
Constant	0.610 (0.718)	0.982 (0.749)	0.982 (0.749)
Observations	89	89	89
R-squared	0.357	0.380	0.380

Notes: OLS estimations. Robust standard errors are reported in parentheses.
Significance levels: * $p < .10$, ** $p < .05$, *** $p < .01$

2.5.2.2 Social Proximity and Peer-effect

Even with a low number of observations, we estimate the effect of the type of nudge (Social Proximity or Peer-Effect) when we interact it with the treatment and the direction of the nudging strategy (More or Less Altruistic), we can perform some estimations. Table 2.10 and 2.11 report the estimations of the model written in equation 2.3. Table 2.10 displays the estimation of the number of marbles given after the first nudging strategy. The reference

group is the one that faced the human, that asked the children to be less altruistic using the Peer-Effect strategy. Compared to this group, the only statistically significant difference is among children that were asked to be more altruistic by a Smart Speaker, using both nudging strategies and by a robot using the Peer-Effect strategy. The larger effect of the smart speaker treatment is when it uses the Peer-Effect strategy. Even if it is not significant at a 10% level, when the human tells the child to be more altruistic, the effect is much larger when it is coupled with the Social Proximity nudge. These results are in line with our previously formulated hypotheses H2.1 and H2.2.

Table 2.10: Estimation with interactions after first nudge

	(1) Given after first nudge
Given	0.828*** (0.117)
Robot	-0.448 (0.683)
Smart Speaker	-0.145 (0.710)
Human \times Less Altruistic \times SP then PE	-0.562 (0.763)
Human \times More Altruistic \times PE then SP	0.295 (1.349)
Human \times More Altruistic \times SP then PE	3.214 (1.963)
Robot \times Less Altruistic \times SP then PE	-0.291 (0.568)
Robot \times More Altruistic \times PE then SP	4.233** (2.105)
Robot \times More Altruistic \times SP then PE	-0.096 (0.664)
Smart Speaker \times Less Altruistic \times SP then PE	-0.304 (0.670)
Smart Speaker \times More Altruistic \times PE then SP	2.802** (1.107)
Smart Speaker \times More Altruistic \times SP then PE	1.302** (0.563)
Constant	0.715 (0.864)
Observations	90
R-squared	0.413
Adj R-squared	0.321

Notes: OLS estimations. Robust standard errors are reported in parentheses. Significance levels: * $p < .10$, ** $p < .05$, *** $p < .01$

Table 2.11 shows the results from the third stage of the game. The fact that smart speakers and robots are more persuasive with the social proximity nudge at this stage can also be because the children have faced the Peer-Effect before, and this result is only due to repetition. A refined design with the implementation of a neutral nudge is needed to

evaluate which effect dominates.

Table 2.11: Estimation with interactions after second nudge

	(1) Given after second nudge
Given	0.296** (0.148)
Given after first nudge	0.531*** (0.134)
Robot	-0.866 (0.691)
Smart Speaker	-0.748 (0.714)
SameLevelNudge=1	0.303 (0.680)
Human \times Less Altruistic \times SP then PE	0.425 (0.647)
Human \times More Altruistic \times PE then SP	1.819 (1.270)
Human \times More Altruistic \times SP then PE	1.117 (1.151)
Robot \times Less Altruistic \times SP then PE	-0.304 (0.737)
Robot \times More Altruistic \times PE then SP	4.601** (1.992)
Robot \times More Altruistic \times SP then PE	2.267* (1.259)
Smart Speaker \times Less Altruistic \times SP then PE	-0.889 (0.928)
Smart Speaker \times More Altruistic \times PE then SP	2.829*** (0.817)
Smart Speaker \times More Altruistic \times SP then PE	1.455* (0.866)
Constant	0.116 (1.015)
Observations	89
R-squared	0.422
Adj R-squared	0.313

Notes: OLS estimations. Robust standard errors are reported in parentheses. Significance levels: * $p < .10$, ** $p < .05$, *** $p < .01$

2.6 Conclusion

The pandemic did not permit us to reproduce the experiment and increase our sample size. However, with our results and experience in this experiment, we can identify issues and deliver a refined design, allowing us to come up with stronger evidence and a more solid external validity. In this section, we will detail those refinements.

As with other experiments using faced interaction in a controlled environment, we lack a large number of observations. That is why even when we designed our experiment, we were planning to reproduce it. The implementation of the experiment has, however, brought to

light problems and possibilities for improvement. This will allow for better external validity.

In order to test the effect of a high or low nudge, we need a better control group, for which we do not intervene. We ask if the child wants to change her allocation of marbles. Also, if we have a sufficiently large number of participants, we can choose the low, high or neutral randomly and do not rely on our last decision rule. Our decision rule was implemented to make the children deviate the most from her previous choice, either to make her more or less altruistic. It will create by treatment (Human, Robot, and Smart Speaker) three subgroups, in which children will be randomly assigned, creating a 3*3 design. The fact that a participant is used to this kind of technology can influence how she will interact with the smart speaker or the robot. Then the implementation of a post(or pre)-experiment survey to collect relevant information such as the presence of a smart speaker or connected device at home, the number of electronic devices, and habituation to technology will be helpful. Also, other studies with children show that their school performance can be correlated to some behaviors.

How can smart devices like robots and smart speakers affect children's altruism? Common intuition suggests that people should nudge children more efficiently. This experiment investigates such effects in a setting where children interact with a human, a robot, or a smart speaker in the context of a dictator game. The interlocutor tries to nudge the children, either to make them more or less altruistic, using nudging strategies that exploit behavioral biases. Our results show that Robots and Smart Speakers are at least as efficient as an adult in influencing the altruism of a sensible agent, a child. As detailed above, further works and replications are needed to confirm our results' channel and mechanism.

Chapter 3

Online Popularity, Fake Followers and Soccer Players' Value

WITH N. SOULIÉ

Abstract

Talent and popularity are valuable assets for celebrities. While talent improvement involves generally efforts, online popularity can be easily manipulated due to firms selling social media fake followers. This paper investigates the economic motives to manipulate social media popularity through fake follower acquisition by professional athletes. Focusing on soccer, we create a unique dataset including 1,075 international players and take advantage of Twitter's suspicious account removal held in July 2018 to proxy fake followers. Empirical explorations provide evidence of consistent patterns of soccer players associated with fake followers. In their cases, Twitter account creation is quickly followed by a transfer into a new club and predominantly occurred before 2015, while Twitter was more popular than Instagram. Results show that fake followers impact significantly players' value, i.e., transfer fees, only if the transfer occurs within 1 to 6 months after the Twitter account creation. These fake followers are associated with an average rise of 6% of transfer fees (\approx €650,000). This result is consistent with the acquisition of fake followers to boost online popularity of a recently created Twitter account and monetize it during transfer bargaining.

3.1 Introduction

An investigation released in 2018 by *The New York Times* documents the existence of online firms that specialize in selling fake followers and names celebrities (politicians, artists, athletes, etc.) who use these services.¹ From an economic viewpoint, these firms create both problems and opportunities. On the one hand, fake follower acquisition is a major concern for social media platforms as it casts doubt on the validity of their metrics (followers, likes, views, etc.). Increasing mistrust in online popularity measurements among users and advertisers can significantly harm the revenue of social media platforms (lower advertising campaigns, decreasing audience size, etc.). As an illustration, social media platforms try to circumvent the expansion of Fake Followers (FFs) by identifying such accounts and suppressing them even if this leads to a decrease in the total number of users. On the other hand, firms selling FFs create an opportunity for people interested in enhancing (unlawfully) their online popularity. Popularity is indeed a major driver of superstars' value or earnings, in addition to performance and bargaining power Adler (1985), Bebchuk et al. (2003), and Rosen (1981). Especially documented in sports, this relationship highlights the economic incentives associated with performance and popularity improvement for professional athletes.² While enhancing sports performance involves time and effort, increasing online popularity can be far easier with the purchase of fake followers. Manipulating online information to influence the perceived quality of products or services has indeed been observed in different contexts, such as restaurants, hotels, the app market or e-commerce (He et al., 2022; Li et al., 2016; Luca and Zervas, 2016; Mayzlin et al., 2014). These manipulations rely on the use of fake reviews or downloads, and take place on different types of online platforms: crowdsourced reviews (TripAdvisor, Expedia, Yelp, etc.), e-commerce (Amazon) or app stores. This paper aims to extend this literature by addressing a new type of quality manipulation on a different platform: the online popularity manipulation of athletes on social media. The aim of this paper is twofold. Firstly, it seeks to provide evidence of the acquisition of fake Twitter followers in soccer for the purpose of manipulating online popularity. Secondly, it aims to highlight the economic motivations of this manipulation by providing measurements of the economic gains of such a practice.

¹The investigation is available here: <https://nyti.ms/2Fm5rCC> (last retrieved: March 2022).

²See for instance Lucifora et al. (2003) for soccer, Vincent et al. (2009) for hockey, Treme et al. (2009) for American football or Ertug et al. (2013) for basketball.

Two literature strands are relevant to this study. The first - the economics of superstars - highlights the positive effects of talent and popularity on celebrities' earnings Adler (1985) and Rosen (1981), and thus provides the main incentive scheme underlying the manipulation of (online) popularity. While Rosen (1981) highlights the positive impact of talent on popularity and earnings, Adler (1985) underlines that popularity can generate large earnings. Investments made by performers in their popularity (TV shows, magazines, etc.), whether or not they are legal, can have a crucial role in their becoming a superstar and then largely affects celebrities' earnings. Even if the causality is hard to determine, those positive relationships between popularity, talent and earnings have been observed in many sports: hockey (Vincent et al., 2009), American football (Treme et al., 2009), basketball (Ertug et al., 2013) and soccer (Carrieri et al., 2018; Franck et al., 2012; Lehmann et al., 2008; Lucifora et al., 2003). The positive impact of popularity on value and the possibility of artificially inflating popularity online raise the question of possible manipulation by celebrities.

The second relevant strand of the literature is the one on online information manipulation. Mayzlin et al. (2014) and Luca and Zervas (2016) highlight the manipulation of online reviews by hotels and restaurants in order to enhance their own online reputation, or damage that of a competitor, and thereby boost their business. These authors also provide useful methodological approaches for the detection, measurement and analysis of online information manipulation. Focusing on online hotel reviews, Mayzlin et al. (2014) show that independent hotels have more fake positive reviews when they are located closely to competitors, suggesting that these hotels post the fake reviews themselves. In the case of restaurants, Luca and Zervas (2016) provide evidence that relatively new ones with a low reputation are associated with fake positive reviews so as to increase their online reputation. Again, this tends to show that these restaurants post the fake positive reviews themselves. In a more general way, Mayzlin et al. (2014) argue that when a valuable characteristic of a product is not, or imperfectly, observable, people who can take advantage of its higher levels can be interested in influencing it through fake information. Professional soccer players can thus be another field of research for the online information manipulation as their popularity is a valuable asset that can be approximated by their number of followers on social media and manipulated through the acquisition of fake ones.

For this purpose, we created a unique dataset of 1,075 international soccer players com-

ing from 43 countries worldwide for the period 1997-2019. It contains more than 500,000 observations about players' careers (statistics, clubs, transfers, achievements, etc.) collected from Transfermarkt.com (www.transfermarkt.com). This website also provides information about players' social media presence (Facebook, Twitter and Instagram), that we complement with manual research of social media accounts in a number of cases. In a similar way to Luca and Zervas (2016), we rely on the detection of fake accounts made by the platform - Twitter - to proxy the number of FFs associated with players' Twitter accounts. More precisely, we use the variation in the number of followers on players' Twitter accounts that occurred on the 12 July 2018 due to the large purge of suspicious accounts carried out by Twitter.³ According to Twitter, these suspicious accounts were removed from account statistics, in particular to 'ensure that malicious actors aren't able to artificially boost an account's credibility permanently by inflating metrics like the number of followers'.⁴

Following the methodological approach used in the literature on cheating (DellaVigna et al., 2010; Duggan et al., 2002; Jacob et al., 2003), we do not directly observe the acquisition of FFs, but our analysis shows consistent patterns for those players associated with FFs. They tend to move to another club few times after creating their Twitter accounts. Once Instagram became the most popular social media platform for soccer fans around 2015, the numbers of FFs associated with newly created Twitter accounts was no longer significant. This suggests it is not worth purchasing FFs any more. The results show that FFs do not have a constant impact on players' valuations, that is, transfer fees, and thus this does not capture an invariant behaviour or characteristic of players. FFs only have a positive and significant impact on the transfer fee if the transfer happens between one and six months after the creation of a Twitter account. In this case, FFs are associated with an average 6 per cent rise in the transfer fee (\approx €650,000). This economic gain far exceeds the costs of buying FFs and is consistent with the purchase of FFs to boost and monetize the online popularity of a newly created Twitter account for transfer negotiations.

The paper is organised as follows. The two relevant literature strands in this paper –

³Twitter announced in July 2018 that suspicious accounts would soon be removed from followers counts (see <https://twitter.com/Twitter/status/1017077041865412609>, last retrieved: March 2022). Articles in *The New-York Times* (<https://www.nytimes.com/2018/07/12/technology/twitter-followers-nyt.html?action=click&module=RelatedCoverage&pgtype=Article®ion=Footer>, last retrieved: March 2022) or BBC News (<https://www.bbc.com/news/world-us-canada-44815550>, last retrieved: March 2022) document the removal consequences, which has led to a 6% decrease in the total number of followers on Twitter on July 12 2018.

⁴Details about Twitter' policy against malicious behaviours using fake accounts are available here: https://blog.twitter.com/official/en_us/topics/company/2018/how-twitter-is-fighting-spam-and-malicious-automation.html, last retrieved: March 2022. This post is also accessible by clicking on the hyperlink "here" in the document "Confidence in follower counts" attached in the tweet announcing the removal of suspicious account.

the economics of superstars and online information manipulation – are presented in Section 2. The third section introduces the data. Section 4 exposes the empirical evidence of FF acquisition by some international soccer players. In Section 5, the econometric model and descriptive statistics are provided. The results are presented in Section 6. The last section concludes and discusses the implications of this study.

3.2 Popularity, Value and Online Information Manipulation

Two literature strands are relevant to this study. The first stream deals with the impact of popularity, talent and bargaining power on superstars' earnings Adler (1985), Bebchuk et al. (2003), and Rosen (1981). This framework fits particularly well for studying professional athletes' earnings. Significant impacts of popularity, talent and bargaining power on athletes' earnings have been shown in many sports. This provides the main economic mechanism that underlines the manipulation of online popularity. As popularity is a valuable asset for superstars (athletes, singers, writers, top executives, etc.), inflating it – unlawfully or not – can lead to higher income for them. The second relevant literature strand addresses cheating behaviours with a particular interest in online information manipulation. When they can generate economic gains, cheating behaviours have been observed in different offline contexts (DellaVigna et al., 2010; Duggan et al., 2002; Jacob et al., 2003). More recently, scholars' attention has been attracted to online information manipulation by firms, especially in the hotel, restaurant or e-commerce sectors. These studies provide evidence of the use of fake online reviews by firms to manipulate their own reputation or that of one of their competitors (Luca and Zervas, 2016; Mayzlin et al., 2014).

3.2.1 Superstars' Value: Talent, Popularity and Bargaining Power

In his analysis of superstars' earnings, Rosen (1981) argues that small differences in individual talent can lead to large differences in earnings. This outcome applies to situations where three main assumptions are satisfied. First, people are able to identify who are the most talented performers. Second, performers are not perfect substitutes and consumers prefer to be served by 'the best' ones. The last assumption considers that large-scale economies exist in the distribution or access to performers (TV, Internet, book, stadium, etc.) so that the best performers can draw large audiences. In such winners-take-all context, the best

performers win the competition and take most of the market.

Even if talent is also a prerequisite to becoming a superstar, Adler (1985) argues that popularity can lead to large differences in earnings between evenly talented individuals. Conversely to Rosen (1981), he assumes that information about individuals' performance is imperfect, and therefore, costly to acquire. Discussions with acquaintances will be the main channel of information acquisition about performance. Superstars may appear thus because they might be known initially by slightly more people than other artists or athletes of similar talent. The most popular stars might be more frequently promoted by people and their performances will be more appreciated. As a consequence, more notorious performers can face large difference in earnings due to positive externalities associated with popularity. Investments made by performers to increase their popularity (TV shows, magazines, social media, etc.) can therefore be crucial to becoming a superstar and earning a high salary (Adler, 2006). Influencers on social media fit particularly well into this framework (Anand et al., 2022), as their remunerations depend on advertising contracts that are based on the size of their audiences (i.e. followers).

More recent studies claim that bargaining power can be an important determinant of high earnings in addition to talent and popularity. Focusing on large companies, Bebchuk et al. (2003) argue that the rising earnings of top managers results notably from information asymmetry between them and shareholders. As shareholders can only imperfectly observe top managers' productivity, the latter take advantage of this asymmetry to fix very high earnings for themselves. The role of bargaining power in rising wages has mainly been explored in the top executive context. For instance, Malmendier et al. (2009) study the impact of CEOs' awards on their incomes. They show that becoming an awarded or superstar CEO leads to growing compensations (stocks, options, etc.). This result is observed specifically in the case of firms with weaker shareholder protection. The main explanation is that the new status of awarded CEOs allow them to negotiate higher compensations.

Admittedly, the main limit of superstar analyses in arts and entertainment relies on the absence of a robust measure of talent or quality, apart economic success (Connolly et al., 2006). It is almost impossible in these domains to use a valid and quantifiable talent measure as talent has a strong subjective component. This limitation is alleviated in the context of sport, which also matches many other prerequisites of superstar theories. First

of all, professional athletes account for some of the top income individuals (Atkinson et al., 2011) and their wages are regularly negotiated. Statistics about performance (shoot, goals, tackle, etc.) or popularity (mentions in newspapers, followers, game audience, etc.) are easily accessible. Lastly, sports games are widely broadcast (TV, radio, streaming, etc.).

Among popular sports, soccer has frequently been studied to test the existence of superstar effects.⁵ Consistent with Rosen's theory, Lucifora et al. (2003) show a concave relationship between the earnings and performance of 533 players in the Italian premier (*Serie A*) and second (*Serie B*) leagues in 1994-95.⁶ This effect holds, however, only for forward players with very high rates of goals and assists per game. In this study, popularity also has a positive and significant effect on earnings.⁷ Garcia del Barrio et al. (2007) focus on the determinants of players' market value in the Spanish premier league (*LaLiga*). They collected data about 369 players for the 2001-2002 season. Performances are approximated by a composite index computed by sports journalists and experts. Popularity is measured by the number of websites quoting a player. The results show a positive impact of performance. In line with Adler's theory, they highlight a more than proportional effect of popularity on the market values for the most famous players. Franck et al. (2012) investigate the impacts of talent and popularity on players' market values in the German premier league (*Bundesliga*). They created a panel dataset (2001-2006) that includes 605 players and a wide range of statistics to measure performances (goals, assists, shots, blocks, saves and fouls). Popularity is captured through the numbers of times players are mentioned in 20 German newspapers. Market values are extracted from Transfermarkt.com and *Kicker*. The results show that both performance (goals and assists) and popularity have positive effects on the market value of players. Consistent with Rosen's and Adler's theories, these effects are stronger for the top players.

Some studies, even if they find positive impacts of popularity and performance on players' values, do not fit exactly within Rosen's and Adler's frameworks. Focusing on the German premier league, Lehmann et al. (2008) measure the impact of performance and popularity on the annual salaries of 264 soccer players in the German premier league (1999-2000). Results provide evidence of positive impacts of both performance (goal, assist, shot on goal and

⁵See also Vincent et al. (2009) for hockey, Treme et al. (2009) for American football or Ertug et al. (2013) for basketball.

⁶Earnings are measured through players' gross salaries excluding bonuses and signing-on fees.

⁷Popularity is measured with dummy variables. Two types of players are considered as popular. Players with a goals per game ratio between 0.25 and 0.4, and players with a goals per game ratio above 0.4.

successful tackles) and popularity (number of mentions in newspapers) on players' earnings. They do not find, however, more-than-proportional increase effects on players' salaries of both performance and popularity, as proposed in Rosen's and Adler's theories. According to Lehmann et al. (2008), this discrepancy with the literature on this topic is due to the relative lack of wealth of the German league compared to the Italian, English or Spanish ones. These latter are more likely to host the soccer superstars as shown by Lucifora et al. (2003), Garcia del Barrio et al. (2007) or Carrieri et al. (2018).

Following Bebchuk et al. (2003), Carrieri et al. (2018) introduce bargaining power as a determinant of soccer players' wages in addition to talent and popularity. They argue that soccer players can benefit from information asymmetry with respect to club owners. More precisely, talented and/or popular players can negotiate higher wages with club owners by threatening them to move to a different team. Players' agents may help them in this purpose. Carrieri et al. (ibid.) focus on players appearing in the Italian premier league (*Serie A*) during the 2013-2014 season. They collected players' annual wage information from on the most read Italian sports newspaper (*La Gazzetta dello Sport*) and information about performance (assist, goal, grade, etc.) from two websites dedicated to soccer (Transfermarkt.com and Soccerways.com). Popularity is approximated using the number of annual search queries on Google for each player. For each player, bargaining power is measured using the total value of the players managed by his agent. Carrieri et al. (ibid.) assume that agents with a large total market value of players are able to bargain advantageous deal for them with clubs' owners. Their results highlight the positive and significant effects of both performance and popularity. They also find a positive effect of bargaining power on players' wages. This illustrates the important role of the player's agent in wage negotiations with the club's owner, and provides evidence that powerful agents are able to negotiate better deals for their clients.

As summarised by Carrieri et al. (ibid.), talent and popularity are two valuable assets for soccer players as teams' owners want to acquire talented players in order to both increase revenue (from tickets, merchandising, TV rights, etc.) and improve the team's performance. This result provides us with the main economic incentive underpinning the manipulation of online popularity. Higher popularity may lead to a higher value or income for a player.

Table 3.1: Summary of literature review

Article	Performance	Popularity	Bargaining	Data
Lucifora et al. (2003)	Goals, assists, strikes and appearances	High scorer	Not included	Italian premier and second leagues, 1995-96
Garcia del Barrio et al. (2007)	Composite grades from sport journalists/experts	Number of websites quoting a player	Not included	Spanish premier league, 2001-02
Lehmann et al. (2008)	Goals, assists, shots on goal and tackles	Citations in sports newspaper (<i>Kicker</i>)	Not included	German premier league, 1999-2000
Franck et al. (2012)	Goals, assists, shots, blocks, saves and fouls	Number of articles mentioning a player across 20 newspapers	Not included	German premier league, 2001-2006
Carrieri et al. (2018)	Goals, assists and average grade by season	Number of annual search queries (Google)	Total market value of player's agent	Italian premier league, 2013-14

3.2.2 Cheating Offline and Online: Performance and Online Information Manipulations

Illegal activities are difficult for economists to analyse as those who engage in these try not to leave a trail (Duggan et al., 2002; Jacob et al., 2003). The lack of formal measure of illegal behaviours involves scholars to adopt approaches that show up indirect evidence of such actions by providing consistent patterns about cheaters and document the economic mechanisms that underpin such behaviours (DellaVigna et al., 2010; Duggan et al., 2002; Mayzlin et al., 2014). The literature has already highlighted cheating practices among professional athletes motivated by economic gains: point shaving by basketball players (Wolfers, 2006)⁸, match rigging by sumo wrestlers (Duggan et al., 2002) or tennis players (Jetter et al., 2017). While previous articles show athletes willingly underperform to gain direct or delayed economic benefits, we focus here on another type of cheating among professional athletes or celebrities, namely online reputation manipulation.

Strategic manipulation of online information has been already documented in empirical studies for products and services: hotels (Mayzlin et al., 2014), restaurants (Luca and Zervas, 2016) or e-commerce (He et al., 2022).⁹ The variety of contexts in which such practices

⁸Some scholars discuss the findings of Wolfers (2006), see for instance Bernhardt et al. (2010).

⁹Another article can be added to this literature: the study of Li et al. (2016) about app developers who 'buy downloads' to gain visibility in app stores. Developers can buy these 'fake downloads' on online platforms that reward smartphone users for this purpose (in-app products, gift cards, etc.). The authors show a positive impact of 'fake downloads' on app rankings. This manipulation is driven by economic gains, as a higher ranking or increased visibility is expected to lead to more downloads from truly interested (and

can be observed is not surprising. As mentioned by Mayzlin et al. (2014), when a valuable characteristic of a product is not, or is imperfectly, observable, people who can take advantage of its higher level of quality of it can be interested in influencing it through fake information. In their paper, Mayzlin et al. (ibid.) investigate possible online review manipulation in the hotel sector. They collected online reviews posted on Expedia and TripAdvisor websites in 2011 for 2,931 American hotels. These platforms have different reviewer identity verification processes, making it more costly for a hotel to post fake reviews on Expedia than TripAdvisor. Indeed, anyone can post a review on TripAdvisor, while only people who have purchased a hotel stay on Expedia can post a review. Mayzlin et al. (ibid.) use the differences in review distributions across the two platforms to identify possible hotel review manipulations. Based on difference-in-difference models, the results show that independent hotels have lower (1 or 2 stars) or higher (5 stars) online reviews when they are close to competitors (<0.5 km). This is consistent with hotels posting fake positive reviews for themselves, and negative ones for competitors. They did not find such effects for affiliated or multi-unit hotels. The authors argue that, for such organisations, the potential reputation damage involved in manipulation detection would affect all affiliated or owned establishments, outweighing the benefits generated by fake positive reviews for only one establishment.

Luca and Zervas (2016) address a similar issue involving online restaurant reviews. They collected more than 300,000 online reviews posted on Yelp between 2004 and 2012 for 3,625 restaurants in Boston. They rely on Yelp's filtering system, which identifies possible fake reviews, to measure this phenomenon. About 50,000 reviews (16% of their sample) were identified as fake by Yelp. The authors highlight two main practices of online reputation manipulation. The first one consists of restaurants in manipulating their own reputation by posting fake positive reviews. This is particularly the case for restaurants facing a decrease in their reputation and also for newly created restaurants with low reputation (i.e., low ratings). The survival of newly created restaurants can be at stake if its online reputation is particularly low. This practice tends to be implemented more by independent restaurants than by affiliated restaurants. If the manipulation is detected, this fraud will indeed negatively affect all affiliates in the chain and then generates large 'cumulative' prejudices for the chain, according to the authors. Moreover, chain restaurants tend to rely on other forms of promotion and branding to build their reputation (Luca, 2011). The second strategy

'monetisable') users.

outlined by the authors consists of posting fake negative reviews about competitors. They find evidence of such a practice among independent restaurants offering the same type of food located within one kilometre.

Focusing on a major e-commerce platform (Amazon.com), He et al. (2022) investigate the use of fake product reviews. They used requests for fake reviews posted on Facebook groups as an indicator of possible manipulations. During a four-month period in 2019, they watched over these groups and detected approximately 1,500 products associated with requests for fake reviews. During the same period, they collect information (ratings, number of reviews, prices, sale ranks, etc.) on Amazon for these 1,500 products and 200,000 similar products (i.e., competitors). This dataset makes it possible to compare the evolution of products with and without requests for fake reviews in Facebook groups. The results highlight that products associated with requests for fake reviews exhibit an increase in their number of reviews, ratings and sales. These effects hold, however, only in the short-term and start to disappear after two weeks. In the long-term, ratings and sales reduce significantly as the products receive many one-star ratings from unsatisfied customers. It seems, then, that fake product reviews are used by providers of low-quality products to temporarily boost their visibility and sales.

Most of the literature about online information manipulation is empirical and mainly focus on the use of fake reviews to influence the perceived quality of products or services. Closer to our paper, the theoretical work of Anand et al. (2022) provides a game-theoretic model of fake follower acquisition by influencers on social media. The remuneration of influencers by marketers is indeed driven by the size of their audience. The larger their audience, the better a marketer will pay them to endorse its products, and reach potential customers. The model of Anand et al. (ibid.) shows that only influencers with intermediate follower counts benefit from buying fake followers. The potential reputation costs are too high for the highest type of influencer in case of manipulation detection, while for the lowest type, the costs of acquiring fake followers and interactions exceed the potential gains.

Online information manipulation should be observed many situations where a valuable characteristic can be difficult to observe (Mayzlin et al., 2014). In the soccer context, popularity is a valuable asset for players that can be approximated by the number of people

following them on social media. However, this metric can be manipulated through the acquisition of FFs. The following sections aim to provide evidence of such practices in the soccer milieu and highlight their economic motives.

3.3 Data and Empirical Explorations

This section provides descriptions of the data used in this study and proposes preliminary analysis of soccer players' careers and social media presence. We first describe the main sources of information used to create the dataset (Transfermarkt.com, Twitter.com and Instagram.com). Then, we document the removal of suspicious accounts held by Twitter on 12 July 2018 in the second subsection. We have taken advantage of this event to measure the number of FFs associated with soccer players. We test in the third subsection a possible alternative scenario explaining the existence of FFs. Fake accounts may result from the social spamming strategy of malicious actors, rather than being acquired deliberately to manipulate online popularity.

3.3.1 Player Career and Performance

We focus in this study on international soccer players for two main reasons. Firstly, they are talented players who attract the attention of fans and the media. Online popularity can thus be a valuable asset for them (Carrieri et al., 2018; Garcia del Barrio et al., 2007). Secondly, they also attract attention from specialized websites, and then, various statistics about them (value, performance, clubs, etc.) are relatively more available. We include in this study players from 43 national teams: the 32 national teams which qualified for the 2018 World Cup (15 June-15 July 2018) and 11 national teams which did not qualify.¹⁰ The full sample contains 1,075 international players. The specialized website on soccer, Transfermarkt, has been used to collect data about players' performances and careers. This website has been chosen because it is considered to be a reliable source of information about players (statistics, careers, achievements, etc.), and especially their values (Deutscher et al., 2016; Franck et al., 2012; Müller et al., 2017). It includes basic information about

¹⁰The 32 national teams qualified for the 2018 World Cup were: Argentina, Australia, Belgium, Brazil, Colombia, Costa Rica, Croatia, Denmark, Egypt, England, France, Germany, Iceland, Iran, Japan, Morocco, Mexico, Nigeria, Panama, Peru, Poland, Portugal, Russia, Saudi Arabia, Senegal, Serbia, South-Korea, Spain, Sweden, Switzerland, Tunisia and Uruguay. In addition, we selected national teams not qualified for the 2018 World Cup in each of the five main geographical soccer areas (Europe, Africa, North and Central America, South America and Asia-Pacific) distinguished by the International Federation of Football Association (FIFA): Algeria, Austria, Cameroon, Chile, China, Italy, Ivory Coast, Netherlands, New Zealand, Romania and USA.

players (age, nationality, size, position, etc.) and also a lot of information about their careers (clubs, achievements, agents, sponsors, etc.), performance (goals, assists, injuries, etc.), value (transfer fee, loan fee, etc.), and even links toward their social media accounts (Facebook, Twitter and Instagram) in some cases. Table C.1 in Appendix C.1 provides some summary statistics about the players included in our sample. The collected data for the 1,075 players span over the period 1997-2020. They include about 520,000 player-match day observations (club and national team, competition and friendly games), 6,000 achievements (competition winner, award, etc.), 6,500 transfers and 20,000 market value updates.

A particular interest in this article is dedicated to transfer information. In this paper, the metrics used to measure the value of players are the fees associated with their transfers and loans. When a player switches from one club to another, it is generally associated with financial negotiations (wages, transfer fees, contract duration, etc.) between the player and the buying club, and also the selling club if the player's contract is not finished. Five different types of transfers are typically observed in soccer: transfer with fees, loan with fees, loan (without fees), free transfer and end of loan. More precisely, *transfers with fees* occur when a player leaves a club before the end of his contract. The buying club must pay a fee to compensate the early end of the contractual relationship between the player and the selling club. *Loans* and *loans with fees* are contracts allowing a player to play for another club for a short period and then come back to their original club. In most cases, these contracts last for one season and can be associated or not with fees paid by the destination club. When a club recruits a player once his last contract has ended, this is a *free transfer*. In this case, the new club does not have to pay compensation to the player's former club, but it is often associated with a signing bonus for the player. Finally, *end of loan* refers to players who finish their loan contracts and go back to their original clubs. This is similar to a free transfer as no charges are associated with the two clubs involved.

Table 3.2 below provides statistics about the different types of transfer. In this article, we only take into consideration in this article transfers that occurred once the players reached 18 years of age. In some cases, the precise nature of a transfer is not detailed on Transfermarkt.com, and thus a category of 'undefined transfer' has been created. The statistics are provided for all our observations ('Full sample') and for the subsample of players who have a Twitter account ('Twitter').

	Full sample		Twitter	
	No. of transfer	%	No. of transfer	%
Transfer with fee	1,758	26.8	1,180	27.5
Loan with fee	248	3.8	172	4.0
Free transfer	929	14.2	526	12.3
Loan	1,028	15.7	707	16.5
End of loan	1,221	18.6	841	19.6
Undefined transfer	1,365	20.8	858	20.0
Total no. of transfers	6,549		4,284	

Table 3.2: Descriptive statistics on transfer types

During players' early careers (18-20 years old), loans and transfers with fees tend to occur in the same proportion (see Fig. C.1 in Appendix C.2). Then, transfers with fees become more prominent. Undefined transfers mostly concern young players. These transfers thus seem to refer in most cases to young players who move from the youth team to the reserve team, or to the professional team. In most cases, this is associated with the signing of a new contract. As they may not be so popular or talented at that time and/or play in minor championships, information about the new contracts is not publicly advertised.

	Transfer fee (million €)		Loan fee (million €)	
	Full sample	Twitter	Full sample	Twitter
Mean	8.7	10.2	1.6	1.9
Standard dev.	15.3	17.2	3.0	3.4
Minimum	0.006	0.006	0.008	0.015
Maximum	222	222	35	35
Total no. of transfers	1,758	1,180	248	172

Table 3.3: Descriptive statistics on transfer fees

Statistics about fees highlight the economic importance of soccer. The cumulative amount of transfer fees for these 1,075 players, even if some of them are far from the end of their career, is approximately €14 billions. Transfer fees are for most small- and medium-sized clubs is their main source of income. Some clubs ground their financial sustainability on *players' trading*, that is, discovering and recruiting talented young players in order to sell them for much more money than they buy them.¹¹ Future sales are even included in the club's budget. For example, in the French premier league, clubs are audited by an independent organisation at the beginning of each season about how they plan to manage their budget, and especially how much revenue they expect from selling players. Clubs found to

¹¹See, for instance, this article in *The New-York Times*: <https://www.nytimes.com/2008/07/19/business/19soccer.html>, last retrieved: March 2022.

be in breach of the rules can have a number of punishments enforced upon them, including transfer embargoes, limitations on the number of first-team players, and can be even relegated to a lower league.¹² Quansah et al. (2021) and Poli, Besson, et al. (2022) detailed how much clubs depend on the income generated on the transfer market to balance their books. The remuneration of players' agents is composed of a wage paid by the player but also by a profit share of his player's transfer fee. The higher the value of the transactions, the better they are paid, perhaps as a direct percentage of the transfer or through assignments.

3.3.2 Social Media and Online Popularity

Information about players' online popularity (followers, friends, etc.) has been collected on Twitter and Instagram from the 14 June to 9 September 2018, twice a day (at 11 a.m. and 11 p.m.).¹³ This information was gathered both using social media API and direct scraping of players' accounts. For each social media platform, we identify the player's account using the following four-step protocol: (i) check if the link toward the player's social media account is available on Transfermarkt, (ii) searches for the player's account on Google using a request of the type: '*Firstname Name Twitter or Instagram*'¹⁴, (iii) check if the player appears among the followers of his club's or national team's social media accounts, and (iv) check if the player appears among the followers of players of his club or national team. If no account has been found during all the previous steps, we consider that the player does not have an account on the corresponding social media platform. We can miss players' social media accounts even if we follow the entire protocol. In this case, it is likely that this player is not at all popular online, as we found with this protocol we still found players who only had a few hundred followers.

Using this protocol, we identified 664 Twitter accounts (61.8%) and 824 Instagram accounts (76.7%) in our sample of players. Most of these are verified profiles, meaning that the account owner's identity has been checked by the social media platform, with respectively 83.6% and 83.4% for Twitter and Instagram.¹⁵ It is noticeable that 561 players (52.2%)

¹²See, for instance, the case of French soccer team of Bordeaux: <https://www.besoccer.com/new/bordeaux-and-angers-spared-relegation-from-ligue-1-by-french-authorities-1017828>, last retrieved: March 2022.

¹³Facebook was initially included. As we observed that very few players have an account on this social media platform, we decided to only take Twitter and Instagram into account in the study.

¹⁴We begin by searching for the name and first name using the Western alphabet for all players. If we do not find the account and that player's name is not originally written using the Western alphabet, we search for his name and first name with its original alphabet. The Transfermarkt website is very useful here as it provides the player's name and first name using their original alphabet (Arabic, Chinese, Cyrillic, Japanese, etc.).

¹⁵We include unverified profiles in the analysis if they satisfy the following conditions: player names and first names are used for the social media's screen name, accounts are described as belonging to the corresponding soccer player, the player is clearly identifiable

have an account on both social media. In Figure 3.1, we plot the statistics about the year of creation for Twitter and Instagram accounts and the number of followers in July 2018. We can see that Twitter account creation occurred earlier, with the first ones in 2009 and a peak in 2011. Instagram accounts started to be created in 2011 with a peak in 2013. Although Twitter seems to be the favourite social media platform used by soccer players at the beginning of the 2010s, Figure 3.1b shows that Instagram is more widely used a few years later.




Figure 3.1: Descriptive statistics about Twitter and Instagram

3.3.3 Twitter and Fake Followers

FFs are typically used to implement two malicious strategies: social media spamming and an artificial increase in follower counts.¹⁶ Social media spamming consists of following a popular person by using a fake account and posting ad-related content on their wall that will be potentially viewed by all their followers. In the case of the metric manipulation, a firm selling FFs will be contacted and a certain number of them will be purchased by the account owner or her relatives. Social media login information must be given to the firm, so as it can manage and control that the requested amount of FFs have really started to follow the targeted account.

These two practices, and the FFs generated for it, are thus major threats for social media platforms. Social spamming downgrades users' experience and information, and can lead them to reduce and even stop using the social media platform. Regarding FFs used

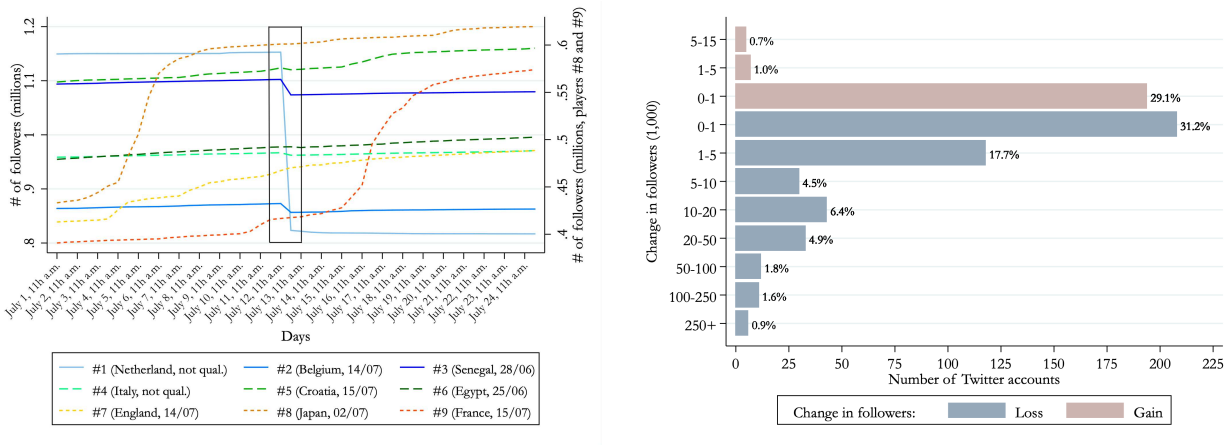
in 'personal' pictures in a soccer and/or personal context and the profile has at least 1,000 followers.

¹⁶For more details, see Twitter's policy against spamming and malicious behaviours (https://blog.twitter.com/official/en_us/topics/company/2018/how-twitter-is-fighting-spam-and-malicious-automation.html, last retrieved: March 2022).

to unlawfully enhance online popularity, this practice creates mistrust about social media metrics among advertisers and users. Users can feel fooled when discovering or understanding that the statistics of the people they follow are false and thus reduce their use of social media. Losing advertisers' trust is perhaps even more harmful for social media platforms. Advertisers can feel overcharged when paying for ad-campaigns due to the presence of FFs. They can then decide to reduce, and even stop, this expenditure, which is the main source of revenue for social media platforms.

In order to mitigate these threats, Twitter removed accounts suspected of being used for spamming or metrics manipulation. In particular, a large purge of fake accounts was carried out by Twitter on 12 July 2018.¹⁷ The twice-a-day collection of social media metrics allows us to accurately measure the number of suspicious accounts associated with the international soccer players in our sample. Figure 3.2a depicts typical trends observed for followers around 12 July 2018. It shows that fake follower removal did not impact all players. Some were not affected at all (small dashed lines) or only very marginally (long dashed lines). However, other players experienced significant and even a critical decrease in their followers (solid lines). Figure 3.2b provides summary statistics on the evolution of followers for our sample of players on 12 July 2018. It shows that the sample can be roughly divided into three tiers. The first tier of the sample gains some followers (in most cases, less than a thousand), the second tier faced very few losses (less than a thousand) and the last tier experienced large loss of followers. The decrease can reach more than 10,000 of followers for 105 players (15.7% of our sample). Figures 3.2a and 3.2b show that large FF losses are not widespread among players. However, some players are associated with large numbers of FFs, which raises the question of possible strategical use in these cases.

¹⁷On 11 July 2018, Twitter announced the removal of fake accounts would occur the next day (see <https://twitter.com/Twitter/status/1017077041865412609>, last retrieved: March 2022). This tweet includes a link toward a post presenting Twitter's policy about 'Confidence in follower counts' (see https://blog.twitter.com/official/en_us/topics/company/2018/Confidence-in-Follower-Counts.html, last retrieved: March 2022). The post contains a link toward more detailed descriptions about how Twitter fights spam and malicious behaviours (see https://blog.twitter.com/official/en_us/topics/company/2018/how-twitter-is-fighting-spam-and-malicious-automation.html, last retrieved: March 2022).



(a) Typical evolution (team, last WC2018 match)

(b) Summary statistics

Figure 3.2: Twitter's removal of FFs on 12 July 2018

3.3.4 Alternative Scenario

Suspicious accounts are used for two main goals: social media spamming and online popularity inflation. Both cases lead to an artificial increase of followers associated with an account, but their originations are clearly opposite. For the acquisition of FFs, the rise stems from a decision taken by the account owner or someone who manages it (the player's agent, communication advisers, club's communication department, etc.). In the case of social media spamming, the account owner suffers from the decision of a malicious actor. Social media spamming indeed consists of posting ad-related contents at cheap cost in order to benefit from the large audience of the targeted account.

In this paper, we investigate the possible manipulation of online popularity; social media spamming is thus a phenomenon that is not in the scope of the study. But as it relies on the same means – fake accounts – we need first to assess whether fake accounts associated with a player result from social spamming behaviour or not. If suspicious accounts removed by Twitter mainly consist of fake accounts created for social media spamming, we might observe a strong correlation between the number of followers of a player and his number of FFs. These correlations are provided in Table 3.4. The number of followers has been corrected with the removal of FFs. Moreover, the sample has been divided into deciles based on the number of followers to have comparable groups of players.

Decile	Corr.	Fake followers		Followers	
		Min.	Max	Min.	Max.
1	.24	0	99	11	2,967
2	.12	0	188	3,026	7,587
3	.07	0	2,591	7,656	15,418
4	.36	0	1,300	15,922	34,379
5	-.03	0	2,709	34,482	77,685
6	.32	0	4,992	78,022	154,131
7	.26	0	12,396	155,085	310,604
8	.32	0	32,355	311,016	778,131
9	.38	0	35,878	794,583	1,646,341
10	.97	114	1,215,302	1,659,446	73,326,669

Table 3.4: Correlation between followers and FFs by decile


For the first nine deciles, Table 3.4 shows that the correlation between fake and 'real' followers is very low and irregular. Moreover, these deciles systematically include some players who are not associated at all with fake accounts. This evidence is not consistent with the wide use of these fake accounts for social media spamming purposes. The last decile, however, exhibits a large and positive correlation. A closer look to this decile shows that this very high level of correlation holds only for the last 5 percent of the sample (see Table C.2 in Appendix C.3). The most popular players exhibit a pattern that matches social media spamming. This result leads us to withdraw the 5 percent most popular players in the remainder of the article and, therefore, to focus on players with less than 3.2 million followers on Twitter.

3.4 Transfers, Social Media and Fake Followers

Social media spamming only explains the presence of FFs for the most popular players online. For the remaining players, there may be other explanations. Further empirical explorations provide us with three main types of evidence that FFs are used to unlawfully inflate players' popularity in our case: i) no more accounts with significant numbers of FFs on Twitter once Instagram becomes the most popular social media platform for players, ii) the strong tie between Twitter account creation and subsequent transfers afterwards, which illustrates the use of social media to monetize popularity, and iii) the strong relationship between a high number of FFs and subsequent transfers involving financial negotiating after the creation to a Twitter account. These findings are detailed in the following subsections.

3.4.1 The Rise of Instagram

A striking feature of Figure 3.3a is that there are no more accounts associated with significant numbers of FFs after 2014. This result echoes the previous statistics we observed in Figure 3.1 (see subsection 3.3.2). In this figure, we observe that followers on Instagram tend to outperform those on Twitter after 2014. Instagram then seems to become the most popular social media platform for soccer players. FFs seem to have disappeared because it is no longer worth investing in them on Twitter as it is no longer the reference for player popularity. The figures in Appendix C.4 focus on players who use both social media platforms and provide additional evidence of this phenomenon. Although these players mostly created their Twitter account first (Figure C.3), most of them have more followers on Instagram in 2018 (Figure C.4). This means that, at some point, the numbers of followers on Instagram have exceeded those on Twitter.¹⁸ Figure 3.3b illustrates this trend by plotting a fitted regression of the number of followers on both social media.¹⁹ This again shows that, on average, Twitter accounts have been created before Instagram ones. We can also see that followers on Instagram outperform those on Twitter in 2015. The rise of Instagram as the main social media platform for soccer players in 2015 is therefore consistent with the disappearance of FFs, and provides evidence of their use for the purpose of online popularity manipulation.



(a) Twitter account creation date and FFs

(b) Estimated average number of followers

Figure 3.3: Twitter account creation, FFs and Instagram growth

¹⁸Figures C.5 and C.6 in Appendix C.4 provide statistics about players using only Twitter or Instagram. These figures confirm that Instagram becomes more popular among soccer players about 2014. When they use only one social media platform, it is far more frequently Instagram. The numbers of followers on Instagram moreover outperform those on Twitter after 2015 as for the players that use both social media platforms.

¹⁹For this, two points are taken into consideration for each player. We know precisely the numbers of followers for both social media platforms in July 2018. Moreover, we assume that the number of followers is equal to 0 on the day of an account's creation.

3.4.2 Social Media and Transfer

A player's professional life is quite cyclical. It is paced by phases of competitions and so-called *mercato*, that is, periods during which players can transfer to another club. These latter occur twice a year with the *Summer mercato* in July and August, and the *Winter mercato* in January. Figure 3.4 provides statistics on Twitter account creation by month from 2009 to 2018 for players accounting respectively for less than 1,000, between 1,000 and 5,000 or more than 5,000 FFs. As FFs are generally sold by thousands, we consider that having less than 1,000 is similar to not having FFs.²⁰ Figure 3.4b highlights that accounts with FFs are over-represented in June, November and December.²¹ These months are special in soccer as they precede the *mercato* periods (July, August and January). It seems that Twitter account creation, and especially for accounts with FFs, occurs more frequently few times before a *mercato*. Table 3.5 shows, moreover, that transfers with fees are over-represented among the transfers which occur after the creation of Twitter accounts. These figures and the table do not indicate, however, if players actually move during the *mercato* following the creation of their Twitter account.

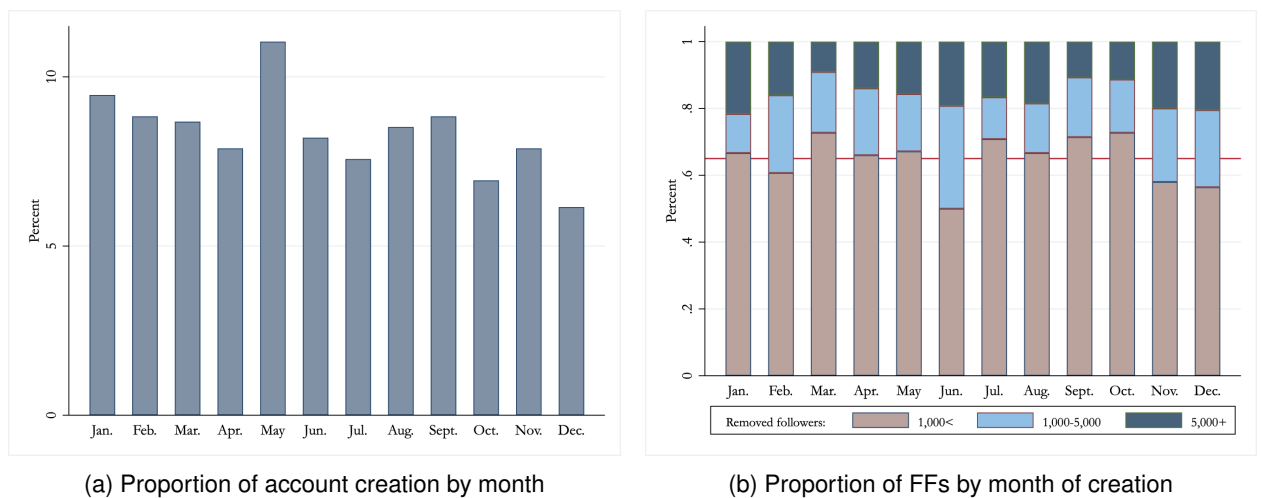


Figure 3.4: Number of Twitter accounts created by month (2009-2018)

²⁰73.6% of players accounting for between 1 and 1,000 FFs have indeed less than 500 FFs.

²¹Proportional difference tests show significant over-representation of accounts with FFs at 5% threshold in June ($z = 2.18$; $pr = 0.015$) and at 15% threshold in November and December (resp. $z = 1.12$; $pr = 0.13$ and $z = 1.19$; $pr = 0.12$). There is also a significant under-representation of accounts with FFs in March ($z = -1.12$; $pr = 0.13$)

	All transfers		Transfers next to Twitter creation	
	2009-19	2009-14	2009-19	2009-14
Transfer with fee	28.0	25.5	40.9	40.3
Loan with fee	4.2	3.6	2.3	2.1
Free transfer	12.4	10.3	14.1	13.7
Loan	17.0	18.1	7.2	7.1
Loan end	20.4	20.3	10.0	10.7
Undefined transfer	18.0	22.2	25.5	26.1
Total no. of transfer	3,561	1,645	432	337

Table 3.5: Twitter account creation and transfer types

Figure 3.5 fills this gap by providing information on the time between Twitter account creation and the player’s next transfer by transfer type. As, on average, a player stays 505 days at a club, Figure 3.5a therefore shows a strong relationship between Twitter account creation and all types of transfer. This trend clearly results from transfers that involve financial negotiating, that is, transfers with fees, loans with fees and free transfers (see Figures 3.5b and 3.5d)²². Loans, which do not involve transfer fees and wage bargaining, do not exhibit different behaviours according to the number of FFs associated with a player. Significant numbers of FFs are mainly associated with accounts created shortly before financial negotiations between a player, the buying club and possibly the selling club.²³ It provides new evidence that popularity is a valuable asset for professional athletes, and especially soccer players, during salary bargaining (Carrieri et al., 2018; Franck et al., 2012; Garcia del Barrio et al., 2007; Lehmann et al., 2008). High adoption rates of social media, respectively 61.8% and 76.7% for Twitter and Instagram, underline players’ interests into monetising their online popularity during transfer negotiations.

3.4.3 Online Popularity Manipulation Strategy

The evidence presented leads us to consider that most FFs have been acquired purposefully to manipulate online popularity. Firstly, the social media spamming explanation does not fit to our data, except for the 5 percent most popular players. The latter have therefore been withdrawn from our analysis. Secondly, recent Twitter accounts are no longer associated with FFs. This evolution results from the rise of Instagram, which offers a larger audience than Twitter. Twitter’s FFs are then no longer attractive for manipulating online popularity.

²²A free transfer involves that the player (and his agent) negotiating his wage with the buying club. Undefined transfers imply that a new contract has to be agreed and thus include wage negotiations in most cases.

²³These Figures are also provided for the sub-period 2009-2014 in Appendix C.5. These Figures highlight the same trends described here for this specific period.

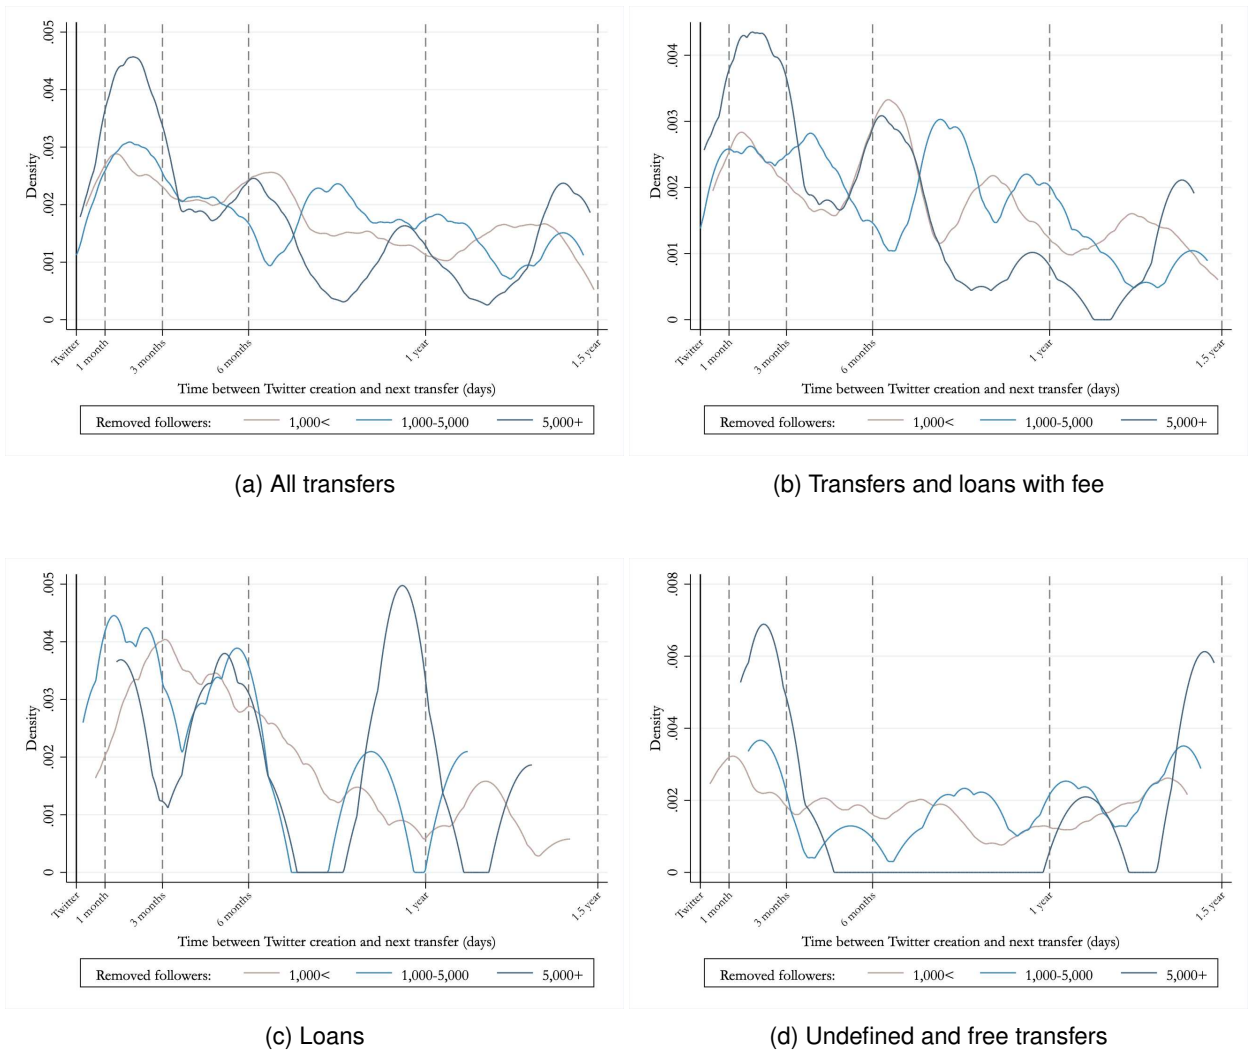


Figure 3.5: Distributions of transfers next to Twitter account creation (2009-2019)

Thirdly, Twitter account creation appears to be strongly related to players' transfers shortly after, especially for players with large numbers of FFs. For the latter, transfers involving fees and/or wages negotiations are over-represented.

This evidence allows us to draw the following scenario about the manipulation of online popularity. Players who want to be transferred create their Twitter account to monetize their online popularity during transfer negotiations. Indeed, the transfer negotiation includes the player's salary in the new club and possibly the compensation for the selling club for breaking the contract. The transfer fee paid by the buying club reflects its expectation of the income which will be generated by the recruited players (Carmichael et al., 1999; Dobson et al., 1999). The higher a player's performance and popularity are, the better income the buying club can expect from the transfer. Skilled players lead to better results and match quality, generating more revenue associated with achievements and attendance. Popular athletes generate revenue for their clubs due to merchandising, ticketing, TV rights, etc. A club can

monitor the performance, however, popularity is more difficult to accurately measure. In this case, some players, agents or selling clubs can be interested in boosting popularity artificially. This might enhance the buying club's expected revenue associated with the recruitment of a player and, therefore, increases its willingness to pay a high transfer fee. This strategy is more likely to occur soon after the creation of the Twitter account, when the number of followers is limited and manipulation can lead to a significant increase in online popularity.

3.5 Econometric Model and Descriptive Statistics

The transfer fee is the main variable of interest in this study. Recruiting a player before the end of his contract involves the buying club paying compensation to the player's current club. The fee is then bargained between the two clubs, including, to some extent, the player's considerations.²⁴ Two conditions must be satisfied for the transfer to occur. First, the transfer fee must reach or exceed the selling club's requirement to let its player leaves. Second, the transfer fee must not exceed the buying club's expected income flows which will be generated by the potential new player's performance and popularity. This framework thus relies on a *hedonic pricing* approach to players' values by clubs and is frequently used in the studies on transfer fee valuations (Carmichael et al., 1999; Dobson et al., 1999).

Using OLS regression, our empirical strategy consists of estimating a transfer fee hedonic price equation that depends on performance, popularity and the bargaining power of players. Of particular interest is given to the impact of FFs for transfers occurring next to the Twitter accounts creation. The model specification is as follows:

$$Fee_{ij} = \beta_0 + \beta_1 FFs_i * Timing_{ij} + \beta_2 Pop_i + \beta_3 Perf_{ij} + \alpha_{ij} + \epsilon_{ij} \quad (3.1)$$

With α_{ij} a series of Fixed Effects (FE) controlling for the year of the transfer, the player's usual position, his age, the destination league and the buying club.²⁵ In the case of age, the corresponding quadratic term is included in the estimations to allow for non-linear effects.²⁶

²⁴Even if, from a technical view point, transfer fees are basically negotiated by the selling and buying clubs, the players' opinion about joining the new club or staying at his current club affects this negotiation. It is difficult for a club to keep a player who really wants to join another club, as he can threaten to stop playing or decrease his performance.

²⁵Buying club FE are introduced for the 130 clubs that appear at least three times in our sample.

²⁶As mentioned by Lucifora et al. (2003), North American studies tend not to jointly use age and experience (or appearance) due to the draft system and the fairly uniform age of entry into the professional league just after college graduation. In this context, age and experience then tend to be co-linear. In European soccer, however, players start their professional careers at many different ages.

The error term is represented by ϵ_{ij} . Performance variables are lagged in order to avoid possible reverse causality problems. These measure the player's performance during the 365 days before the transfer. We add variables that capture players' experience with their clubs and national teams.

Table 3.6 presents summary statistics for the variables used in our estimations. The observation level is per transfer. Our main independent variables of interest are the number of FFs on Twitter, reported as *FF*, and the *Timing* which states the timing of the transfer in relation to the creation of the Twitter account. This variable is composed of seven dummy variables. The first one *Before* is equal to 1, if the transfer happened before the creation of the Twitter account. The five following dummies refer to transfers that occurs next to the creation of Twitter accounts and indicate how much time elapses between the creation of the Twitter account and this transfer. More precisely, *0-1 month after* is equal to 1 if the next transfer after the creation of the Twitter account occurs within one month. Following the same logic, *1-6 months after*, *6-12 months after* and *12+ months after* equal 1 if the next transfer after the creation of the Twitter account occurs, respectively, within 1-6 months, 6-12 months or more than 12 months. The last dummy *After* equals 1 if the transfer happens after the creation of the Twitter account, but it is not closest one.

Players are also characterised using different sets of variables. Regarding bargaining power, we include the total value in Euros of the players managed by the player's agent. Loan with fees indicates if the transfer is a loan associated with a fee (*vs.* transfer with fees). Age and age squared are computed at the time of the transfer, as position (goalkeeper, defender, midfielder and forward) is the usual position of the player on the field. The championship is the destination championship after the transfer. All performance measurements (goals, assists, etc.), achievements, number of games or justice issues (doping) are computed for a period of one year before the transfer.

Table 3.6: Summary statistics

Variable	Description	Mean	S. D.	Min.	Max.
Dependent Variable					
Transfer Fee	Log of transfer fee (€)	14.93	1.56	8.99	18.79
Variables of Interest					
No Insta.	Equals 1 if the player has no Instagram account	0.16	0.37	0	1
Insta. Foll.	Log of followers on Instagram in July 2018	10.56	4.92	0	16.74
Twitter Foll.	Log of followers on Twitter (minus FFs) in July 2018	11.31	2.16	2.4	14.97
FF	Number of FFs on Twitter in thousands	3.6	6.96	0	62.32
Before	Equals 1 if transfer happens before Twitter creation	0.29	0.46	0	1
0-1 month after	Equals 1 if transfer happens 0 to 1 month after Twitter creation	0.01	0.12	0	1
1-6 months after	Equals 1 if transfer happens 1 to 6 months after the creation	0.04	0.2	0	1
6-12 months after	Equals 1 if transfer happens 6 to 12 months after the creation	0.04	0.19	0	1
12+ months after	Equals 1 if transfer happens 12 months or more after the creation	0.12	0.32	0	1
After	Equals 1 if transfer occurs after the transfer next to Twitter creation	0.49	0.5	0	1
Control Variables					
Age	Age of the player at transfer date	24.14	3.39	18	37
Age sq	Squared value of the age	594.04	169.91	324	1369
Value Managed	Log of the total value of the players managed by player's agent	3.94	1.82	-2.3	7.13
Loan with Fee	Equals 1 if the transfer is a paid loan	0.13	0.34	0	1
Defender	Equals 1 if the player is a defender	0.31	0.46	0	1
Midfielder	Equals 1 if the player is a midfielder	0.27	0.45	0	1
Forward	Equals 1 if the player is a forward	0.35	0.48	0	1
England	Equals 1 if the player is transferred to England (<i>EPL</i>)	0.2	0.4	0	1
Spain	Equals 1 if the player is transferred to Spain (<i>LaLiga</i>)	0.12	0.33	0	1
Italy	Equals 1 if the player is transferred to Italy (<i>Serie A</i>)	0.18	0.38	0	1
Germany	Equals 1 if the player is transferred to Germany (<i>Bundesliga</i>)	0.14	0.35	0	1
France	Equals 1 if the player is transferred to France (<i>Ligue 1</i>)	0.1	0.3	0	1
Africa	Equals 1 if the player is transferred to Africa	0.03	0.17	0	1
North America	Equals 1 if the player is transferred to Central or North America	0.04	0.19	0	1
Asia	Equals 1 if the player is transferred to Asia	0.06	0.24	0	1
Middle East	Equals 1 if the player is transferred to Middle East	0.06	0.24	0	1
East Europe	Equals 1 if the player is transferred to Eastern Europe	0.03	0.16	0	1
West Europe	Equals 1 if the player is transferred to Western Europe (Top 5 excl.)	0.18	0.38	0	1
Goals	Number of goals during last year	6.15	7.28	0	44
Assists	Number of assists during last year	4.01	4.41	0	31
Conceded	Number of goals conceded during last year	62.96	17.82	0	122
Captain	Number of games the player was captain during last year	1.55	6.11	0	53
Yellow	Number of yellow cards given to the player during last year	5.06	3.52	0	18
Red	Number of red cards given to the player during last year	0.16	0.42	0	3
Trophies Club	Number of trophies won with club during last year	0.4	0.74	0	5
Trophies Nat	Number of trophies won with national team during last year	0.03	0.16	0	1
Trophies Young	Number of trophies won with club youth team during last year	0.01	0.11	0	1
On	Log of the no. of entries in games during last year	1.45	0.93	0	3.61
Club Nat	Log of the no. of national games with club during last year	3.28	0.59	0	3.97
Club Inter	Log of the no. of international games with club during last year	0.92	1.01	0	2.94
Nat	Log of the no. of games with national team during last year	1.35	0.97	0	3.14
Justice	Equals 1 if suspended during last year (doping, etc)	0.01	0.12	0	1
Observations	80		1,032		

3.6 Results and Discussion

3.6.1 Main Results

Our period of interest starts in 2009, once players in our sample started to create Twitter accounts (See Section 3.3.2). In our models, we include transfers and loans with fees that happened once players reached the age of 18. Table 3.7 reports the results of our regression according to different time periods and measures of popularity. Two periods are indeed considered: 2009-2019, which corresponds to the longest period we can work on, and 2009-2014, which is suspected to be particularly relevant for the manipulation of Twitter followers (see Figure 3.3). We use two measures of popularity for players; Instagram and Twitter followers. All the players we are focusing on have a Twitter account, but not all have an Instagram one (see Section 3.3.2). When measuring the popularity of players using Twitter, we use the log of followers on this social media platform (see Models 1 and 3). If we use Instagram as a measure of popularity, we use the log of the number of Instagram followers for those who have it and the log of Twitter followers – using an interaction term – for those who do not (see Models 2 and 4).

In line with previous empirical studies, Table 3.7 highlights the impact of online popularity on players' valuation. The number of followers on both Twitter and Instagram have positive and significant effects on transfer fees for the period 2009-2019.²⁷ For this period, Instagram has a stronger effect than Twitter. An increase of 1% of Instagram followers generates an average increase in transfer fees of 0.066% compared to 0.042% for Twitter. The results for 2009-2014 confirm that Twitter was the dominant social media platform during this period after which Instagram becomes the main one. Model 4 shows that the positive impact of Twitter is stronger during this period (0.072%) compared to 2009-2019, meaning that the impact of Twitter on transfer fees declines after 2014. On the other hand, Instagram followers are not significant (Model 4) for 2009-2014, suggesting a stronger impact in more recent years.

Besides the impact of online popularity, we focus on a possible impact of FFs on players' valuation. FFs do not have an effect on transfers happening before the creation of Twitter accounts or after the one next to the creation of the Twitter account. They do not capture,

²⁷Including the transfers before players reached 18 years of age enlarges the period under consideration to 2002-2019 and adds 96 transfers to our analysis. These models give very similar results to the ones presented in Table 3.7. Although online followers seem to be a proxy of players' popularity even before the emergence of Twitter, we choose to start to considering transfers only from 2009, i.e., once Twitter started to be adopted by the players in our sample.

Table 3.7: Main specifications

Dependant variable: Transfer Fee (log)	2009-2019		2009-2014	
	Twitter (1)	Insta-Twitter (2)	Twitter (3)	Insta-Twitter (4)
Online Popularity:				
Twitter Foll.	0.042* (0.023)		0.072* (0.039)	
Insta. Foll.		0.066** (0.030)		0.063 (0.046)
No Insta=0 × Twitter Foll.		0.003 (0.029)		0.037 (0.048)
No Insta=1 × Twitter Foll.		0.062** (0.028)		0.097** (0.044)
Creation Timing and FF:				
Before × FF	0.003 (0.012)	0.002 (0.012)	0.009 (0.015)	0.008 (0.015)
0-1 month after × FF	-0.011 (0.016)	-0.011 (0.015)	-0.020 (0.026)	-0.023 (0.026)
1-6 months after × FF	0.015*** (0.006)	0.016*** (0.005)	0.018** (0.007)	0.019** (0.007)
6-12 months after × FF	-0.004 (0.015)	-0.004 (0.014)	-0.004 (0.017)	-0.004 (0.017)
12+ months after × FF	0.011 (0.008)	0.012 (0.007)	0.005 (0.011)	0.005 (0.011)
After × FF	0.002 (0.007)	0.001 (0.007)	-0.016 (0.015)	-0.017 (0.015)
Constant	6.997*** (1.691)	6.508*** (1.808)	8.975*** (2.266)	8.958*** (2.416)
Position FE	Yes	Yes	Yes	Yes
Performance FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
League FE	Yes	Yes	Yes	Yes
Club FE	Yes	Yes	Yes	Yes
Observations	1,032	1,032	422	422
R-squared	0.738	0.741	0.798	0.800

Cluster-robust standard errors (at player level) are shown in parentheses.

* p<.10, ** p<.05, *** p<.01

in other words, a specific behaviour of players that affects all his transfers. FFs only have a significant effect if the transfer follows the creation of the Twitter account and occurs within 1 to 6 months of this. Whatever the specification, the coefficients are very similar within the two-time spans considered in the models. Focusing on 'Twitter' specification (i.e., Models 1 and 3), one thousand additional FFs is associated with an average rise of 1.5% and 1.8% in transfer fees respectively for the periods 2009-2019 and 2009-2014. For the period 2009-2019, the mean value of transfer fees and FFs for players transferred between 1 and 6 months after creating their Twitter accounts are respectively €10.8 millions and roughly 4,000. The average effect of FFs is thus about €650,000, which represents 6% of the average transfer fee. This gain far exceeds the cost of thousands of FFs.

These results are consistent with a scenario in which a player, his agent and/or the selling club unlawfully inflates the player's online popularity to monetize it during a transfer negotiation. A transfer generates indeed two main streams of revenue: the selling club and agent are paid with the transfer fee, while the player will earn the wage negotiated with the buying club, in addition to possible bonuses. According to the literature on soccer players' value, both streams of revenue – wages and transfer fees – are positively impacted by popularity (Carrieri et al., 2018; Garcia del Barrio et al., 2007; Lehmann et al., 2008; Lucifora et al., 2003). A (fictitious) rise in online popularity will lead the buying club to overestimate the expected revenue associated with popularity (attendance, audience, merchandising, etc.) generated by a player. The buying club will thus be prone to offering a higher wage and agreeing to pay a higher transfer fee compared to a situation where the player appears less popular. The player, agent and selling club thus have an incentive to unlawfully inflate the player's online popularity to gain higher wages and transfer fees. They can all agree to implementing the manipulation or only some of them can, creating a positive externality for the others. From the selling club's and the player's agent perspectives, the higher the transfer fee is, the more they earn. They have a direct advantage for manipulating the player's popularity to bargain a higher transfer fee with the buying club. This implies that the selling club or the agent is in charge of the player's social media account, which is a frequent situation.²⁸ From the player's perspective, unlawfully increasing his online popularity might basically help him (or his agent) to bargain a higher wage with the buying club. This (fictitious) rise in online popularity will meanwhile positively affects the transfer

²⁸See, for instance, this BBC article: <https://www.bbc.com/sport/football/56911978>.

fee received by the selling club and his agent. Both the transfer fee and the wage are usually negotiated in a similar time span. Moreover, this higher transfer fee may make the selling club and the agent more likely to support and accept the transfer.

The results concerning control variables are consistent with previous studies (see Appendix C.6). For the period 2009-2019, players' experiences have strong positive effects whether this is with their clubs (national or international competitions) or their national teams (Carrieri et al., 2018; Franck et al., 2012; Lehmann et al., 2008). Age and age squared have a significant impact on transfer fees, although the former has a positive effect while the latter is negative. This inverted U-shape relationship between age and value has been frequently observed (Carrieri et al., 2018; Franck et al., 2012). This reflects that greater experience (tactics, ability, knowledge of the game, etc.) – measured by age and appearances – is offset by worsening physical performance (slower speed, more frequent injuries, etc.) when players get older. There are also positive impacts on transfer fees in 2017, 2018 and 2019 compared to 2009. These results may be driven by the rising trend of transfer fees observed from 2017 in the big-five European leagues, and especially in the English *Premier League* (Poli, Ravenel, et al., 2019). Loan with fee has a negative and significant impact on transfer fees as it corresponds to short-term transfers – one year in most cases – compared to the usual contract duration for transfers (three years or more). Winning achievement(s) with their clubs has a positive effect on transfer fees. Attracting a player from a successful team seems to involve higher costs. The number of goals has a positive effect, confirming that players with a high ability to score are costly to recruit (Carrieri et al., 2018; Franck et al., 2012). Frequent scorers directly enhance their teams' performance in addition to their own popularity. Scoring goals enhances a club's revenue and thus, increases their willingness to pay for this kind of player. The results show a positive impact of an agent's total market value of their portfolio of players on transfer fees. In line with the findings of (Carrieri et al., 2018), it appears that, in addition to performance and popularity, an agent's market power is important for negotiating interesting deals with buying clubs. Most of these results hold for the period 2009-2014, except those related to transfer year and players' age, as the period under consideration is very short for exhibiting time effects.

3.6.2 Robustness Check

As we can expect in this kind of framework, the effect of online popularity and then fake followers can differ depending on where is located the player on the distribution of market values. Most studies on earnings or values of high-skilled and/or popular workers (e.g. singers, CEOs, athletes, etc.) highlight skewed distribution of earnings and non-linear effects of workers' characteristics (popularity, performance, etc.) on the dependent variable resulting from superstar effects. In this context, quantile regression can complete the Ordinary Least Square (OLS) regressions to study the determinants of earnings and values. According to (Cameron et al., 2009), conditional quantile regression is more robust to outliers and does not involve strong assumptions about parametric distribution of regression errors. This model allows moreover to examine the effects of regressors at different points in the dependent variable's distribution, providing thus more information about this relationship. Figure 3.6 displays the distribution of the transfer fees we observe in our data. The quantile regression model permits to take into account this distribution that do not perfectly follow a normal distribution.




Figure 3.6: Distribution of Transfer fees

With the quantile regression model we assume that the θ^{th} quantile, $\theta \in (0, 1)$, of the conditional distribution of the dependent variable (Fee) is a linear function of the regressors (x):

$$Fee_{ij} = \beta_0^\theta + \beta^\theta x_{ij} + u_{ij}^\theta \quad (3.2)$$

Where β_0^θ is a constant term, $\beta^\theta x_{ij}$ a vector of parameters and u_{ij}^θ an error term. Let's $Q^\theta(y_{ij}|x_{ij})$ denotes the conditional quantile function and assumes that the distribution of the error term u_{ij}^θ satisfy $E(u_{ij}^\theta|x_{ij}) = 0$, then:

$$Q^\theta(Fee_{ij}|x_{ij}) = \beta_0^\theta + \beta^\theta x_{ij} \quad (3.3)$$

For a given quantile θ , quantile regression estimators are computed by minimising over β^θ the following asymmetrically weighted sum of absolute error function:

$$Q(\beta^\theta) = \sum_{i:Fee_i \geq \beta^\theta x_i} \theta |y_{ij} - (\beta_0^\theta + x_i \beta^\theta)| + \sum_{i:Fee_i < \beta^\theta x_i} (1 - \theta) |Fee_i - (\beta_0^\theta + \beta^\theta x_i)| \quad (3.4)$$

The minimisation problem is solved by using linear programming and the estimators are asymptotically normally distributed under general conditions (Cameron et al., 2009).

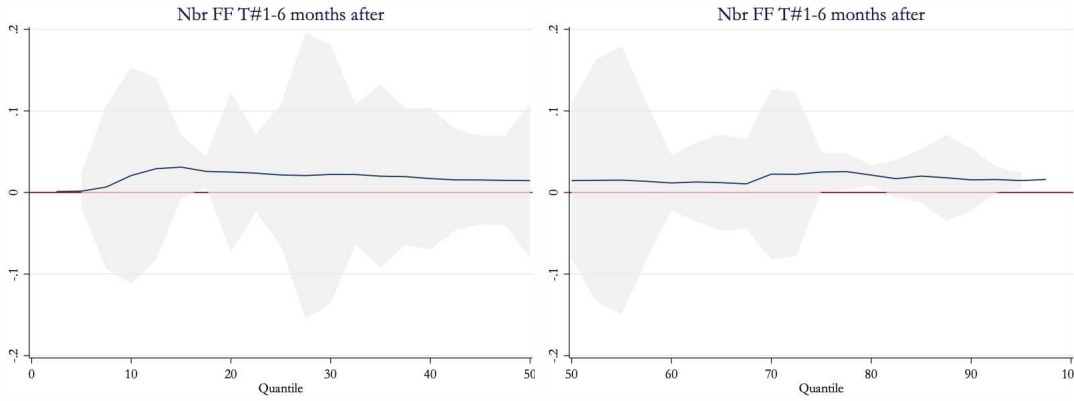


Figure 3.7: Effect of Fake Followers by quantile

Notes: The grey area represent the 95% interval of confidence of the effect of the coefficient associated to the variable $FF * 1-6$ months after.

As reported in section 3.6.1, Fake followers have an effect on transfer fee for the one that happen between 1 and 6 months after the creation of the twitter account. Figure 3.7 shows the effect of the variable of interest $1-6$ months after * FF that result from the quantile regression model (3.4). Table 3.8 report the results of our quantile regression according to the specification of equation 3.4. Table 3.8 shows estimated coefficients at the 10th, 20th, 35th, 50th, 75th, 80th, and 90th percentiles of the transfer fee distribution. First of all

we observe that performance and popularity remain substantially positive and significant along the whole earnings distribution. Our interest is the impact of the Number of Fake Followers, combined with the timing of the transfer on the transfer fees. At the low end of the distribution, the coefficients for *1-6 months after * FF* are insignificant; however, they are positively significant for the 0.75 and 0.8 quantiles. At the 0.75 and 0.8 we observe that an increase in the number of Fake Followers of one thousand when the transfer is between 1 and 6 months after the twitter account creation is associated with an increase in transfer fees ranging from 2.1% to 2.5%.

Table 3.8: Quantile Regression Specification

Quantile	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Transfer Fee (log)	0.1	0.2	0.25	0.5	0.75	0.8	0.9
Online Popularity:							
Twitter Foll.	0.032 (0.061)	0.010 (0.048)	0.024 (0.036)	0.039* (0.023)	0.058** (0.025)	0.064** (0.032)	0.085** (0.037)
Creation Timing x FF:							
Before × FF	-0.022 (0.043)	0.004 (0.021)	-0.001 (0.018)	0.004 (0.015)	-0.005 (0.013)	-0.004 (0.020)	-0.009 (0.013)
0-1 month after × FF	0.044 (0.126)	0.038 (0.291)	0.038 (0.217)	-0.012 (0.493)	-0.040 (0.068)	-0.039 (0.105)	-0.058 (4.641)
1-6 months after × FF	0.021 (0.067)	0.025 (0.050)	0.021 (0.044)	0.014 (0.048)	0.025** (0.012)	0.021*** (0.006)	0.015 (0.019)
6-12 months after × FF	-0.036 (0.273)	-0.023 (0.072)	-0.016 (0.100)	-0.008 (0.033)	0.001 (0.022)	-0.002 (0.077)	-0.018 (0.034)
12+ months after × FF	0.009 (0.021)	0.014 (0.018)	0.011 (0.013)	0.009 (0.008)	0.002 (0.017)	0.002 (0.009)	-0.004 (0.011)
After × FF	-0.009 (0.025)	0.003 (0.021)	0.001 (0.012)	-0.007 (0.005)	-0.002 (0.007)	-0.006 (0.013)	-0.016 (0.021)
Constant	5.187 (5.000)	5.166* (2.838)	4.353* (2.490)	6.695*** (1.405)	8.219*** (1.671)	7.486*** (2.584)	7.476*** (2.472)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Club FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Position FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Performance FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
League FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,032	1,032	1,032	1,032	1,032	1,032	1,032

Robust standard errors in parentheses

* p<.10, ** p<.05, *** p<.01

The results display in Table 3.8 illustrate that around and below the median transactions are not influence by the number of Fake Followers the player have on its Twitter account. It may be because these transaction are the most familiar for people in the soccer industry. Then manipulation occur for others transactions, the more unlikely one. For the really low one, popularity and then online popularity do not matter, these are low quality and

low reputation players. For the really high one, because of the attention given to these transactions and those players, there is no room for manipulation of the online popularity. If there is, it will be detected and so have no effect, or even a negative one.

3.7 Conclusion

Social media metrics can be easily and cheaply manipulated nowadays through, for instance, the acquisition of FFs, views or comments. Individuals who can monetize their popularity, such as professional soccer players, can be interested in making economic gains by unlawfully increasing their online popularity. We have taken advantage of the removal of a large number of fake accounts by Twitter on 12 July 2018, to measure FFs associated with Twitter accounts for a sample of international soccer players. We complement this information with data about players' performance and careers, and have created an original dataset including more than 500,000 observations (game statistics, achievements, transfers, market values, etc.). Empirical explorations provide evidence of specific behaviour of players with FFs. These accounts were mostly created in 2009-2014, when Twitter was more popular than Instagram for soccer players and their fans. Once Instagram becomes more popular, Twitter account creations are no more associated with significant numbers of FFs. It seems that it is no longer worth 'investing' in FFs on Twitter as this is not the reference for players' popularity any more. The strong relationship between the creation of Twitter accounts and transfers – which illustrates the interest about popularity monetisation – is particularly salient for players with significant numbers of FFs. Transfers involving financial bargaining (transfers with fees, loans with fees or free transfers) are indeed over-represented just after the creation of Twitter accounts. This trend is particularly strong for players with significant numbers of FFs, who tend to frequently move to another club shortly after the creation of their Twitter accounts. It means that their Twitter accounts are created close to the bargaining period for these transfers. This evidence suggests that newly created Twitter account can be boosted by the acquisition of FFs to monetize it during negotiations with a potential buying club.

Our results show that FFs on Twitter have an impact on players' valuations, that is, transfer fees, but only for transfers that occur within 1 to 6 months next to the creation of Twitter accounts. These FFs are associated with an average rise of 6% in the transfer fee. The average economic gain associated with this manipulation ($\approx \text{€}650,000$) largely overtakes

the costs of buying thousands of FFs. These findings are consistent with the purposeful manipulation of online popularity to gain economic benefit. The aim of the manipulation is basically to make the buying club overestimate the revenue generated by the player's popularity (attendance, audience, merchandising, etc.). The buying club may thus be prone to offering higher wages and transfer fees to recruit the (fictitiously) popular player. The manipulation positively affects the two streams of revenue associated with a transfer, that is, wages and transfer fees (Carrieri et al., 2018; Garcia del Barrio et al., 2007; Lehmann et al., 2008; Lucifora et al., 2003). In this context, a player, his agent and the selling club thus have an incentive to implement the manipulation to directly benefit from economic gains. All of them can agree to implement this strategy as they will directly retrieve an economic gain, through higher transfer fees and wages. Even if only one actor decides to do so, it will create a positive externality for the others as wages and transfer fees tend to be negotiated at the same time.

This article provides evidence of online popularity manipulation for economic purposes and measures the 'return on investment' of such unlawful behaviour. In this sense, it makes contributions to two existing strands in the literature. Firstly, it provides a new illustration of (online) popularity as a valuable asset in the case of professional athletes. The number of followers indeed has a positive impact on a player's value (Carrieri et al., 2018; Garcia del Barrio et al., 2007; Lehmann et al., 2008; Lucifora et al., 2003). Secondly, we add to the literature on online information manipulation (He et al., 2022; Li et al., 2016; Luca and Zervas, 2016; Mayzlin et al., 2014). We document another kind of online information manipulation and provide some measures of its large positive 'return on investment'. Regarding policy implications, our results support the current trend in social media to hide account metrics (number of followers, views, interactions, etc.). While it may be detrimental for users to detect popular people, this policy may alleviate the incentives underpinning the manipulation of social media metrics. A second implication is related to the platforms that sell fake metrics (followers, reviews, likes, comments, views, etc.). As they have no purpose other than increasing in the best case the noise in these metrics, and in the worst case, to mislead users, the lawfulness of such practice may be questioned by regulators.

This paper has a number of limitations. Firstly, even if our empirical explorations provide evidence of specific behaviours of players with FFs, we do not know exactly when these fake

accounts started to follow soccer players on Twitter. Secondly, our econometric approach does not take into account the possible selection effect among players about whether they are transferred or not during a *mercato* period. Lastly, we have identified manipulation of Twitter followers, but such manipulation can also occur with Instagram followers.

Appendix A

Appendix - The Work/Leisure Trade-off in Online Labor Markets: The Case of Amazon Mechanical Turk

A.1 Supplementary Appendix A: Amazon Mechanical Turk operating

A.1.1 Amazon Mechanical Turk operation

Employees have access to a dashboard that displays three sections: “total earnings”, “HIT status” and “HIT totals”. The “total earnings” displays the total earnings a worker has received from the realisation of Human Intelligence Tasks, the gains made from bonuses and the sum of these two. Then, the “HIT status” displays a list of daily activities and the daily revenue, along with the number of visits that were submitted, approved, rejected or waiting for the given day. And finally, the “HIT totals” displays information about the HITs which have been accepted or are in process (including the percentage of successes that occurred, were returned or abandoned and the percentage of jobs that were approved, rejected or pending those presented).

Employers (companies, researchers or independent developers that need jobs performed) can use the Amazon Mechanical Turk API to programmatically integrate the results of that work directly into their business processes and systems. When employers set up their job, they must specify how much they are paying for each HIT accomplished, how many workers they want to work on each HIT, the maximum time a worker has to work on a single task and how much time the workers have to complete the work, as well as the specific details about the job they want to be completed.

Tasks on AMT can be characterized in four large categories, “processing photos and videos”, “data cleaning or verification of data”, “information collection” and “data processing”. We describe these more precisely in the Table A.1.

Table A.1: Description of tasks available on Amazon Mechanical Turk

Task	Description
<i>Processing photos and videos</i>	Amazon Mechanical Turk provides a platform for processing images, a task well suited to human intelligence. Requesters have created tasks asking workers to label objects found in an image, select the most relevant picture in a group of pictures, screen inappropriate content, and classify objects in satellite images. Also, crowd workers have completed tasks of digitizing text from images such as scanned forms filled out by hand.
<i>Data cleaning or verification of data</i>	Companies with large online catalogs use Mechanical Turk to identify duplicates and verify details of item entries. Some examples of fixing duplicates are identifying and removing duplicates in Yellow Pages directory listings and online product catalog entries. Examples of verifying details include checking restaurant details (phone number and hours) and finding contact information from web pages (author name and email).
<i>Information collection</i>	Diversification and scale of personnel of Mechanical Turk allow collecting an amount of information that would be difficult outside of a crowd platform. Mechanical Turk allows requesters to amass a large number of responses to various types of surveys, from basic demographics to academic research. Other uses include writing comments, descriptions and blog entries to websites and searching data elements or specific fields in large government and legal documents.
<i>Data processing</i>	Companies use Mechanical Turk’s crowd labor to understand and respond to different types of data. Common uses include editing and transcription of podcasts, translation, and matching search engine results.

A.2 Supplementary Appendix: Data

A.2.1 Amazon Mechanical Turk

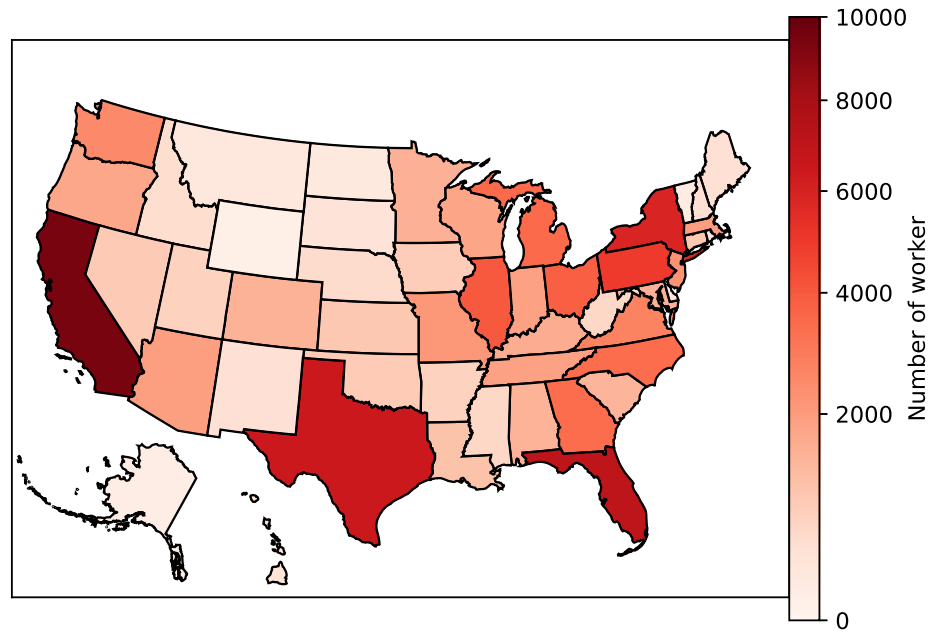


Figure A.1: Distribution of Mechanical Turk workers in the US by state

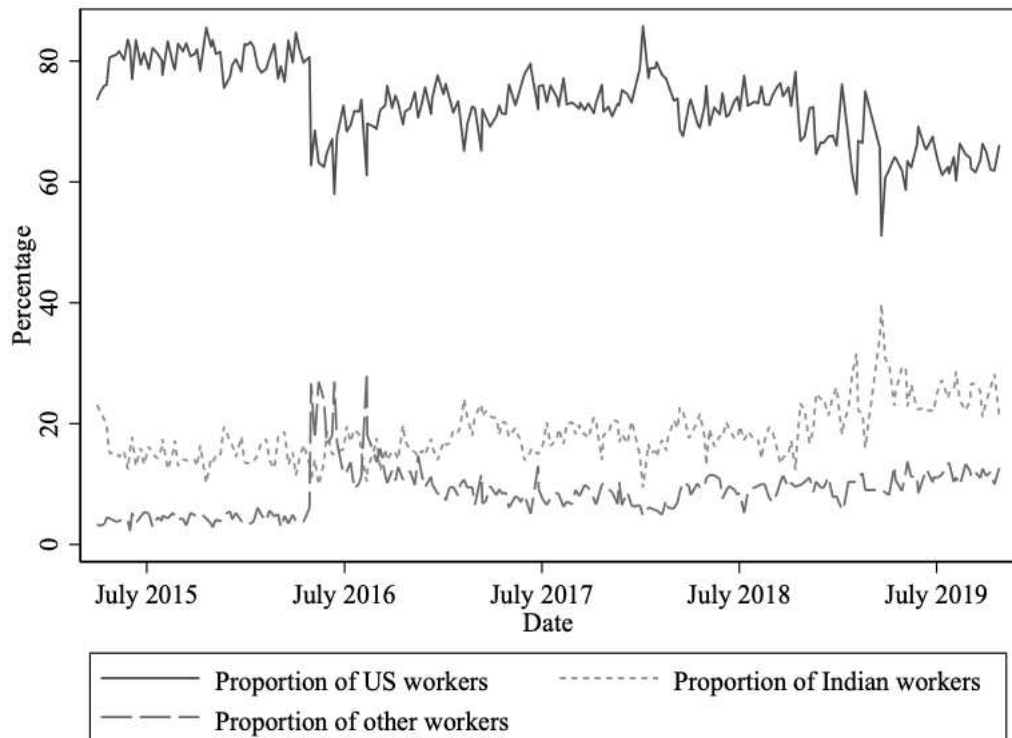


Figure A.2: Proportion of workers by country

Table A.2: Description of categorical variables

	Observations	Proportion
Time spent on AMT per week		
Less than 1 hour	2,518	3%
1-2 hours	7,080	9%
2-4 hours	11,814	15%
4-8 hours	18,409	23%
8-20 hours	21,262	27%
20-40 hours	12,503	15%
More than 40 hours	5,914	7%
Total	79,500	100%
Amount earned on AMT per week		
Less than \$1	1,961	3%
\$1-\$5	8,232	11%
\$5-\$10	11,148	14%
\$10-\$20	13,216	17%
\$20-\$50	18,288	23%
\$50-\$100	14,059	17%
\$100-\$200	8,600	11%
\$200-\$500	3,396	4%
More than \$500	600	1%
Total	79,500	100%
Education level		
Some High School	756	1%
High School Graduate	7,142	9%
Some college, no degree	16,636	22%
Associates degree	7,655	10%
Bachelors degree	34,393	42%
Graduate degree, Masters	11,433	14%
Graduate degree, Doctorate	1,485	2%
Total	79,500	100%

Table A.3: Summary statistics: Demographics at an individual level, 96 observations per day

Variable	Mean	Std. Dev.	Min.	Max.
Overall				
United-States	0.729	-	0	1
India	0.184	-	0	1
Others	0.088	-	0	1
Age	34.582	11.14	7	109
Male	0.513	-	0	1
Household income (\$)	46,624.016	30,878.546	5,000	100,000
Household size	2.986	1.521	1	6
Observations	157,951			
United States				
Age	35.598	11.369	9	109
Male	0.462	-	0	1
Household income (\$)	53,148.046	29,699.418	5,000	100,000
Household size	2.829	1.511	1	6
Observations	115,096			
India				
Age	31.682	10.219	7	109
Male	0.643	-	0	1
Household income (\$)	23,680.213	23,361.19	5,000	100,000
Household size	3.629	1.407	1	6
Observations	28,999			

A.2.2 Pokémon Go

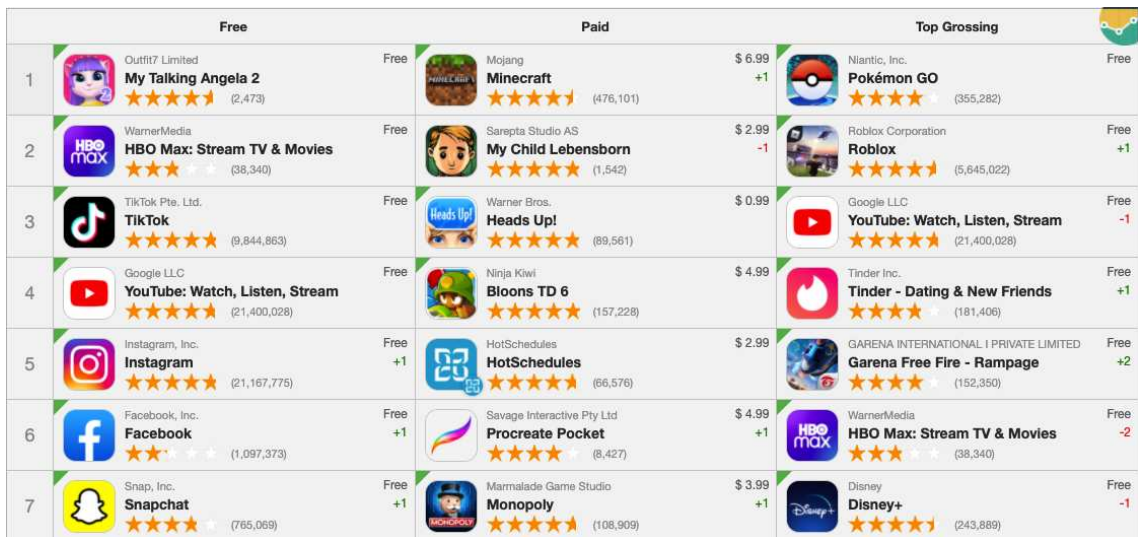


Figure A.3: Top iPhone apps in the US in July 2021

Source : App Store.

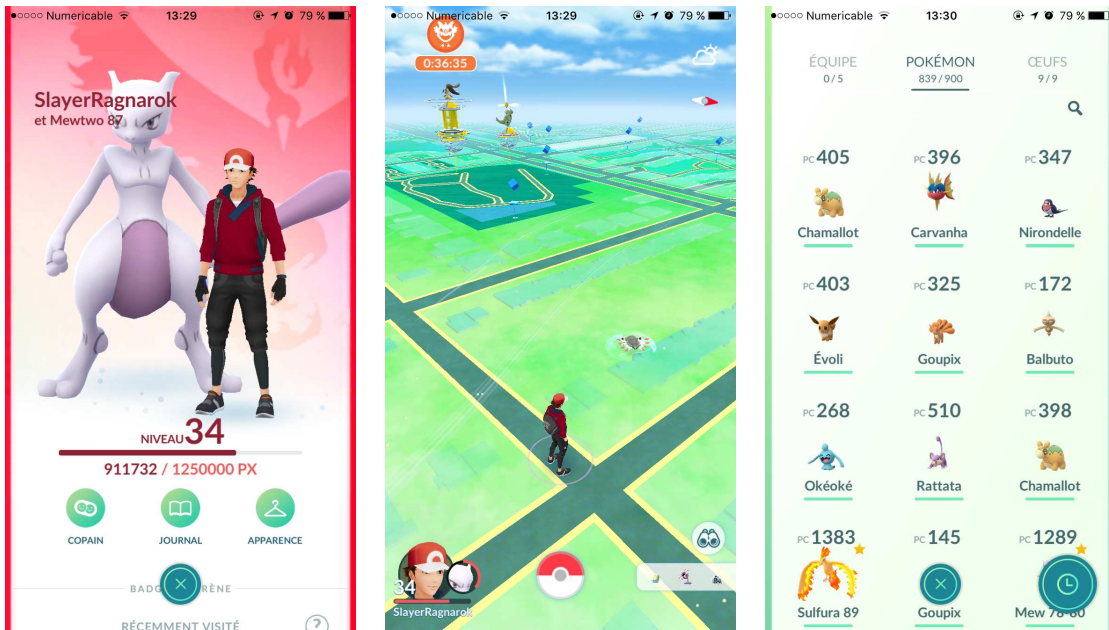


Figure A.4: Overview of the Pokémon Go user interface

Notes: The screenshot on the left shows the player's information. The middle screen shows the real-world map that the player must move around to play. The screenshot on the right shows the Pokémon captured by the player.

A.3 Supplementary Appendix: Robustness Checks & IV test

A.3.1 Placebo Test

We consider the potential for false significance in our estimates as a result of spurious relationships or serial correlations in our dependent variable. One of the major concerns is the validity of the leisure activity index, *PokeGo*. To confirm that the effect of the leisure activity measured by the variable is well defined and not random we perform a placebo test, and more precisely a permutation test (Abadie et al., 2010). The permutation test procedure involves estimating the distribution of “placebo treatment effects”, the dispersion of which provides evidence on the statistical uncertainty underlying the point estimates of interest. The 95% confidence interval implied by the permutation test is simply the range of placebo treatment estimates such that 2.5% of the estimates fall both above and below the interval. As expected, the results are not significant (in 95% of cases). We use a standard random number generator to create a hundred variables, but we only keep the first six variables (*Placebo Test 1 to 6*) in order to illustrate the test. These variables take one value per day and are in the same range as the Pokémon Go variable (0 to 100). Under this placebo test, since these variables do not carry any real information about the usage of an application the corresponding coefficient should not identify any effect in the placebo models and should show that there is no impact (in 95% of cases). Table A.4 presents the set of regressions, and shows that these placebo test variables have no effect. And so, it confirms that the Pokémon Go variables and the part-time job search variables are effective and measure a real interest and their effect is not due to randomness.

Table A.4: Main specification with placebo explicative variables

PropWorker	(1) All	(2) All	(3) All	(4) All	(5) All	(6) All
Placebo Test 1	0.003 (0.006)					
Placebo Test 2		0.005 (0.006)				
Placebo Test 3			0.005 (0.006)			
Placebo Test 4				-0.000 (0.006)		
Placebo Test 5					0.001 (0.006)	
Placebo Test 6						0.000 (0.006)
Constant	92.261*** (1.290)	92.147*** (1.289)	92.211*** (1.265)	92.426*** (1.266)	92.359*** (1.254)	92.408*** (1.271)
Day of the week FE	✓	✓	✓	✓	✓	✓
Week FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Observations	1655	1655	1655	1655	1655	1655
R-squared	0.419	0.419	0.419	0.419	0.419	0.419

Note: OLS with time fixed effects (FE) estimations. *PropWorker* is the dependent variable in all columns. Robust standard errors reported in parentheses.

Significance levels: * $p < .10$, ** $p < .05$, *** $p < .01$

A.3.2 IV Test

Table A.5: IV test

	Durbin	Endogeneity		Over-identification		
		Wu-Hausman	C-stat	Basmann	Sargan	Hansen J-stat
PropWorker	.0015	.0019	.0014	.3258	.3161	.3554
PropWorker Man	.1574	.1659	.1323	-	-	-
PropWorker Woman	.0339	.0376	.0471	.1008	.0943	.1099
PropWorker Over 35	.8511	.8541	.9197	-	-	-
PropWorker Under 35	.0021	.0024	.0039	.0035	.0030	.0053
PropWorker Man Over 35	.6553	.6619	.6170	-	-	-
PropWorker Woman Over 35	.5715	.5793	.5117	-	-	-
PropWorker Man Under 35	.0580	.0631	.0632	.0001	.0001	.0000
PropWorker Woman Under 35	.0308	.03426	.0592	.7551	.7502	.7773
PropWorker Large cities	.0172	.0195	.0105	.9008	.8987	.8921
PropWorker Small cities	.1982	.2073	.2376	-	-	-
First Stage				2,122.221		
First Stage GMM				85.564		


Notes: Tests about endogeneity indicate if an IV strategy is needed, if so a test for good identification of the IV is carried out. The test for over-identification is rejected except for the dependant variable PropWorker Man Under 35.

Appendix B

Appendix - Bad Nudge, Kids and Voice Assistants: A Social Preferences lab-in-the-field Experiment

B.1 Figure

B.1.1 Interlocutor



(a) Pepper Robot



(b) Smart Speaker

Figure B.1: Robot and Smart Speaker

B.1.2 Distribution

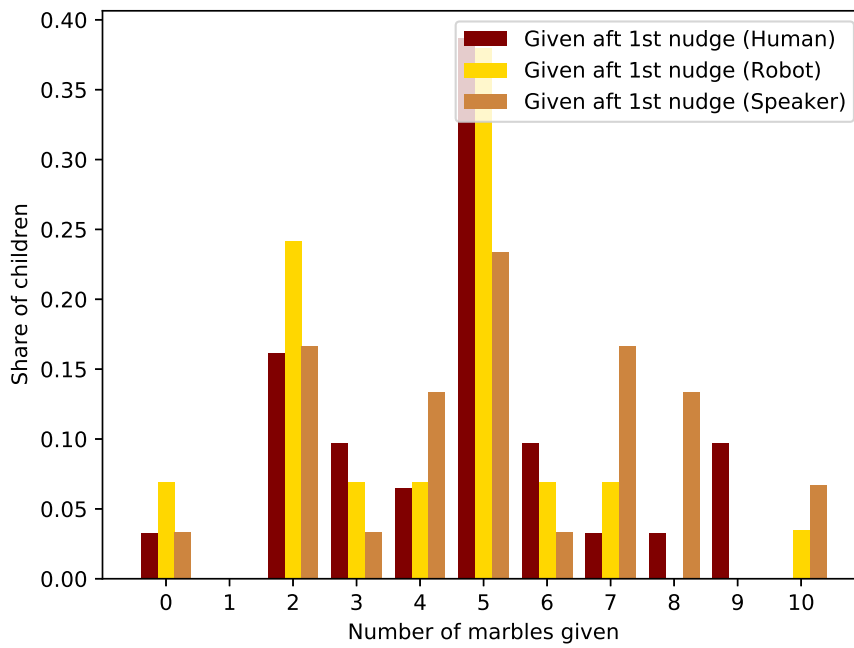


Figure B.2: Distribution of marbles given after the first nudge by treatment




Figure B.3: Distribution of marbles given after the second nudge by treatment

B.2 Statistics on real participants : by treatment

Table B.1: Summary statistics Human

Variable	Mean	Std. Dev.	Min.	Max.	N
Age	8.742	1.692	6	11	31
Male	0.355	-	0	1	31
Nudge Order	0.452	-	0	1	31
More Altruist 1	0.161	-	0	1	31
More Altruist 2	0.323	-	0	1	31
Given	4.839	1.655	0	9	31
Given after first nudge	4.742	2.206	0	9	31
Given after second nudge	4.935	2.19	0	9	31

Table B.2: Summary statistics Robot

Variable	Mean	Std. Dev.	Min.	Max.	N
Age	8.724	1.688	6	11	29
Male	0.414	-	0	1	29
Nudge Order	0.552	-	0	1	29
More Altruist 1	0.207	-	0	1	29
More Altruist 2	0.345	-	0	1	29
Given	4.103	2.093	0	7	29
Given after first nudge	4.103	2.193	0	10	29
Given after second nudge	3.793	2.274	0	8	29

Table B.3: Summary statistics Smart Speaker

Variable	Mean	Std. Dev.	Min.	Max.	N
Age	8.548	1.729	6	11	31
Male	0.452	-	0	1	31
Nudge Order	0.3	-	0	1	30
More Altruist 1	0.267	-	0	1	30
More Altruist 2	0.31	-	0	1	29
Given	4.774	2.68	0	10	31
Given after first nudge	5.233	2.515	0	10	30
Given after second nudge	4.517	2.627	0	10	29

B.3 Possible Selection Bias

Table B.4: Ttest: Participants vs. Non Participants

	Mean Non Participants	Mean Participants	Diff.
Age	8.30	8.67	-0.37
Male	0.45	0.40	0.05
Observations	83	91	

As we had data about all the child of the school, even if they weren't participate to the experiment, we can test if they are some selection bias that with compromise the external validity (out of the group of subject to the entire school). We find no statistically significant differences for the characteristics we had (age and gender).

B.4 Speech of the Interlocutor

B.4.1 French

”Bonjour J’aimerais jouer avec toi. On va faire un jeu! Tu peux prendre des billes, et on te donnera les billes a la fin de la journée.

Voici 2 bols, un bol pour toi et un bol pour quelqu’un d’autre.

Compte a voix haute chaque bille que tu déplaces. Tu peux prendre au maximum 10 billes. Tu peux les mettre dans ton bol, et elles seront pour toi. Tu peux les mettre dans l’autre bol, et on les donnera a quelqu’un d’autre de l’école.

Est-ce que tu as compris les règles ?

Combien de billes peux-tu déplacer ?

A qui vont les billes, que tu mets dans le bol de droite ?

Tu as pris x billes.

Peer effect:

- Si ≤ 3 - D’habitude les enfants prennent 8 billes. Si tu veux, tu peux modifier ton choix.
- Si ≥ 4 - D’habitude les enfants prennent 2 billes. Si tu veux, tu peux modifier ton choix.

Social Proximity:

- Si ≤ 3 - Moi, si on me demande, je prendrais 8 billes. Si tu veux, tu peux modifier ton choix.
- Si ≥ 4 - Moi, si on me demande, je prendrais 2 billes. Si tu veux, tu peux modifier ton choix.

Le jeu des billes est fini. On te donnera tes billes cet après-midi, a la fin de la classe. Merci d’avoir jouez avec moi.”

B.4.2 English

”Hello I would like to play with you. Let’s play a game! You can take some marbles, and we’ll give you the marbles at the end of the day.

Here are two bowls, one for you and one for someone else.

Count out loud each marble you move. You can take a maximum of 10 marbles. You can put them in your bowl, and they are yours. You can put them in the other bowl, and they will be given to someone else in the school.

Do you understand the rules?

How many marbles can you move?

Who do the marbles go to, that you put in the bowl on the right?

- You took x marbles.

according to the child’s choice :

Peer effect:

- If ≤ 3 - Children usually take 8 marbles. If you want, you can change your choice.
- If ≥ 4 - Children usually take 2 marbles. If you want, you can change your choice.

Social Proximity:

- If ≤ 3 - I, if asked, would take 8 marbles. If you want, you can change your choice.
- If ≥ 4 - I, if asked, would take 2 marbles. If you want, you can change your choice.

The marble game is over. You will be given your marbles this afternoon, at the end of the class. Thank you for playing with me.”

Appendix C

Appendix - Online Popularity, Fake Followers and Soccer Players' Value

C.1 Sample descriptive statistics

Table C.1: Summary statistics of players

	Full sample	Twitter
Age (mean/sd/min/max)	26.9 (4.0/17/45)	27.1 (3.9/18/45)
Position:		
Goalkeeper	12.9	12.3
Defender	33.6	33.9
Midfielder	27.3	25.8
Forward	26.2	28.0
Nationality:		
Africa	18.6	17.9
Asia-Pacific	14.3	7.7
Central/North America	8.7	12.9
Eastern Europe	9.7	5.4
Middle East	4.8	3.6
South-America	13.4	15.7
Western Europe	30.6	32.3
Club:		
Africa	4.0	2.4
Asia-Pacific	7.8	3.9
Central/North America	5.9	7.7
Eastern Europe	8.2	3.3
Middle East	7.8	7.2
South-America	5.3	6.0
Western Europe	61.0	69.6
Total	1,075	664

C.2 Transfer: years and players' ages

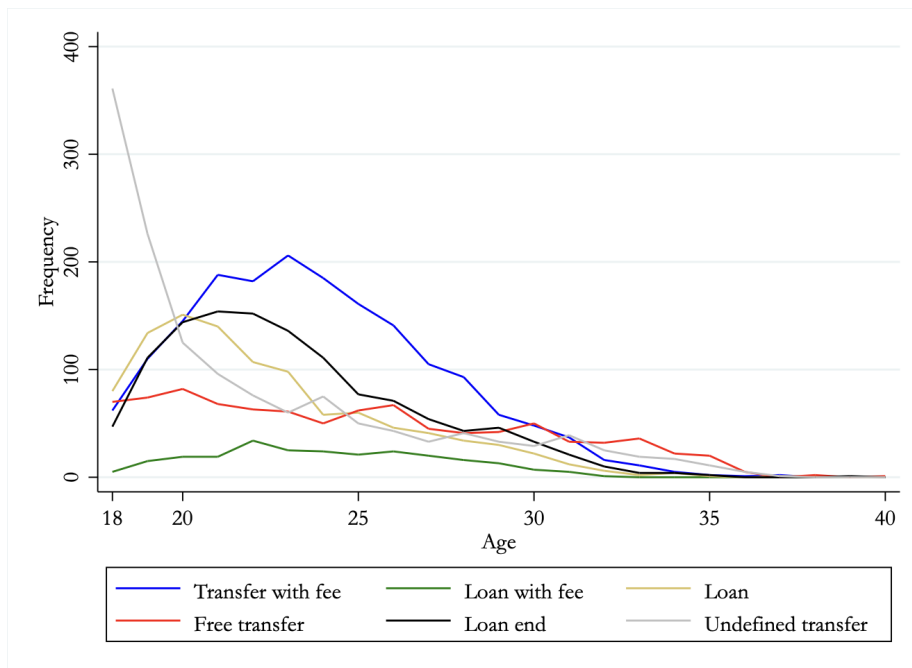


Figure C.1: Transfer types and players' age

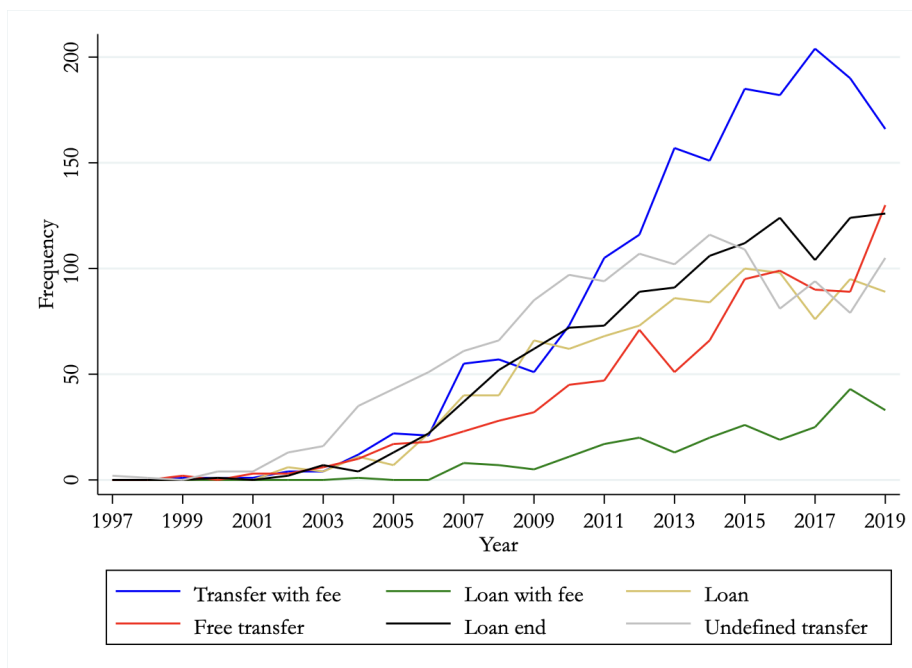


Figure C.2: Transfer types by year

C.3 'Real' and fake followers

Ventile	Corr.	Fake followers		Followers	
		Min.	Max	Min.	Max.
1	.01	0	1	11	1,175
2	.28	0	99	1,178	2,967
3	.02	0	172	3,026	4,413
4	.04	0	188	4,489	7,587
5	.40	0	351	7,656	11,087
6	-.15	0	2,591	11,157	15,418
7	.36	0	514	15,922	23,154
8	-.07	0	1,300	23,167	34,379
9	-.16	0	1,092	34,482	54,128
10	-.29	0	2,709	54,130	77,685
11	.10	0	1,553	78,022	103,031
12	.14	0	4,992	103,091	154,131
13	.24	0	7,610	155,085	217,129
14	-.01	0	12,396	220,446	310,604
15	-.07	0	13,990	311,016	520,065
16	.36	0	32,355	521,627	778,131
17	-.07	0	28,539	794,583	1,223,102
18	.22	0	35,878	1,241,758	1,646,341
19	.26	114	62,855	1,659,446	3,198,009
20	.97	19,635	1,215,302	3,208,290	73,326,669

Table C.2: Correlation between followers and FFs by ventile

C.4 Account creation date and number of followers on Twitter and Instagram

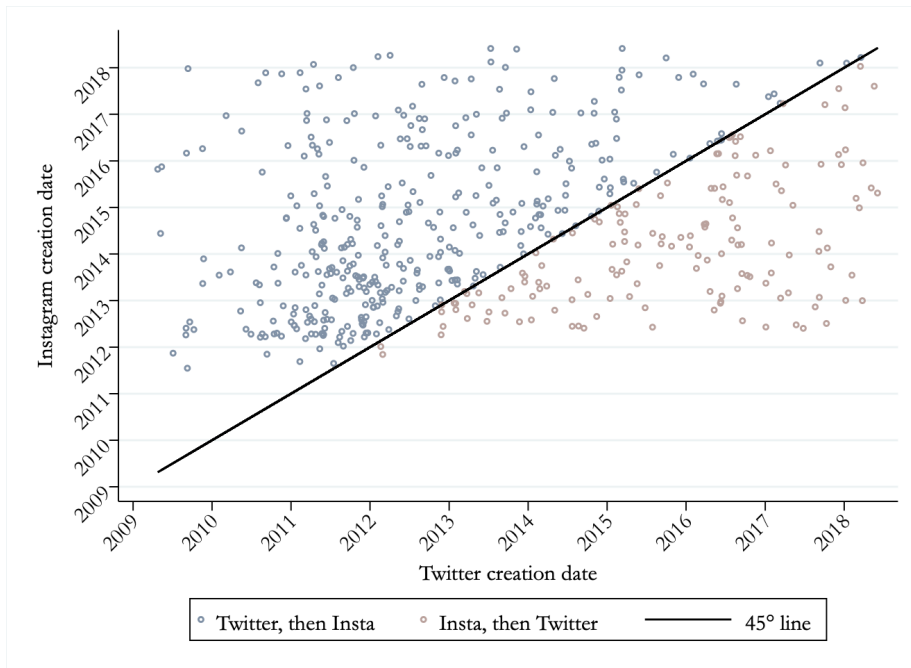


Figure C.3: Twitter and Instagram account creation dates

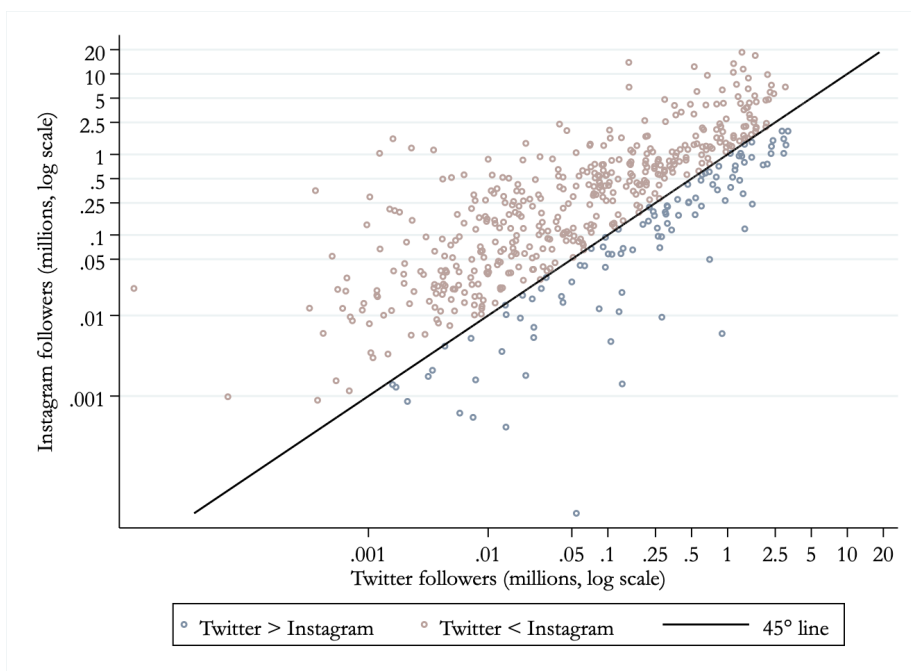


Figure C.4: Number of followers on Twitter and Instagram (July 2018)

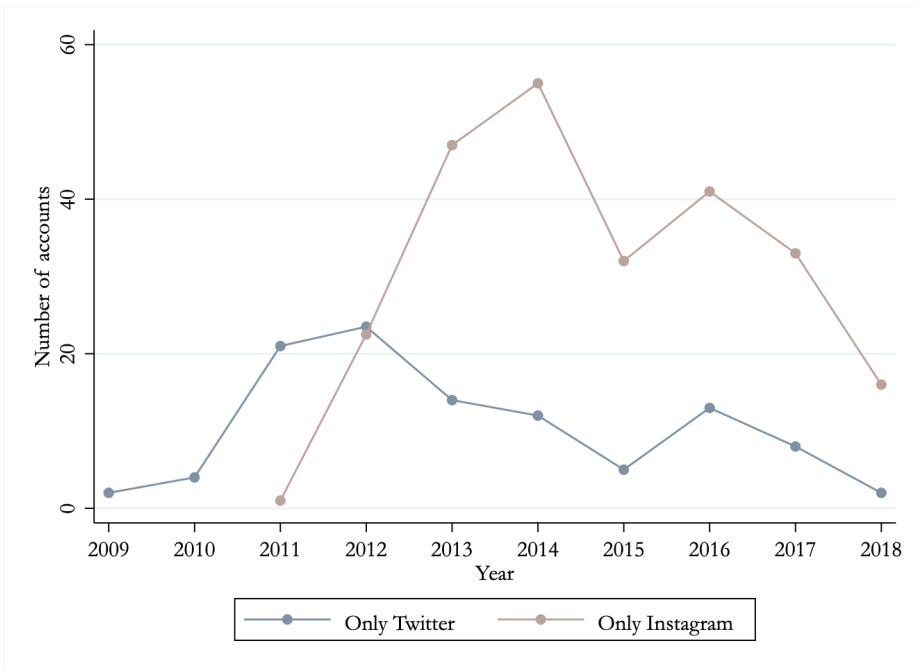


Figure C.5: Only Twitter or Instagram account creation dates

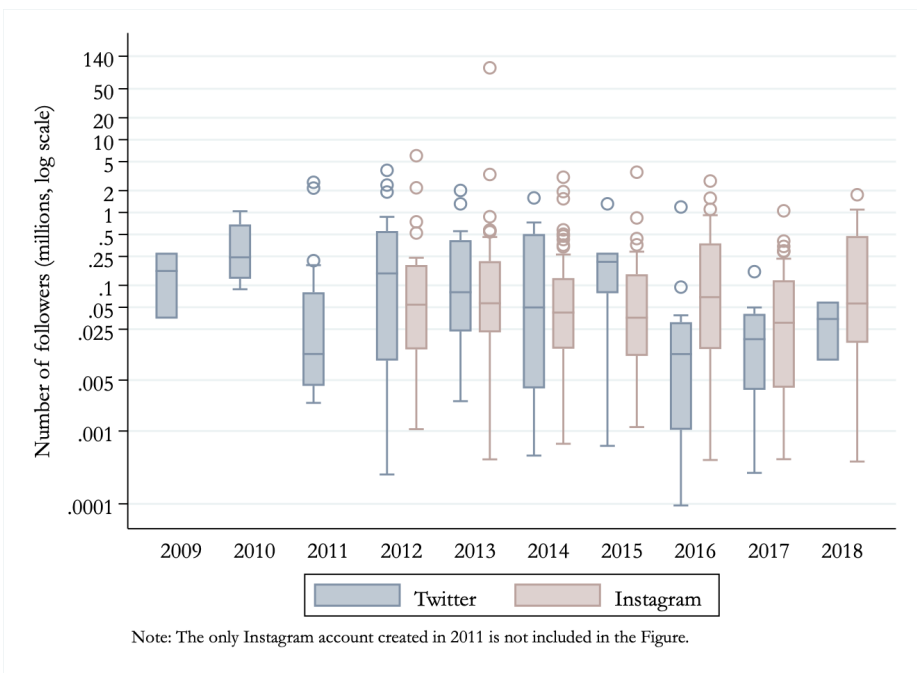


Figure C.6: Number of followers for only Twitter or Instagram (July, 2018)

C.5 Twitter account creation and transfers

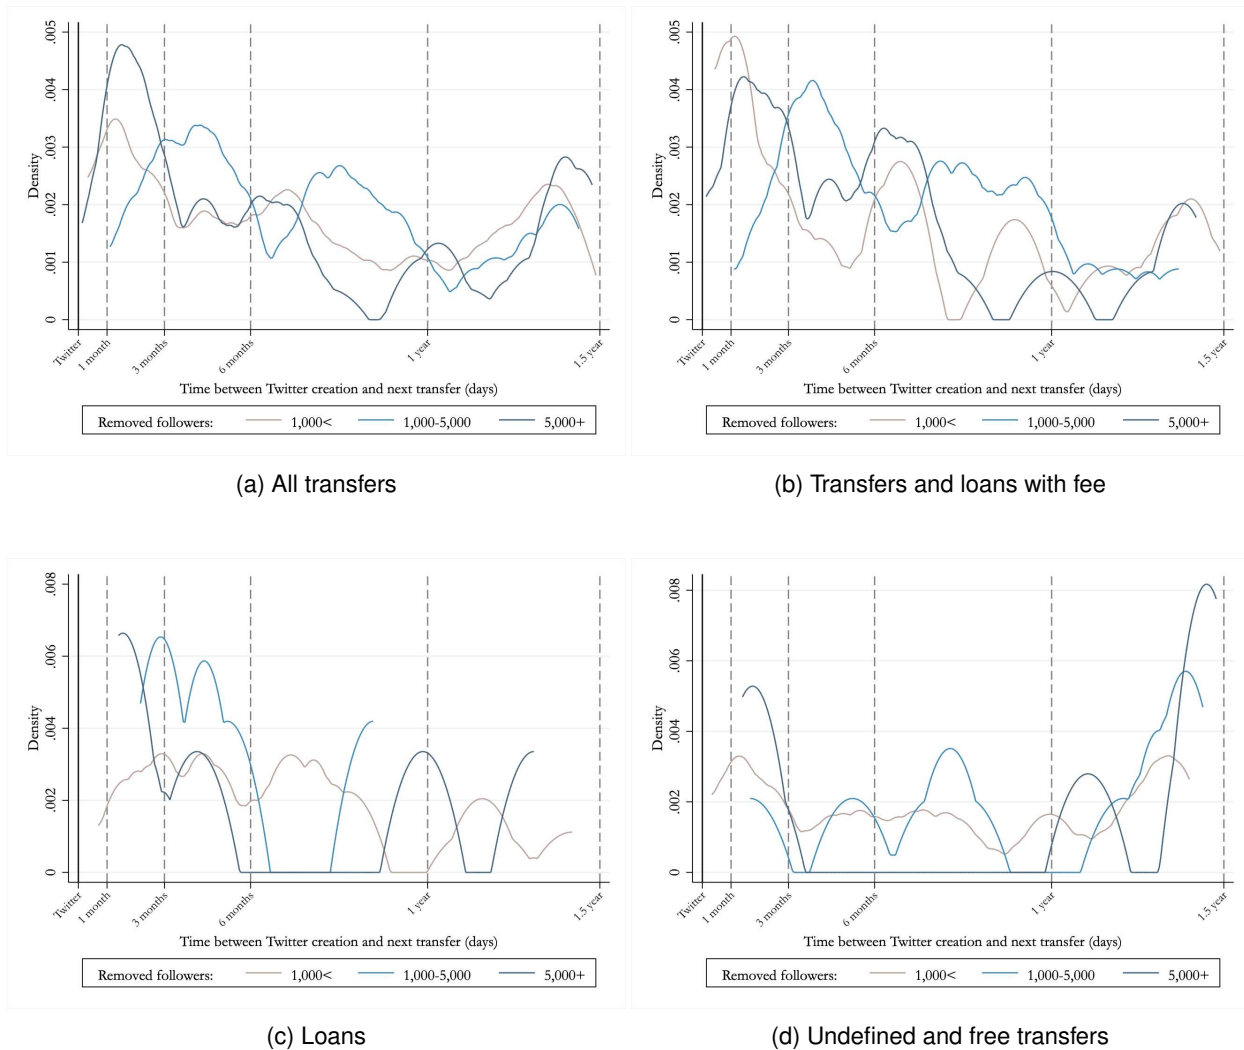


Figure C.7: Transfers following Twitter account creation (2009-2014)

C.6 Results

Table C.3: Main specification: control variables results

Dependant variable: Transfer Fee (log)	2009-2019		2009-2014	
	Twitter (1)	Insta-Twitter (2)	Twitter (3)	Insta-Twitter (4)
Age	0.402*** (0.138)	0.414*** (0.145)	0.175 (0.174)	0.148 (0.192)
Age sq	-0.008*** (0.003)	-0.009*** (0.003)	-0.004 (0.003)	-0.003 (0.004)
Defender	0.016 (0.152)	0.001 (0.150)	-0.136 (0.276)	-0.167 (0.272)
Midfielder	0.149 (0.174)	0.114 (0.172)	-0.091 (0.316)	-0.118 (0.312)
Forward	0.137 (0.202)	0.092 (0.201)	-0.225 (0.393)	-0.281 (0.392)
Loan with Fee	-1.489*** (0.098)	-1.499*** (0.098)	-1.314*** (0.218)	-1.331*** (0.217)
Africa	-0.547 (0.368)	-0.510 (0.375)	-0.909* (0.545)	-0.898* (0.540)
North America	-0.198 (0.224)	-0.212 (0.219)	0.106 (0.632)	0.113 (0.609)
Asia	0.163 (0.261)	0.184 (0.253)	-0.099 (0.482)	-0.083 (0.482)
Middle East	0.063 (0.208)	0.043 (0.208)	0.241 (0.344)	0.202 (0.346)
East Europe	-0.450* (0.232)	-0.402* (0.224)	-0.221 (0.333)	-0.174 (0.329)
West Europe	-0.648*** (0.139)	-0.640*** (0.140)	-0.537*** (0.203)	-0.516*** (0.205)
France	0.142 (0.189)	0.143 (0.181)	0.172 (0.336)	0.189 (0.321)
Germany	-0.634*** (0.222)	-0.598*** (0.224)	-0.581* (0.333)	-0.522 (0.338)
Italy	0.014 (0.170)	-0.009 (0.169)	0.266 (0.438)	0.262 (0.436)
Spain	0.187 (0.191)	0.227 (0.190)	0.784** (0.333)	0.813** (0.339)
England	0.312** (0.157)	0.327** (0.155)	0.553* (0.292)	0.576* (0.295)
Trophies Club	0.075* (0.043)	0.070* (0.042)	-0.093 (0.088)	-0.094 (0.089)
Trophies Nat	-0.074 (0.159)	-0.068 (0.153)	-0.056 (0.277)	-0.036 (0.282)
Trophies Young	0.213 (0.236)	0.211 (0.232)	0.401 (0.284)	0.446 (0.295)

Continued next page

Value Managed	0.051**	0.042*	0.043	0.035
	(0.022)	(0.022)	(0.039)	(0.041)
On	0.017	0.013	0.028	0.027
	(0.051)	(0.050)	(0.098)	(0.097)
Justice	-0.435	-0.424	-0.447	-0.419
	(0.267)	(0.279)	(0.383)	(0.382)
Goals	0.013**	0.013**	0.016	0.015
	(0.006)	(0.006)	(0.012)	(0.011)
Assists	0.008	0.007	0.027*	0.027*
	(0.010)	(0.009)	(0.016)	(0.016)
Conceded	0.002	0.002	0.002	0.002
	(0.002)	(0.002)	(0.004)	(0.004)
Captain	-0.001	-0.001	-0.000	-0.002
	(0.005)	(0.004)	(0.013)	(0.012)
Yellow	-0.009	-0.009	-0.008	-0.006
	(0.010)	(0.009)	(0.018)	(0.018)
Red	0.085	0.082	-0.040	-0.040
	(0.071)	(0.070)	(0.122)	(0.122)
Club Nat	0.297***	0.304***	0.433***	0.436***
	(0.092)	(0.092)	(0.147)	(0.148)
Club Inter	0.085**	0.078**	0.088	0.083
	(0.034)	(0.034)	(0.067)	(0.066)
Nat	0.151***	0.149***	0.210***	0.219***
	(0.043)	(0.044)	(0.075)	(0.075)
Constant	6.997***	6.508***	8.975***	8.958***
	(1.691)	(1.808)	(2.266)	(2.416)
Time FE	Yes	Yes	Yes	Yes
Club FE	Yes	Yes	Yes	Yes
Obs	1032	1032	422	422
R-squared	0.738	0.741	0.798	0.800

Cluster-robust standard errors (at player level) are shown in parentheses.

* $p < .10$, ** $p < .05$, *** $p < .01$

Table C.3 shows most of the coefficients for the control variables. Performances are computed for the year before the transfer and league indicates the championship where the player goes to. Time FE and Club FE are not reported in this table.

Bibliography

- Abadie, Alberto et al. (2010). "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program". In: *Journal of the American Statistical Association* 105.490, pp. 493–505.
- Adler, Moshe (1985). "Stardom and talent". In: *American Economic Review* 75.1, pp. 208–212.
- (2006). "Stardom and talent". In: *Handbook of the Economics of Art and Culture, Volume 1*. Ed. by Victor A. Ginsburgh et al. Elsevier.
- Aguiar, Mark et al. (2021). "Leisure luxuries and the labor supply of young men". In: *Journal of Political Economy* 129.2, pp. 337–382.
- Akerlof, George A (1997). "Social distance and social decisions". In: *Econometrica: Journal of the Econometric Society*, pp. 1005–1027.
- Althoff, Tim et al. (2016). "Influence of Pokémon Go on Physical Activity: Study and Implications". In: *Journal of Medical Internet Research* 18.12, e315.
- Anand, Abhinav et al. (2022). "Influencer marketing with fake followers". In: *IIM Bangalore Research Paper* 580.
- Atkinson, Anthony B. et al. (2011). "Top incomes in the long run of history". In: *Journal of Economic Literature* 49.1, pp. 3–71.
- Autor, David H (2001). "Why do temporary help firms provide free general skills training?" In: *The Quarterly Journal of Economics* 116.4, pp. 1409–1448.
- Baker, Scott R et al. (2017). "The Impact of Unemployment Insurance on Job Search: Evidence from Google Search Data". In: *Review of Economics and Statistics* 99.5, pp. 756–768.
- Bauer, Michal et al. (2014). "Parental background and other-regarding preferences in children". In: *Experimental Economics* 17.1, pp. 24–46.
- Bebchuk, Lucian Arye et al. (2003). "Executive Compensation as an Agency Problem". In: *Journal of Economic Perspectives* 17.3, pp. 71–92. DOI: 10.1257/089533003769204362. URL: <https://www.aeaweb.org/articles?id=10.1257/089533003769204362>.
- Benenson, Joyce F et al. (2007). "Children's altruistic behavior in the dictator game". In: *Evolution and Human Behavior* 28.3, pp. 168–175.
- Benson, Alan et al. (2020). "Can Reputation Discipline the Gig Economy? Experimental Evidence from an Online Labor Market". In: *Management Science* 66.5, pp. 1802–1825.
- Bernhardt, Dan et al. (2010). "Point Shaving in College Basketball: A Cautionary Tale for Forensic Economics". In: *Economic Inquiry* 48.1, pp. 14–25. DOI: 10.1111/j.1465-7295.2009.00253.x.
- Blum, Bernardo S et al. (2006). "Does the internet defy the law of gravity?" In: *Journal of international economics* 70.2, pp. 384–405.
- Bohnet, Iris et al. (2008). "Betrayal aversion: Evidence from brazil, china, oman, switzerland, turkey, and the united states". In: *American Economic Review* 98.1, pp. 294–310.
- Borchert, Kathrin et al. (2018). "Unemployment and Online Labor". In: *ZEW-Centre for European Economic Research Discussion Paper* 18.23.
- Boxell, Levi et al. (2017a). "Greater Internet use is not associated with faster growth in political polarization among US demographic groups". In: *Proceedings of the National Academy of Sciences* 114.40, pp. 10612–10617.
- (2017b). *Is the internet causing political polarization? Evidence from demographics*. Tech. rep. National Bureau of Economic Research.
- Bruce, Allison et al. (2002). "The role of expressiveness and attention in human-robot interaction". In: *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No. 02CH37292)* 4, pp. 4138–4142.
- Buhrmester, Michael et al. (2011). "Amazon's Mechanical Turk: A New Source of Inexpensive, Yet High-Quality, Data?" In: *Perspectives on Psychological Science* 6.1, pp. 3–5.
- Burch, Gordon et al. (2018). "Can You Gig it? An Empirical Examination of the Gig Economy and Entrepreneurial Activity". In: *Management Science* 64.12, pp. 5497–5520.

- Calvano, Emilio et al. (2020). "Artificial Intelligence, Algorithmic Pricing, and Collusion". In: *American Economic Review* 110.10, pp. 3267–97.
- Camerer, Colin F et al. (1995). "Anomalies: Ultimatums, dictators and manners". In: *Journal of Economic perspectives* 9.2, pp. 209–219.
- Cameron, Colin A. et al. (2009). *Microeconomics using Stata*. College Station, Texas: Stata Press Publication.
- Carmichael, Fiona et al. (1999). "The labour market in association football: Who gets transferred and for how much?" In: *Bulletin of Economic Research* 51.2, pp. 125–150.
- Carrieri, Vincenzo et al. (2018). "What makes you 'super-rich'? New evidence from an analysis of football players' wages". In: *Oxford Economic Paper* 70.4, pp. 950–973. DOI: 10.1093/oep/gpy025. URL: <https://doi.org/10.1093/oep/gpy025>.
- Charness, Gary et al. (2008). "What's in a name? Anonymity and social distance in dictator and ultimatum games". In: *Journal of Economic Behavior & Organization* 68.1, pp. 29–35.
- Chevalier, Judith A et al. (2006). "The effect of word of mouth on sales: Online book reviews". In: *Journal of marketing research* 43.3, pp. 345–354.
- Choi, Hyunyoung et al. (2012). "Predicting the present with Google Trends". In: *Economic record* 88, pp. 2–9.
- Chong, Alberto et al. (2009). "Television and Divorce: Evidence from Brazilian Novelas". In: *Journal of the European Economic Association* 7.2-3, pp. 458–468.
- Comino, Stefano et al. (2019). "Updates Management in Mobile Applications: iTunes versus Google Play". In: *Journal of Economics & Management Strategy* 28.3, pp. 392–419.
- Connolly, Marie et al. (2006). "Rockonomics: The economics of popular music". In: *Handbook of the Economics of Art and Culture* 1, pp. 667–719.
- Cook, Cody et al. (2021). "The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Rideshare Drivers". In: *The Review of Economic Studies* 88.5, pp. 2210–2238.
- DellaVigna, Stefano et al. (2010). "Detecting Illegal Arms Trade". In: *American Economic Journal: Economic Policy* 2.4, pp. 26–57. DOI: 10.1257/po1.2.4.26. URL: <http://www.aeaweb.org/articles?id=10.1257/po1.2.4.26>.
- Deutscher, Christian et al. (2016). "Does Performance Consistency Pay Off Financially for Players? Evidence From the Bundesliga". In: *Journal of Sports Economics* 17.1, pp. 27–43. DOI: 10.1177/1527002514521428.
- Difallah, Djellel et al. (2018). "Demographics and Dynamics of Mechanical Turk Workers". In: *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pp. 135–143.
- Dobson, Stephen et al. (1999). "The determination of player transfer fees in English professional soccer". In: *Journal of Sport Management* 13, pp. 259–279.
- Dube, Arindrajit et al. (2020). "Monopsony in Online Labor Markets". In: *American Economic Review: Insights* 2.1, pp. 33–46.
- Duch-Brown, Néstor et al. (2017). "The impact of online sales on consumers and firms. Evidence from consumer electronics". In: *International Journal of Industrial Organization* 52, pp. 30–62.
- Duggan, Mark et al. (2002). "Winning Isn't Everything: Corruption in Sumo Wrestling". In: *American Economic Review* 92.5. doi: <http://www.aeaweb.org/articles?id=10.1257/000282802762024665>, pp. 1594–1605.
- Ellison, Glenn (2011). "Is peer review in decline?" In: *Economic Inquiry* 49.3, pp. 635–657.
- Ertug, Gokhan et al. (2013). "Getting What You Need: How Reputation and Status Affect Team Performance, Hiring, and Salaries in the NBA". In: *Academy of Management Journal* 56.2, pp. 407–431. DOI: 10.5465/amj.2010.1084. URL: <https://doi.org/10.5465/amj.2010.1084>.
- Faccio, Mara et al. (2020). "Death by Pokémon GO: The Economic and Human Cost of Using Apps while Driving". In: *Journal of Risk and Insurance* 87.3, pp. 815–849.
- Farronato, Chiara et al. (2018). *The welfare effects of peer entry in the accommodation market: The case of airbnb*. Tech. rep. National Bureau of Economic Research.
- Forsythe, Robert et al. (1994). "Fairness in simple bargaining experiments". In: *Games and Economic behavior* 6.3, pp. 347–369.
- Franck, Egon et al. (2012). "Talent and/or popularity: What does it take to be a superstar?" In: *Economic Inquiry* 50.1, pp. 202–216.
- Gans, Joshua S et al. (2021). "Exit, tweets, and loyalty". In: *American Economic Journal: Microeconomics* 13.2, pp. 68–112.
- Garcia del Barrio, Pedro et al. (2007). "Hidden Monopsony Rents in Winner-take-all Markets—Sport and Economic Contribution of Spanish Soccer Players". In: *Managerial and Decision Economics* 28. doi: <https://doi.org/10.1002/mde.1313>, pp. 57–70.
- Gentzkow, Matthew (2006). "Television and Voter Turnout". In: *The Quarterly Journal of Economics* 121.3, pp. 931–972.

- (2007). “Valuing new goods in a model with complementarity: Online newspapers”. In: *American Economic Review* 97.3, pp. 713–744.
- Gentzkow, Matthew and Jesse M Shapiro (2011). “Ideological segregation online and offline”. In: *The Quarterly Journal of Economics* 126.4, pp. 1799–1839.
- (2008). “Preschool Television Viewing and Adolescent Test Scores: Historical Evidence from the Coleman Study”. In: *The Quarterly Journal of Economics* 123.1, pp. 279–323.
- Ghose, Anindya et al. (2012). “Designing Ranking Systems for Hotels on Travel Search Engines by Mining User-Generated and Crowdsourced Content”. In: *Marketing Science* 31.3, pp. 493–520.
- Goldfarb, Avi et al. (2019). “Digital economics”. In: *Journal of Economic Literature* 57.1, pp. 3–43.
- Gomez-Herrera, Estrella et al. (2019). “Is There a Gender Wage Gap in Online Labor Markets? Evidence from Over 250,000 Projects and 2.5 Million Wage Bill Proposals”. In: *Max Planck Institute for Innovation & Competition Research Paper* 19.7.
- Graham, Mark et al. (2017). “Digital Labour and Development: Impacts of Global Digital Labour Platforms and the Gig Economy on Worker Livelihoods”. In: *Transfer: European Review of Labour and Research* 23.2, pp. 135–162.
- Gummerum, Michaela, Yaniv Hanoch, et al. (2010). “Preschoolers’ allocations in the dictator game: The role of moral emotions”. In: *Journal of Economic Psychology* 31.1, pp. 25–34.
- Gummerum, Michaela, Monika Keller, et al. (2008). “To give or not to give: Children’s and adolescents’ sharing and moral negotiations in economic decision situations”. In: *Child Development* 79.3, pp. 562–576.
- Hara, Kotaro et al. (2018). “A Data-Driven Analysis of Workers’ Earnings on Amazon Mechanical Turk”. In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1–14.
- Harbaugh, William T and Kate Krause (2000). “Children’s altruism in public good and dictator experiments”. In: *Economic Inquiry* 38.1, pp. 95–109.
- Harbaugh, William T, Kate Krause, and Timothy R Berry (2001). “GARP for kids: On the development of rational choice behavior”. In: *American Economic Review* 91.5, pp. 1539–1545.
- Harbaugh, William T, Kate Krause, and Lise Vesterlund (2002). “Risk attitudes of children and adults: Choices over small and large probability gains and losses”. In: *Experimental Economics* 5.1, pp. 53–84.
- He, Sherry et al. (2022). “The Market for Fake Reviews”. In: *Marketing Science* forthcoming. doi: <https://doi.org/10.1287/mksc.2022.135>.
- Hill, Kim et al. (2004). “Economic experiments to examine fairness and cooperation among the Ache Indians of Paraguay”. In: *Foundations in Human Sociality: Economic experiments and ethnographic evidence from fifteen small-scale societies*, pp. 382–412.
- Hummel, Dennis et al. (2019). “How Effective Is Nudging? A Quantitative Review on the Effect Sizes and Limits of Empirical Nudging Studies”. In: *Journal of Behavioral and Experimental Economics* 80.
- Jacob, Brian A. et al. (2003). “Rotten Apples: An Investigation of the Prevalence and Predictors of Teacher Cheating”. In: *The Quarterly Journal of Economics* 118.3. doi: <https://doi.org/10.1162/00335530360698441>, pp. 843–877.
- Jetter, Michael et al. (2017). “Good Girl, Bad Boy? Evidence Consistent with Collusion in Professional Tennis”. In: *Southern Economic Journal* 84.1. doi: <https://doi.org/10.1002/soej.12213>, pp. 155–180.
- Kagel, John H et al. (2020). *The handbook of experimental economics, volume 2*. Princeton university press.
- Kearney, Melissa S et al. (2015). “Media Influences on Social Outcomes: The Impact of MTV’s 16 and Pregnant on Teen Childbearing”. In: *American Economic Review* 105.12, pp. 3597–3632.
- Kleinberg, Jon et al. (2018). “Human decisions and machine predictions”. In: *The quarterly journal of economics* 133.1, pp. 237–293.
- Köbis, Nils et al. (2021). “Bad machines corrupt good morals”. In: *Nature Human Behaviour* 5.6, pp. 679–685.
- Kosse, Fabian et al. (2020). “The formation of prosociality: causal evidence on the role of social environment”. In: *Journal of Political Economy* 128.2, pp. 434–467.
- Kroft, Kory et al. (2014). “Does online search crowd out traditional search and improve matching efficiency? Evidence from Craigslist”. In: *Journal of Labor Economics* 32.2, pp. 259–303.
- Krupka, Erin L et al. (2013). “Identifying social norms using coordination games: Why does dictator game sharing vary?” In: *Journal of the European Economic Association* 11.3, pp. 495–524.
- Kuhn, Peter et al. (2014). “Is internet job search still ineffective?” In: *The Economic Journal* 124.581, pp. 1213–1233.
- La Ferrara, Eliana et al. (2012). “Soap Operas and Fertility: Evidence from Brazil”. In: *American Economic Journal: Applied Economics* 4.4, pp. 1–31.

- Lehmann, Erick E. et al. (2008). "What Does it Take to be a Star? The Role of Performance and the Media for German Soccer Players". In: *Applied Economics Quarterly* 54.1, pp. 59–70. DOI: 10.3790/aeq.54.1.59. URL: <http://dx.doi.org/10.3790/aeq.54.1.59>.
- Li, Xing et al. (2016). "Paying Incumbents and Customers to Enter an Industry: Buying Downloads". SSRN Working Paper. URL: <http://dx.doi.org/10.2139/ssrn.2834564>.
- List, John A et al. (2021). *How experiments with children inform economics*. Tech. rep. National Bureau of Economic Research.
- Luca, Michael (2016). "Reviews, reputation, and revenue: The case of Yelp. com". In: *Com (March 15, 2016). Harvard Business School NOM Unit Working Paper* 12-016.
- (2011). "Reviews, reputation, and revenue: The case of Yelp.com". Harvard Business School Working Paper, No. 12-016.
- Luca, Michael and Georgios Zervas (2016). "Fake it till you make it: Reputation, competition, and Yelp review fraud". In: *Management Science* 62.12, pp. 3412–3427. DOI: 10.1287/mnsc.2015.2304. URL: <https://doi.org/10.1287/mnsc.2015.2304>.
- Lucifora, Claudio et al. (2003). "Superstar Effects in Sport: Evidence From Italian Soccer". In: *Journal of Sports Economics* 4.1, pp. 35–55. DOI: 10.1177/1527002502239657. URL: <http://dx.doi.org/10.1177/1527002502239657>.
- Luo, Xueming et al. (2019). "Frontiers: Machines vs. humans: The impact of artificial intelligence chatbot disclosure on customer purchases". In: *Marketing Science* 38.6, pp. 937–947.
- Ma, Shuzhong et al. (2021). "The Effect of Online Search on International Trade". In: *Applied Economics*, pp. 1–16.
- Malmendier, Ulrike et al. (2009). "Superstar CEOs". In: *Quarterly Journal of Economics* 124.4, pp. 1593–1638. DOI: 10.1162/qjec.2009.124.4.1593.
- Malti, Tina et al. (2009). "Children's moral motivation, sympathy, and prosocial behavior". In: *Child development* 80.2, pp. 442–460.
- Mayzlin, D. et al. (2014). "Promotional Reviews: An Empirical Investigation of Online Review Manipulation". In: *American Economic Review* 104.8, pp. 2421–2455. DOI: 10.1257/aer.104.8.2421. URL: <https://doi.org/10.1257/aer.104.8.2421>.
- Müller, Oliver et al. (2017). "Beyond crowd judgments: Data-driven estimation of market value in association football". In: *European Journal of Operational Research* 263, pp. 611–624. DOI: 10.1257/aer.104.8.2421. URL: <https://10.1016/j.ejor.2017.05.005>.
- Murnighan, J Keith et al. (1998). "Ultimatum bargaining by children and adults". In: *Journal of Economic Psychology* 19.4, pp. 415–445.
- Paolacci, Gabriele et al. (2010). "Running Experiments on Amazon Mechanical Turk". In: *Judgment and Decision Making* 5.5, pp. 411–419.
- Peer, Eyal, Serge Egelman, et al. (2020). "Nudge me right: Personalizing online security nudges to people's decision-making styles". In: *Computers in Human Behavior* 109, p. 106347.
- Peer, Eyal, Joachim Vosgerau, et al. (2014). "Reputation as a Sufficient Condition for Data quality on Amazon Mechanical Turk". In: *Behavior Research Methods* 46.4, pp. 1023–1031.
- Poli, Raffaele, Roger Besson, et al. (2022). "Econometric Approach to Assessing the Transfer Fees and Values of Professional Football Players". In: *Economies* 10.1, p. 4.
- Poli, Raffaele, Loïc Ravenel, et al. (2019). *Financial analysis of the transfer market in the big-5 European leagues (2010-2019)*. CIES Football Observatory Monthly Report n°47. URL: <https://www.football-observatory.com/IMG/sites/mr/mr47/en/>.
- Quansah, Tommy et al. (2021). "The Importance of Club Revenues for Player Salaries and Transfer Expenses—How Does the Coronavirus Outbreak (COVID-19) Impact the English Premier League?" In: *Sustainability* 13.9, p. 5154.
- Rahwan, Iyad et al. (2019). "Machine behaviour". In: *Nature* 568.7753, pp. 477–486.
- Rosen, Sherwin (1981). "The economics of superstars". In: *American Economic Review* 71.5, pp. 845–858.
- Roth, Aaron M et al. (2019). "A robot's expressive language affects human strategy and perceptions in a competitive game". In: *2019 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, pp. 1–8.
- Sari, Florent (2019). "Spatial Disparities in Accessibility to Recreational Amenities: the Case of Pokémon GO". In: *The Annals of Regional Science*, pp. 1–31.
- Schanke, Scott et al. (2021). "Estimating the impact of "humanizing" customer service chatbots". In: *Information Systems Research* 32.3, pp. 736–751.
- Schilbach, Frank (2019). "Alcohol and Self-Control: A Field Experiment in India". In: *American Economic Review* 109.4, pp. 1290–1322.
- Seamans, Robert et al. (2014). "Responses to entry in multi-sided markets: The impact of Craigslist on local newspapers". In: *Management Science* 60.2, pp. 476–493.

- Shirado, Hirokazu et al. (2017). "Locally noisy autonomous agents improve global human coordination in network experiments". In: *Nature* 545.7654, pp. 370–374.
- Stanton, Christopher T et al. (2016). "Landing the first job: The value of intermediaries in online hiring". In: *The Review of Economic Studies* 83.2, pp. 810–854.
- Sunstein, Cass R (2020). "Sludge audits". In: *Behavioural Public Policy*, pp. 1–20.
- Sutter, Matthias (2007). "Outcomes versus intentions: On the nature of fair behavior and its development with age". In: *Journal of Economic Psychology* 28.1, pp. 69–78.
- Sutter, Matthias et al. (2019). "Economic behavior of children and adolescents—A first survey of experimental economics results". In: *European Economic Review* 111, pp. 98–121.
- Thaler, Richard H (2018a). "From cashews to nudges: the evolution of behavioral economics". In: *American Economic Review* 108.6, pp. 1265–87.
- (2018b). "Nudge, not sludge". In: *Science* 361.6401, pp. 431–431.
- Thaler, Richard H and Shlomo Benartzi (2004). "Save more tomorrow™: Using behavioral economics to increase employee saving". In: *Journal of political Economy* 112.S1, S164–S187.
- Thaler, Richard H and Cass R Sunstein (2009). *Nudge: Improving decisions about health, wealth, and happiness*. Penguin.
- Treme, Julianne et al. (2009). "Widely Received: Payoffs to Player Attributes in the NFL". In: *Economic Bulletin* 23.9, pp. 1631–1643.
- Tucker, Catherine et al. (2011). "How does popularity information affect choices? A field experiment". In: *Management Science* 57.5, pp. 828–842.
- Vincent, Claude et al. (2009). "Determinants of Pay in the NHL: A Quantile Regression Approach". In: *Journal of Sports Economics* 10.3, pp. 256–277.
- Vollmer, Anna-Lisa et al. (2018). "Children conform, adults resist: A robot group induced peer pressure on normative social conformity". In: *Science Robotics* 3.21.
- Waldfoegel, Joel et al. (2006). "Does information undermine brand? Information intermediary use and preference for branded web retailers". In: *The Journal of Industrial Economics* 54.4, pp. 425–449.
- Wallsten, Scott (2013). *What are we not doing when we're online*. Tech. rep. National Bureau of Economic Research.
- Ward, Michael R (2011). "Video Games and Crime". In: *Contemporary Economic Policy* 29.2, pp. 261–273.
- Wolfers, Justin (2006). "Point Shaving: Corruption in NCAA Basketball". In: *American Economic Review* 96.2, pp. 279–283.
- Yeung, Karen (2017). "'Hypernudge': Big Data as a mode of regulation by design". In: *Information, Communication & Society* 20.1, pp. 118–136.
- Zhang, Laurina (2018). "Intellectual property strategy and the long tail: Evidence from the recorded music industry". In: *Management Science* 64.1, pp. 24–42.

Résumé en Français

Cette thèse traite des impacts de la numérisation sur les comportements économiques. Le chapitre 1 traite de la réduction des coûts de transport et de recherche et de son impact sur le marché du travail en ligne. Ces dernières années, on a assisté à un passage croissant du travail en personne au travail à distance, ce qui a créé de nouveaux emplois et un nouvel environnement de travail dans lequel les travailleurs peuvent être plus enclins à gérer leurs heures de travail et les compromis entre travail et loisirs. J'utilise l'introduction du jeu mobile Pokémon Go, largement répandu, pour observer son impact sur les comportements des travailleurs sur la plateforme de travail en ligne, Amazon Mechanical Turk. Je montre qu'une augmentation de l'utilisation relative de Pokémon Go entraîne une diminution quotidienne de la proportion de travailleurs américains sur Amazon Mechanical Turk. Le chapitre 2 étudie les résultats d'une expérience de laboratoire sur le terrain menée dans une école primaire française en juillet 2019, où un haut-parleur intelligent, un robot et un adulte tentaient d'influencer les enfants dans leur choix de partager des billes avec d'autres enfants. Nous avons adapté un jeu de dictateur pour le public d'enfants, puis estimé l'impact de deux stratégies de nudging différentes (proximité sociale et effet de pair) sur le résultat du jeu de dictateur. Au cours de l'interaction avec les enfants, les nudges étaient moins efficaces lorsqu'ils étaient mis en œuvre par l'adulte que par les assistants vocaux, ce qui met en lumière le potentiel de ces dispositifs émergents lorsqu'il s'agit de manipuler leur public vulnérable. Le chapitre 3 étudie les motifs économiques de la manipulation de la popularité sur les médias sociaux par l'acquisition de faux followers par les joueurs de football professionnels. Nous tirons parti de la suppression des comptes suspects de Twitter tenue en juillet 2018 pour approximer les faux followers. Les résultats montrent que les faux followers ont un impact significatif sur la valeur des joueurs, c'est-à-dire sur les frais de transfert, uniquement si le transfert a lieu dans les 1 à 6 mois suivant la création du compte Twitter.