

KU LEUVEN

FACULTY OF ECONOMICS
AND BUSINESS

Essays on the Economics of the Pharmaceutical Industry



Dissertation presented to
obtain the degree of Doctor in
Business Economics

by

Melissa Newham

Number 737

2021

KATHOLIEKE UNIVERSITEIT LEUVEN

DOCTORAL THESIS

**Essays on the Economics of
the Pharmaceutical Industry**

Author:

Melissa Newham

Supervisor:

Prof. Jo Seldeslachts

*A thesis submitted in fulfillment of the requirements
for the degree of Doctor in Business Economics*

at the

Department of Management, Strategy and Innovation

of the

Faculty of Economics and Business

April 2021

Since the dissertations in the series published by the Faculty of Economics and Business are the personal work of their authors, only the latter bear full responsibility.

Daar de proefschriften in de reeks van de Faculteit Economie en Bedrijfswetenschappen het persoonlijk werk zijn van hun auteurs, zijn alleen deze laatsten daarvoor verantwoordelijk.

Supervisor:

Prof. Jo Seldeslachts
KU Leuven & DIW Berlin

Doctoral committee:

Prof. Albert Banal-Estañol
Universitat Pompeu Fabra & City University London

Prof. Tomaso Duso
DIW Berlin & Technische Universität (TU) Berlin

Prof. Fiona Scott Morton
Yale University

Prof. Otto Toivanen
Aalto University & KU Leuven

Prof. Frank Verboven
KU Leuven

Abstract

This dissertation consists of four essays which advance our understanding of the pharmaceutical industry and how to effectively regulate this sector. Each chapter aims to provide new insights, which are relevant for policymakers, based on rigorous economic analysis of real-world data. The chapters in this dissertation focus on the US pharmaceutical industry; the largest market for pharmaceutical products in the world in terms of value. The first essay in this dissertation studies voting behavior in the FDA's advisory committees which vote on questions related to the approval of new drugs. The second essay investigates the determinants of the market entry decisions of pharmaceutical firms that produce generic medicines. In particular, it investigates whether ownership links between generic manufacturers and the originator (or brand) firm created by "common investors" (investors that own shares in both the brand and generic firm) affect the likelihood that a generic drug will be launched on the market. The third essay investigates the presence of ownership links created by large shareholders in the pharmaceutical industry more broadly, using tools from network analysis, and discusses how common ownership may affect product market outcomes in the sector. The final essay studies the extent to which promotional gifts and other transfers made to physicians by pharmaceutical companies influences the prescription decisions of physicians and how this affects healthcare costs.

Acknowledgements

This thesis would not have been possible without the support and encouragement of a number of people. First and foremost, I thank my supervisor Jo Seldeslachts for his unwavering support and belief in me. Throughout my PhD, Jo Seldeslachts provided me with the freedom to pursue my research interests and work independently, while at the same time always being available to provide considered feedback and advice. I could not have asked for a better thesis supervisor, co-author and life coach to walk by my side during this journey.

I am deeply grateful to the members of my committee, Albert Banal-Estañol, Tomaso Duso, Fiona Scott Morton, Otto Toivanen and Frank Verboven, for their support throughout my PhD. I have been incredibly fortunate to benefit from your insightful comments and helpful career advice. I thank my co-author Rune Midjord for countless productive discussions regarding our joint paper. I fondly recall working together in cafes in Copenhagen. I am grateful to the Flemish Science Foundation for trusting in and financing my research vision through the FWO-ICM fellowship (project 1103419N).

This thesis was written over the course of several years spent working in different locations. I have benefited greatly from being exposed to different research environments and communities. I thank the faculty members, PhD students and other members of staff at the department of Management, Strategy and Innovation at KU Leuven for their support and helpful assistance. I would like to offer special thanks to Reinhilde Veugelers for her enthusiasm for my research and career advice.

In Berlin, my home has been the Firms and Markets department at DIW Berlin. I thank all of my colleagues in Berlin for the role they have played in inspiring and encouraging me through the good times, and the more challenging times. In particular, I thank Tomaso Duso for his constant support, Hannes Ullrich for always taking the time to read my work and provide comprehensive feedback, and to my office mate Pauline Affeldt for her calm advice and friendship throughout this time.

I would like to extend my sincere gratitude to the academics who have hosted me abroad. In South Africa, I offer special thanks to Simon Roberts and the team at CCRED in Johannesburg. I thank Lukasz Grzybowski for hosting me at the University of Cape Town. In the US, I thank Federico Ciliberto at the University of Virginia and Fiona Scott Morton, Mike Sinkinson and Florian Ederer at Yale University. I also want to acknowledge Massimo Motta and Adina Claiici, who I

first met while working at the Chief Economist Team in Brussels before starting my PhD, and who have remained sources of support and inspiration for me during my studies.

Finally, and most importantly, I wish to thank my friends and family. In particular, I thank Annekatrin Schrenker, Boryana Ilieva, Jan Malek, Joanna Piechucka, Kai Barron, Kevin Tran, Marrit Teirlinck, Martin Gross, Max Schäfer and Nuria Boot. I thank the friends in my eco's group from UCT who started out in economics with me and continue to inspire me with the type of work and research that they do.

A very big thank you goes to my parents, Andrew and Priska Newham, and my brother Simon Newham, for their immense love and support throughout this journey. I could not have done it without you!

Contents

Abstract	iii
Contents	vi
List of Tables	ix
List of Figures	xi
1 Introduction	1
1.1 Overview of the pharmaceutical industry	2
1.2 Outline of the dissertation	5
2 Do Expert Panelists Herd? Evidence from FDA Committees	9
2.1 Introduction	9
2.2 Model	14
2.2.1 Herd voting	19
2.2.2 The likelihood function	20
2.3 Data	21
2.3.1 FDA advisory committees	21
2.3.2 Data collection and variables	23
2.4 Descriptive analysis	27
2.4.1 Vote outcomes	27
2.4.2 Voting behavior	28
2.5 Estimation and identification	32
2.6 Estimation results	35
2.6.1 Estimates	35
2.6.2 Herd votes	38
2.7 Information aggregation	40
2.7.1 Probability of a correct assessment	40
2.7.2 Votes and FDA approvals	43
2.8 Conclusion	44
2.9 Appendix	46

2.9.1	Seating order	46
2.9.2	Correlation in vote questions part of the same meeting	47
2.9.3	Alternative specification of the prior	48
2.9.4	Additional regressions	49
2.9.5	Counterfactual with 100% herd types	50
3	Common Ownership and Market Entry: Evidence from the Pharmaceutical Industry	51
3.1	Introduction	51
3.2	Literature	55
3.3	Theoretical framework	57
3.4	Data	61
3.5	Individual entry	68
3.6	Market outcomes	80
3.7	Conclusion	85
3.8	Appendix	87
3.8.1	Common ownership	87
3.8.2	Dataset construction	89
3.8.3	Formal model and proofs	91
3.8.4	Robustness	97
4	Common Ownership in the US Pharmaceutical Industry: A Network Analysis	104
4.1	Introduction	104
4.2	Background	107
4.2.1	Pharmaceutical industry	107
4.2.2	Institutional investors and common ownership	109
4.2.3	Institutional investors' influence in pharma	111
4.3	Data and descriptive statistics	113
4.4	Network analysis	116
4.4.1	Common ownership links between top firms	117
4.4.2	Investor networks	128
4.4.3	Centrality in the brand and generic networks	131
4.5	Antitrust implications	138
4.5.1	Innovation	139
4.5.2	Entry	141
4.5.3	Pricing	142
4.6	Conclusion	144

5	The Interaction between Industry Payments to Physicians, Insurance and Drug Costs: Evidence from Medicare Part D	145
5.1	Introduction	146
5.2	Contribution to the literature	149
5.3	Data	151
5.3.1	Medicare Part D	152
5.3.2	Diabetes	153
5.3.3	Payments to physicians	155
5.3.4	Dataset construction	156
5.3.5	Descriptive analysis	158
5.4	Empirical strategy	160
5.4.1	Causal model	160
5.4.2	Using machine learning for prediction	162
5.4.3	Estimators	164
5.4.4	Comparison with matching	166
5.5	Estimation and results	167
5.5.1	Estimation and out-of-sample performance of the ML models	168
5.5.2	Effect of industry payments	170
5.6	Discussion	175
5.7	Concluding remarks	177
5.8	Appendix	178
5.8.1	Tables and figures	178
5.8.2	Treatments for diabetes	183
5.8.3	Illustrative model	187
	Bibliography	190

List of Tables

2.1	Voting question characteristics	26
2.2	Voter characteristics	27
2.3	Descriptive statistics for vote outcomes	28
2.4	Summary statistics	30
2.5	Reduced-form results	32
2.6	Estimation results	38
2.7	Simulated outcomes with Bayesian updating	40
2.8	Simulated outcomes with naïve updating	40
2.9	Alignment of FDA decisions and unanimous voting outcomes	44
2.10	Summary statistics for seat no. by voter characteristics	46
2.11	Summary statistics for seat no. for 10 most frequent voters	46
2.12	Estimation results	48
2.13	Additional reduced-form results	49
3.1	Top 5 largest investors (2013)	68
3.2	Example data structure	71
3.3	Summary statistics	72
3.4	Main specification	74
3.5	Cross-correlations between common ownership measures	75
3.6	Categorical specification	76
3.7	Instrumental variables regression	77
3.8	First-stage IV regressions	79
3.9	Summary statistics at the market-level	82
3.10	Market outcomes - δ_S (market)	84
3.11	Market outcomes - Share links $\delta_S > 0.5$	85
3.12	Robustness - Drug market fixed effects	97
3.13	Robustness - Probit	98
3.14	Robustness - Logit	99
3.15	Robustness - Indicator for top 100 sales	100
3.16	Robustness - Using total US sales for drugs in top 100	101
3.17	Robustness - Broader entrant set	102

3.18	Robustness - Entry ever	103
4.1	Top five investors in top brand firms	110
4.2	Top five investors in top generic firms	111
4.3	Top 10 common investors in brand firms	115
4.4	Top 10 common investors in generic firms	116
4.5	Levels of centrality for the 20 brand companies	133
4.6	Levels of centrality for the 20 generic companies	134
5.1	Payment value (USD) by drug cost	157
5.2	Summary statistics for physician characteristics (2014-2017)	159
5.3	Summary statistics for payments related to anti-diabetic drugs	159
5.4	Out-of-sample performance - Brand prescription rate	169
5.5	Out-of-sample performance - Drug cost per 30-day supply	170
5.6	Results - Brand prescription rate	173
5.7	Results - Drug cost per 30-day supply	174
5.8	Firms and brands with payments to physicians in the sample	178
5.9	Variable definitions	179
5.10	Diabetes treatments in the sample	185

List of Figures

1.1	Schematic of the life-cycle of a pharmaceutical drug	2
2.1	Probability density function of private signals $ \theta = 1$	16
2.2	Cut-off signals (illustration with $\tau = 1$)	18
2.3	Agreement of votes	28
2.4	Information aggregation	41
2.5	Information aggregation with a lower precision of information	43
2.6	Information aggregation with $\tau = 1.26$ and $\lambda=1$	50
2.7	Information aggregation with $\tau = 1$ and $\lambda=1$	50
3.1	Exclusivities and patent protection in pharmaceuticals	62
3.2	Histogram of entry	63
3.3	Decline in brand and total revenues with number of generic entrants	66
3.4	Evolution of common ownership	70
3.5	Evolution of common ownership - Rank measures	70
3.6	iShares U.S. Pharmaceutical ETF (IHE) - Snapshot of Holdings	88
3.7	Profitable entry of g as a function of δ and p_1	92
3.8	Number of entrants in equilibrium as a function of δ	92
4.1	Common ownership network of the top 20 brand firms - Individual ownership	119
4.2	Common ownership network of the top 20 brand firms - Joint ownership	121
4.3	Common ownership network of the top 20 generic firms - Individual ownership	123
4.4	Common ownership network of the top 20 generic firms - Joint ownership	124
4.5	Bipartite network of brands and generics - Individual ownership	126
4.6	Bipartite network of brands and generics - Joint ownership	127
4.7	Investor networks in brand firms	129
4.8	Investor networks in generic firms	130
4.9	Average degree centrality of brand and generic firms over time (2004-2014)	135

4.10	Histogram of degree centrality for brands within the 85-brand firm network in 2004 and 2014	137
4.11	Histogram of degree centrality for generics within the 25-generic firm network in 2004 and 2014	138
5.1	Costs to treat diabetes	154
5.2	Comparison of physicians' prescription outcomes	160
5.3	Distribution of error term - Brand prescription rate	182
5.4	Distribution of error term - Drug cost per 30-day Supply	182
5.5	Prescriptions by drug class	186
5.6	Total drug cost by drug class	186
5.7	Price threshold for brand prescription	189

Chapter 1

Introduction

The pharmaceutical industry makes a significant contribution to global health and economic welfare through the development and production of innovative medical treatments. The COVID-19 pandemic has showcased the ability of the industry (with the support of governments) to conduct research and deliver vital medical treatments and vaccines to fight the virus within record time. However, while the industry is responsible for many breakthrough treatments, the conduct of pharmaceutical firms often causes controversy. For example, firms sometimes provide misleading information concerning the side-effects of drugs, and there is frequent debate over whether drug prices are set too high. A well-functioning pharmaceutical industry is a priority for governments and regulation plays a crucial role in this industry. Almost every activity in the industry is regulated from product development through to manufacturing and marketing.


This dissertation comprises four essays which contribute towards our understanding of the pharmaceutical industry and how to effectively regulate this sector. Each chapter aims to provide new insights, based on rigorous economic analysis of real-world data, that are relevant for policymakers. The chapters in this dissertation focus on the US pharmaceutical industry. The focus on the US is motivated by the fact that the US pharmaceutical industry is the largest in the world. The US accounts for approximately half of global pharmaceutical sales, and 65% of the sales of new medicines launched between 2013-2018 were made in the US. The US pharmaceutical industry also leads in terms of research and development (R&D) expenditures (EFPIA, 2019).

This introduction is split into two sections. The first section aims to provide the reader with a high-level overview of how the pharmaceutical industry works, with an emphasis on the US market, and positions each chapter within this bigger picture. The second section provides a summary of each chapter in this dissertation.

1.1 Overview of the pharmaceutical industry

Figure 1.1 provides a simplified schematic of the typical “life-cycle” of a pharmaceutical product. The first stage in a drug’s life-cycle is the research and development stage. Each drug begins with discovery and development in a laboratory. Once a drug candidate has been identified, researchers undertake preclinical studies using laboratory experiments and testing in animal subjects. If the preclinical results are sufficiently promising, the drug candidate progresses to testing in humans. There are three phases of human clinical trials. In Phase I, the safety of the drug is tested with a small sample (ca. 20-100) of healthy individuals. In Phase II, the effect of the drug is tested in a larger group of people (ca. 50-300). Phase III trials involve large groups of subjects (ca. 300-3,000 or more) and aim to provide a definitive assessment of how effective the drug is. The process of developing and testing a new drug is lengthy, expensive and unpredictable. Among the largest pharmaceutical firms, roughly 20-30% of drugs that begin Phase I clinical trials end up being approved for use (Adams and Brantner, 2006).

Figure 1.1: Schematic of the life-cycle of a pharmaceutical drug



Pharmaceutical companies typically patent newly designed molecules early in the process. Patent protection ensures that if the drug reaches the market it will enjoy a monopoly for a period of time (“market exclusivity”). Patent protection is important in the pharmaceutical industry because the costs of imitation are quite low: once a product is known to be safe and effective, it can be backward-engineered with little difficulty. Patents and other forms of market exclusivity created by regulation allow firms to earn high profits on new drugs which compensate for R&D expenditures and incentivize innovation.

If a drug is successful in the final stage of clinical trials (Phase III), the innovator applies for approval from the relevant regulatory body. In the US, the Food and Drug Administration (FDA) is responsible for approving new drugs. In Europe, it is the European Medicines Agency. Regulators evaluate the results of clinical trials

and weigh this against the known safety risks and existing treatment options. In this way, regulatory agencies serve an important function. In the pharmaceutical industry, there is asymmetric information between the user and the producer: a consumer cannot tell simply by looking at a drug (or even by taking it) if it is a good product. The regulator provides certification of drug quality. Regulation is also important because the social costs of a bad drug can be very high. Approved drugs are then launched on the market and marketed under a “brand-name” for example Aspirin. Once on the market, the safety of drugs is further monitored, and pharmaceutical firms are subject to detailed post-marketing reporting requirements which include the reporting of adverse events.

In the US, firms that wish to market a new drug must file a New Drug Application (NDA) or a Biologics License Application (BLA) with the FDA. To assist in making difficult approval decisions, the FDA makes use of advisory committees. These committees consist of around a dozen medical experts who vote on yes/no questions related to drug approvals. Chapter 2 in this dissertation investigates how the organization of these advisory committees affects how the committee members vote and, in turn, how this affects the final assessment made by the committee. The findings of Chapter 2 are relevant not just for the FDA’s expert committees, but also have broader implications for how voting procedures should be set up in other expert committees.

The approval process for generic drugs is slightly different. A generic drug is a medication created to be the same as an already marketed brand-name drug in all aspects; dosage form, safety, strength, route of administration, quality, performance characteristics, and intended use. Firms that produce generic drugs can enter the market once the regulatory protections afforded to the brand product have expired or have been successfully challenged in court. To apply for the approval of a generic drug, firms must submit an Abbreviated New Drug Application (ANDA) to the FDA. Approval is fairly straightforward as the firm needs only to show that the drug is bioequivalent to the original product and is safely manufactured. Thus, the costs associated with bringing a generic drug to market are far lower than for new drugs. Typically, firms in the pharmaceutical industry specialize either in the development and production of new drugs, or in the production of generic drugs. Thus, one can speak of “generic firms” that primarily develop and launch generic drugs and “brand firms” that have R&D capabilities and launch new drugs.

Once on the market, drugs are sold for a specific price. The pharmaceutical industry is particular in that it is often not the end-user of the product that pays (fully) for it. In the US, private and public health insurers account for about 80% of prescription drug spending (Cubanski et al., 2019). New prescription drugs in the

US tend to have very high prices, higher than in other high-income countries. This is because, unlike in other countries, the US government does not directly regulate or negotiate the price of drugs. Firms can set high prices because, in the case of vital medical treatment, demand is very inelastic. Moreover, while the drug is protected by a patent it faces competition only from non-identical drugs that treat the same underlying condition, which may be poor substitutes. Several studies find that this “brand-brand” competition does not effectively lower list prices (Sarpawari et al., 2019).

High prescription drug prices are a concern for governments. In the US, the entry of generic drugs is crucial for lowering prices. For products with a single generic producer, the generic average market price is 39% lower than the brand average market price before generic competition. With six or more competitors, generic prices show price reductions of more than 95% compared to brand prices.¹ Accordingly, promoting generic entry is an important policy goal for the FDA.²

Generic entry marks the end of the brand firm’s monopoly on the drug and, because generic drugs are available at a lower price, it leads to a massive decline in revenues from the brand drug. Consequently, brand firms employ several strategies to delay or deter generic entry. Chapter 3 investigates one way in which generic entry may be deterred. In particular, Chapter 3 tests whether ownership links between generic firms and brand firms created by “common investors” (investors that have shares in both the brand and generic firm) reduce the likelihood of generic entry. Numerous pharmaceutical firms are owned by the same set of investors. This phenomenon is known as “common ownership”. Chapter 4 documents the evolution and extent of common ownership in the pharmaceutical industry and discusses the implications thereof for firms’ strategies and product market outcomes. The topic of common ownership and its effect on product market outcomes is a major concern for policymakers, particularly in antitrust. It has been described as “the major new antitrust challenge of our time” (Posner et al., 2017). While Chapters 3 and 4 focus on common ownership in the pharmaceutical industry, the findings also contribute to the broader, ongoing debate about the effect of common ownership on competition.

After a new drug is launched on the market, pharmaceutical firms devote enormous resources to marketing and promoting their drug – just as much as they do on

¹See FDA website, New Evidence Linking Greater Generic Competition and Lower Generic Drug Prices. Available at: <https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/generic-competition-and-drug-prices>

²See FDA website, Statement from FDA Commissioner Scott Gottlieb, M.D., on new policy to improve access and foster price competition for drugs that face inadequate generic competition [Press release]. 19 February 2019. Available at: <https://www.fda.gov/news-events/press-announcements/statement-fda-commissioner-scott-gottlieb-md-new-policy-improve-access-and-foster-price-competition>

research and development (Gagnon and Lexchin, 2008). While healthcare payers and consumers pay for prescription drugs, physicians control access. Thus, the largest share of promotional expenditure is spent on “detailing” which is the promotion of prescription drugs directly to physicians. In the US, it is not uncommon for pharmaceutical firms to promote their drugs by making transfers of value to physicians. These transfers can take various forms, for example, gifts, free meals, consulting and speaker fees, funding for further education and research grants. This practice is controversial and has been subject to growing public and academic debate. Some states in the US ban certain types of payments to physicians while in other states the practice is left unregulated. Understanding whether and how detailing affects healthcare costs and outcomes is useful information for regulators seeking to design effective policies or implement bans. To this end, Chapter 5 investigates whether payments from pharmaceutical companies cause physicians to prescribe more expensive medicines, and which factors affect the relationship between payments and prescribing choices.

1.2 Outline of the dissertation

Chapter 2 (“Do Expert Panelists Herd? Evidence from FDA Committees” co-authored with Rune Midjord) studies voting behavior in the FDA’s advisory committees. This chapter addresses the question whether, and to what extent, expert panelists engage in herd behavior when voting on questions related to the approval of new drugs, and how this affects the quality of the final assessment made by the committee. This research contributes to the literature by being the first to empirically investigate herd behavior in voting in expert committees and assess its consequences for information aggregation. Herding occurs when panellist are swayed by the votes of members who voted before them. Herd behavior is notoriously difficult to measure empirically. In many contexts, decisions may be clustered for reasons other than herding. Our approach solves this problem in two ways. Firstly, we develop a structural model to quantify herd behavior. The structural approach allows us to separate confounding factors (including expertise, cautiousness and the strength of the common prior) from members’ inclination to herd. Secondly, we exploit a change in the voting procedure for FDA committees from sequential to simultaneous voting in 2007. Access to simultaneous data allows us to get a grip on key parameters of the model, like members’ expertise, when there is no herding at play.

For our analysis, we construct a database using the verbatim transcripts of FDA committee meetings. We find that experts are indeed susceptible to herd behavior: around half of the panelists are willing to vote against their own assessment if votes

from previous experts indicate otherwise. We find that there is heterogeneity in herd behavior; temporary committee members are more prone to herding than regular (standing) members. Additionally, we show that herd behavior has detrimental consequences for information aggregation. This research has implications for how to optimally structure voting procedures in advisory committees. The main policy implication of this research is that expert committees with similar features to the FDA committees should substitute sequential voting with simultaneous (electronic) voting.

Chapter 3 (“Common Ownership and Market Entry: Evidence from the Pharmaceutical Industry” co-authored with Jo Seldeslachts and Albert Banal-Estañol) addresses the question whether and to what extent common ownership affects the market entry decisions of pharmaceutical firms that produce generic drugs. Research on the topic of common ownership (where two firms are partially owned by the same investor) and its effect on product market outcomes has recently gained the attention of policymakers, the media and academics. Firms that are largely owned by shareholders who also have sizeable stakes in competitors might just simply act in these shareholders’ interest, which leads them - rather than maximizing their own profits - to maximize the return of their shareholders’ portfolios. This research contributes to the literature by being one of the first papers to directly investigate the effect of common ownership on entry.

In the pharmaceutical industry, maintaining monopolized markets is crucial for firms that sell brand-name drugs. With the event of generic entry, the revenues derived from the brand drug can decline by as much as 90%. Moreover, brand losses typically outweigh what generics stand to gain from market entry on average. Thus, entry decisions may crucially depend on whether owners of generic firms also have an interest in brand firms.

To empirically analyze the effect of common ownership on entry decisions, we use data on patent expiration and drug approvals in the US from the FDA Orange Book combined with data on the ownership structure of pharmaceutical companies from Thomson Reuters Global Ownership Database. We find that brand-name drugs that are produced by firms that have significant shareholder overlap with firms that produce generic drugs, face less competition from generic drugs. The effect is large: a one-standard-deviation increase in common ownership with the brand decreases the probability of entry by that generic firm by 15-18%. We also find that common ownership has an economically significant effect on the total number of generic firms in a specific drug market. Our research provides evidence that common shareholders indeed influence strategic decisions of companies. This has implications for competition policy and antitrust. Given the importance of generic

entry in terms of reducing drug prices, common ownership in the pharmaceutical industry has the potential to raise healthcare costs.

Chapter 4 (“Common Ownership in the US Pharmaceutical Industry: A Network Analysis” co-authored with Jo Seldeslachts and Albert Banal-Estañol) applies tools from network analysis to study how common ownership has evolved in the US pharmaceutical sector. Further, we discuss the implications of our findings for competition policy. Given that common investors are both influential and, as we show, have substantial ownership stakes in multiple pharmaceutical firms, common ownership links may affect competition and innovation in the industry. There are surprisingly few papers that make use of network analysis to study and visualize common ownership patterns. This research contributes to the literature by analyzing the structure and characteristics of common ownership networks in the pharmaceutical industry.

Our data comprises of publicly owned pharmaceutical firms that were active in the US pharmaceutical market between 2004 and 2014. The data sources are the same as in Chapter 3. Our empirical analysis yields three main findings. Firstly, we find that brand firms are strongly linked to each other by large institutional investors and that the network created by ownership links has become increasingly dense over time. Secondly, in contrast to this, the network of generic firms is much sparser and stays this way over the time span of our sample. Finally, when considering the common ownership links between brand firms, on the one hand, and generic firms, on the other, we find that brand firms have become more connected to generic firms over time.

Overall, our analysis indicates that common ownership is widespread and increasing for the brand firm and brand-generic firm network. Seen in combination with the results in Chapter 3, the increase in brand-generic connectivity appears to have led to a decrease in generic entry. The increase in connectivity between brand firms may also have led to higher drug prices as commonly owned firms have less incentive to compete. Common ownership may also affect innovation in the industry – both in positive and negative ways. Common ownership between brand companies may, on the one hand, enhance information sharing, generate synergies, and increase the incentives to invest in R&D. On the other hand, common ownership may also incentivize firms to innovate in a way that avoids head-on competition between each other in the innovation space. Given the importance of R&D and effective competition in the pharmaceutical industry, this chapter provides ample reason for policymakers to pay closer attention to common ownership in the pharmaceutical industry.

Chapter 5 (“The Interaction between Industry Payments to Physicians, Insurance and Drug Costs: Evidence from Medicare Part D”) analyzes the impact of

payments to physicians from pharmaceutical companies. To promote their drugs, pharmaceutical firms often make transfers of value to physicians e.g., free meals, consulting fees, research grants. The extent to which these payments influence physicians' behaviour and healthcare outcomes is unclear. This research investigates whether, and to what extent, such payments lead to the prescription of costlier drugs. Further, it tests whether physicians who have patients with lower out-of-pocket expenditures prescribe relatively more expensive medications in response to industry payments. Research shows that physicians consider how much a drug will cost a patient (out-of-pocket costs) when prescribing (Carrera et al. 2018; Lundin, 2000). If patients have very low (or zero) out-pocket-costs, there may be little push back if the physician prescribes a more expensive drug. Hence, when patients have lower out-of-pocket costs, physicians may be more likely to prescribe costlier medication as a consequence of payments. This research contributes to the existing literature by providing an estimate of the causal effect of industry payments on the cost of treatment for diabetes and by showing that this effect varies with a measure of patients' out-of-pocket costs. Moreover, this research develops a novel empirical strategy to identify the aforementioned effects.

In the analysis, data from a federal database on the universe of payments to physicians in the US between 2014 and 2017 is linked to prescribing behavior in Medicare Part D. To identify the effect of industry payments, the analysis uses data on the prescribing patterns of physicians in Vermont, where a strict ban on industry payments to physicians is in place, combined with machine learning techniques to construct the counterfactual outcome for physicians who receive payments in the nearby states of New Hampshire and Maine.

The empirical analysis focuses on the prescription of anti-diabetic medications. The main findings are that receipt of payments related to anti-diabetic medication increases the average brand prescription rate by 5 percentage points and the average drug cost per dose by 21 USD. Physicians with a higher share of patients with a low-income subsidy, and who therefore face lower out-of-pocket expenditures, prescribe relatively more expensive medication in response to receiving a payment. Back-of-the-envelope calculations based on our estimates suggest that banning industry payments would result in a 3% decline in total prescription costs for diabetes. Given that existing research does not point to strong informational benefits of payments to physicians, the main policy implication of this research is that a ban on industry payments is likely to be an effective way to contribute towards healthcare cost containment in the US.

Chapter 2

Do Expert Panelists Herd? Evidence from FDA Committees¹

Chapter Abstract

This chapter develops a structural model to address the question whether, and to what extent, expert panelists engage in herd behavior when voting on important policy questions. The data comes from FDA advisory committees voting on questions concerning the approval of new drug applications. The analysis utilizes a change in the FDA's voting procedure from sequential to simultaneous voting to identify herding. Estimates suggest that around half of the panelists are willing to vote against their private assessment if votes from previous experts indicate otherwise and, on average, 9 percent of the sequential votes are actual herd-votes. Temporary committee members are more prone to herding than regular (standing) members. We find that simultaneous voting improves information aggregation given our estimates.

2.1 Introduction

Many important decisions within public and private organizations are based on recommendations from expert committees. Advisory boards give strategic advice to the management of corporations and expert committees recommend on issues such as climate, national security, education, and medical drugs.² The main advantage of expert committees is their ability to aggregate multiple sources of information and hence allow for more informed decisions. A common way to gauge the information

¹This chapter is published in the DIW Discussion Paper Series as: Newham, M. and Midjord, R. (2020). Do Expert Panelists Herd? Evidence from FDA Committees. *DIW Berlin Discussion Paper No. 1825*.

²In 2006 the United States government maintained 916 federal advisory committees composed of 67,346 members (Brown, 2009).

held by individual committee members is to have a sequential vote (by roll call or going stepwise around the committee table). For example, roll call voting is used by committees of the European Parliament and it is one of the official voting procedures for advisory committees under the United Nations. In the United States, roll call voting is often used by advisory committees under city and town councils and occasionally by committees at the federal level.³ When committee members vote sequentially, the information contained in the vote will be affected, and possibly diluted, if members engage in herd behavior (i.e. if they are swayed by observing the preceding votes). To this end, it is crucial to know whether, and to what extent, expert panelists engage in herd behavior and how this affects information aggregation.

In this paper, we develop a structural model to estimate the prevalence of informational herding (see Banerjee, 1992; Bikhchandani, Hirshleifer, and Welch, 1992; Welch, 1992; Smith and Sørensen, 2000) and investigate its consequences for information aggregation in the U.S. Food and Drug Administration’s (FDA) advisory committees. While other papers have studied herding in different empirical settings⁴ and laboratory experiments,⁵ this is the first paper to estimate a model of herd behavior in advisory committees. Moreover a change in the voting procedure for FDA committees from sequential to simultaneous voting in 2007 provides a remarkably “clean” natural experiment which we take advantage of to identify herd behavior.

The drug approval decisions made by FDA affect millions of users; if beneficial drugs do not win approval patients miss the opportunity of improved medication and if bad drugs are approved the consequences can be fatal.⁶ To assist in making difficult approval decisions, the FDA makes use of advisory committees. The committees consist of around a dozen medical experts who vote in a fixed order on yes/no questions related to drug approvals, for example, “Should omapatrilat be approved for the treatment of hypertension?” In 2007 the FDA changed the voting procedure for their advisory committees from sequential to simultaneous voting citing concerns of “momentum” effects in sequential voting. The concern was that some sequential voters may be influenced by the preceding votes, especially if those votes signal a

³For example, the advisory committees under the Federal Communications Commission and the Health Resources and Services Administration.

⁴To mention a few: presidential primaries (Knight and Schiff, 2010), restaurant dining (Cai, Chen, and Fang, 2009), investment recommendations (Graham, 1999), stock market trading (Cipriani and Guarino, 2014), financial decisions (Bursztyn et al., 2014), movie sales (Moretti, 2011), and movie reviews (Camara and Dupuis, 2014).

⁵See Anderson and Holt (1997) for an early reference and Weizsäcker (2010) for a meta analysis.

⁶One of the most debated FDA decisions is the approval of the painkiller Vio (nytimes.com/topic/subject/vioxx-drug). According to Graham et al. (2005) Vioxx caused an estimated 88,000 to 140,000 excess cases of serious heart disease in the U.S. over its market life.

clear trend.⁷

Our data is gathered from verbatim transcripts of FDA committee meetings held between 1996 and 2014. In our dataset there are 813 voting questions for which 1,378 unique panelists cast a combined total of 10,466 votes. Roughly half of the voting questions take place under sequential voting. Data on individual experts includes their educational background, gender and “voter category.” FDA committees operate with four categories of voting members. *Regular* (standing) members serve four-year terms and have recognized expertise in a relevant field. Additional experts are usually added as *temporary* voting members. Like regular members, temporary members have expertise in a relevant field. Each committee also has a qualified *consumer* and a *patient representative* as voting members. We make use of this information to study heterogeneity in herd behavior across different types of committee members.

In our model, committee members vote on independent yes/no questions e.g. whether the benefits of a new drug outweigh the risks. For each voting question, there is a common prior on the correct answer (the state) being “yes.” The prior contains relevant public information on the question at hand including committee discussions and presentations leading up to the voting stage. On top of this, each expert receives a private signal about the state which depends on the precision of private information and the true state. In this respect, each panelist draws on his/her unique experience, intuition, and analytical skills. In the empirical implementation, we allow for the general precision of the continuous signals to vary across committee members with certain observable characteristics (e.g. educational background and voter category) and recover estimates for members’ ability.

There are two types of committee members in our model. The *herd* type uses public information, his/her private signal, and the *vote-history* to update beliefs about the state and votes “yes” if, and only if, the updated beliefs exceed his/her standard of proof for voting yes. Intuitively, if the vote history is dominated by “yes” votes the herd type updates in a way that favors the “yes” state. However, the order in which the votes are placed matters for the belief updating and if a few “no” votes are the most recent ones and come from members with high expertise then the belief updating may favor the “no” state. By contrast, the *expressive* type considers only public information and his/her private signal (ignoring any preceding

⁷ “There has been much discussion inside and outside FDA regarding sequential versus simultaneous voting...scholars and social scientists have studied the risk of “momentum” in sequential voting, exploring whether some voters may be influenced, perhaps even subconsciously, by the votes that precede theirs, especially if those votes are nearly identical or signal a clear trend. [footnote reference to Banerjee (1992) and Callander (2007)].” In Draft Guidance for FDA Advisory Committee Members and FDA Staff: Voting Procedures for Advisory Committee Meetings (2007).

votes) when casting his/her vote.⁸ A member’s type is private information. In one version of the model we assume herd types are fully Bayesian and take into consideration the probability that committee members before them are also herd types, in another version we assume that they are “naïve” and believe that all preceding votes come from expressive types (following Eyster and Rabin, 2010). We estimate the proportion of herd types and members’ standards of proof, allowing these parameters to vary with observable characteristics.

The problem empiricists face when seeking to measure herd behavior is that private information is not observable so it is difficult to tell when agents have altered their behavior due to observing the actions of others. In many contexts, decisions may be clustered for reasons other than herding. For example, in our context, if a drug is clearly a good drug then we expect many experts will vote in favor of the drug. The structural approach allows us to separate the confounding effects of members’ inclination to herd, their level of expertise, their degree of cautiousness (standard of proof required to vote in favor of a drug) and the strength of the common prior.

To identify the model’s parameters we rely on the fact that we observe committee members voting across multiple voting questions and make use of information on the exact sequence of votes under sequential voting to identify herd voting. Importantly, the natural experiment allows us to observe voting outcomes when herd behavior can be ruled out. Intuitively, herding tends to make committee members follow the vote-trend under sequential polling which makes, all else equal, unanimous outcomes more frequent than absent any herding. However, high precision of the committee members private information also makes unanimous vote outcomes more common. Having access to simultaneous data allows us to get a grip on key parameters of the model, like the precision of members’ private information, when there is no herding at play.

Our results suggest that experts voting on important questions relating to drug approvals are indeed susceptible to herd behavior. On average, the proportion of herd types is close to one half (48% in the Bayesian updating version and 52% in the naïve updating version) i.e. around half of the expert panelists take into account, and are potentially swayed by, the votes that precede theirs. Interestingly, the share of herd types is considerably larger among temporary members than regular members.

The presence of herd types gives rise to the possibility of “herd votes.” A herd

⁸Our model is a model of statistical herding with expressive types added. A closely related concept is reputational herding where agents are motivated by appearing to be well informed about the state i.e. having strong private signals (see Ottaviani and Sørensen (2001) for a model of reputational concerns in committees).

vote occurs when a herd type is swayed by the vote history. For example, when he/she votes “yes” in a sequential set-up whereas he/she *would have* voted “no” under simultaneous voting. Using our model to simulate voting under the two regimes, we find that on average around 9% of the sequential votes are herd votes i.e. cases where members actually change their vote from what it would have been if they had ignored the vote history. This level of herding in the simulated data generates patterns in line with what we see in the real data, namely that the share of unanimous vote outcomes increases markedly under sequential voting.

Our estimates on the accuracy of the committee members private information are relatively high; on average private signals go in the wrong direction only around 20 percent of the time. Regular members have the most accurate private information whereas the signals of consumer and patient representatives are less precise. In general, the committee members are slightly cautious, meaning that it takes more than the preponderance of evidence for a committee member to vote “yes.” Consumer representatives are particularly cautious and at the other end of the spectrum patient representatives are neither cautious or incautious.

We next consider the consequences of herding on information aggregation. Although simultaneous voting is not informationally efficient for all parameter values of our model, we find that switching to simultaneous voting improves the probability that the committee’s assessment matches the state; where the effect is larger if we assume committee members apply naïve updating. Additionally, we use our model to demonstrate that the detrimental consequences of herding are exacerbated when information is less precise. Our main policy implication is therefore to follow the example of the FDA and substitute sequential “go-around” voting with simultaneous (electronic) voting.

Our model and estimation approach is inspired by the methodology in Cipriani and Guarino (2014) which serves as the first paper to estimate herd behavior with a structural model. In their application they estimate herd behavior in financial markets using transaction data from a publicly traded stock. Herding can occur over the course of a day due to uncertainty about whether an informational event has occurred and whether the fundamental value of the stock has increased or decreased. In their model, a market maker interacts with the sequence of traders and sets the price of the asset. Our model simplifies this framework by dispensing with the market maker and price mechanism. On the other hand, we incorporate heterogeneity regarding priors, signals and preferences, and develop a version of the model with naïve updating. Furthermore, we make use of data generated through simultaneous voting whereas in Cipriani and Guarino (2014) all transaction data is assumed to be generated sequentially.

By using a structural model to explain individual committee members' behavior, this paper also draws on the framework put forward in Iaryczower and Shum (2012). They explain decision-making in the U.S. Supreme Court by taking into account differences in the bias or ideology of justices, the information available to the justices and their ability to apply the law to the specifics of the case. Further research in this tradition includes Hansen, McMahon, and Rivera (2014) who explain individual voting behavior on the Bank of England's Monetary Policy Committee and Camara and Kyle (2016) who estimate a voting model to recover FDA committee members' skill and bias associated with financial ties. Iaryczower, Shi, and Shum (2018) use a structural model to quantify the effect of deliberation on the decisions of US appellate courts. Whereas Iaryczower, Shi, and Shum (2018) study the effect of deliberation before voting, we consider voting after collective discussions have taken place.

The FDA's reform of voting procedure in 2007 has been examined in a case study by Urfalino and Costa (2015). They collect data from six committees (202 voting questions) from 2003 to 2010 and report the proportion of unanimous, strong majority, and majority outcomes. Urfalino and Costa (2015) show that under simultaneous voting the proportion of unanimous outcomes is lower while the proportion of strong majority outcomes is higher. The authors suggest that these changes are due to reduced expert conformity following the shift to simultaneous voting. We extend this analysis in our descriptive and reduced-form section.

The rest of the paper is organized as follows. In Section 2.2 we introduce the theoretical model. Section 2.3 describes the data. Section 2.4 undertakes a descriptive and reduced-form (regression) analysis. In Section 2.5 we describe the estimation procedure and Section 2.6 discusses the main results. Section 2.7 considers information aggregation. Section 2.8 concludes.

2.2 Model

We consider advisory committees voting on various yes/no questions. As in the case with FDA committees, we can think of expert panelists polling on issues regarding a specific application, proposal, or scientific question. There are J voting questions and a generic voting question is denoted by $j \in \{1, \dots, J\}$. For each voting question, j , there is a common unobserved state $\theta^j \in \{0, 1\}$ that equals 1 if the correct answer to question j is "yes" and 0 if the correct answer to question j is "no." The state is independently drawn across the J voting questions. Let $\mu_0^j \in (0, 1)$ indicate the common prior belief that $\theta^j = 1$. The common prior contains relevant public information including committee discussions leading up to the voting stage. The

number of voting members on question j is $N^j \geq 2$ and we denote the vote from committee member $i^j \in \{1^j, \dots, N^j\}$ by $v_i^j \in \{1, 0\}$, where $v_i^j = 1$ is a “yes” vote and $v_i^j = 0$ is a “no” vote.

Voting procedure. The voting procedure on question j can be either sequential or simultaneous. If voting on question j is sequential, then the panelists vote sequentially and openly in an exogenously given order. The voting order is such that committee member i^j votes as the i 'th person (i.e. member 1^j votes first, then member 2^j , etc.). Let $h_i^j \equiv v_1^j, \dots, v_{i-1}^j$ denote the voting history observable for member i^j where $h_1^j = \{\emptyset\}$. If voting is simultaneous then the vote-history is empty for all committee members. We indicate by $\xi^j \in \{\textit{simultaneous}, \textit{sequential}\}$ the voting procedure for question j .

Signals. For every voting question j , each committee member i^j receives a private signal about the state. The signals are i.i.d. conditional on the state. The private signal S_i^j has the following linear state-contingent densities (following Cipriani and Guarino (2014)):

$$f^1(s_i^j | \theta^j = 1) = 1 + \tau(2s_i^j - 1)$$

$$f^0(s_i^j | \theta^j = 0) = 1 - \tau(2s_i^j - 1)$$


where $\tau \in (0, \infty)$. (See Figure 2.1)

The parameter τ is a measure of the level of strength in the experts' signals, where a larger τ means higher precision. In the case of the FDA's advisory boards, member i^j 's signal realization can be thought of as a process whereby member i^j considers the results and design of the clinical trials and draws on his/her personal experience, intuition, and analytical skills within a particular field (also allowing for randomness and misconceptions).

When $\tau \leq 1$ the support of the densities is $[0, 1]$. For $\tau > 1$, the support shrinks to $[\frac{\tau-1}{2\tau}, \frac{\tau-1+2\sqrt{\tau}}{2\tau}]$ for f^1 and $[\frac{\tau+1-2\sqrt{\tau}}{2\tau}, \frac{\tau+1}{2\tau}]$ for f^0 .⁹ The signals satisfy the monotone likelihood ratio property. For committee member i^j , the likelihood ratio after receiving signal s_i^j , $\frac{P(\theta^j=1|h_i^j, s_i^j)}{P(\theta^j=0|h_i^j, s_i^j)} = \frac{f^1(s_i^j|\theta^j=1)P(\theta^j=1|h_i^j)}{f^0(s_i^j|\theta^j=0)P(\theta^j=0|h_i^j)}$, is higher than the likelihood ratio before receiving the signal if $s_i^j > \frac{1}{2}$ and lower if $s_i^j < \frac{1}{2}$. In this way, a signal larger than one half is affirmative news and a signal lower than one half is negative news regarding the yes/no question at hand.¹⁰

⁹The intervals ensure that the density functions integrate to one.

¹⁰As explained in Cipriani and Guarino (2014), when $\tau \geq 1$ there are some signal realizations, s_i^j , that are only possible when the state is 1 (or 0), which then reveal the true state with certainty to member i^j . In fact, when $\tau \geq 1$ signal realizations higher than or equal to $\frac{\tau+1}{2\tau}$ are only possible when the state is 1 and signal realizations lower than or equal to $\frac{\tau-1}{2\tau}$ are only possible when the state is 0.

Figure 2.1: Probability density function of private signals $|\theta = 1$ 

Types and payoffs. Committee members want their vote to match the state and may require a higher or lower standard of proof (π) in order to vote yes. We define the payoffs for member i^j as follows:

$$Payoff = \begin{cases} 0, & \text{if } v_i^j = \theta \\ -\pi, & \text{if } v_i^j = 1 \text{ and } \theta^j = 0 \\ -(1 - \pi), & \text{if } v_i^j = 0 \text{ and } \theta^j = 1 \end{cases}$$

where $\pi \in (0, 1)$.

Voters can be of two types depending on the information they use to infer the state: *Herd* types ($t = H$) are Bayesian members who condition their vote on the common prior about the state, their private information, as well as on the history of votes of previous members along the sequence. In a variant of the model, we also consider herd types who follow a naïve updating rule (as in Eyster and Rabin, 2010). *Expressive* types ($t = E$) are myopic voters who disregard the information contained in previous votes and only condition on the common prior and their private information. In effect, the expressive type always stays true to his/her own assessment based solely on the prior and the private signal. Instead of assuming that the expressive type (myopically) ignores any information from the vote history we could provide the expressive type with an additional negative payoff when voting against his/her own judgment that is based solely on the common prior and the private signal. This additional payoff would correspond to a psychological cost from

not following one's own gut feeling (see e.g. Brennan and Pettit, 2000). These versions are equivalent in our setup and we have chosen the “myopic version” for simplicity.¹¹

Types are distributed independently across voters and voting questions and the probability that a committee member is the herd type is λ . A member's type is private information. When $\pi > 0.5$ ($\pi < 0.5$) we say that the experts are cautious (incautious) and it requires relatively more (less) affirmative evidence to vote “yes” (“no”). The parameters λ , π and τ are common knowledge. In the empirical analysis we consider heterogeneity in τ , λ and π across committee members with differing observable characteristics.

The motivation behind the structure of payoffs is the assumption that experts want to answer the FDA's questions correctly and to the best of their abilities.¹² Furthermore committing type 1 or type 2 voting errors may not weight the same, giving rise to a threshold of doubt that is different from one half.

We can formally characterize the voting behavior of member i^j when he/she is the expressive type. Let $\bar{s}_{i,t=E}^j$ indicate the cut-off signal such that $P(\theta^j = 1 | \bar{s}_{i,t=E}^j) = \pi$.¹³ Using Bayes' rule and the law of total probability:

$$P(\theta^j = 1 | \bar{s}_{i,t=E}^j) = \pi \iff \bar{s}_{i,t=E}^j = \frac{\mu_0^j - \pi}{2\tau(2\mu_0^j\pi - \pi - \mu_0^j)} + \frac{1}{2} \quad (2.1)$$

The cutoff signal from equation (1) characterizes voting behavior of the expressive type: vote yes if $s_i^j > \bar{s}_{i,t=E}^j$ and vote no if $s_i^j < \bar{s}_{i,t=E}^j$.¹⁴

¹¹While anecdotal, discussions with experts who have previously served on an FDA Advisory Committee during the era of sequential voting indicate that some experts made a point of deciding on their vote before the voting starts, while others were open to adjusting their vote during the voting procedure. One expert stated that her vote was determined before voting (S. Caprio, personal communication, February 10, 2020). Another expert explained that most of the time his mind was made up before casting his vote, but after hearing the potentially novel perspectives earlier on in the voting sequence there would be a 10% chance that he would change his vote based on these discussions, however any change in vote was just as likely to move against the “herd” as with it (T. Carpenter, personal communication, February 12, 2020).

¹²Discussions with experts indicate that experts are motivated to partake in committee meetings to gain insight into the drug review process, learn how to critique an application and to “give back” to the government or “be a good citizen”. Financial compensation is not an incentive to partake in meetings. With their vote, experts communicate what they believe is the correct answer to the question to the best of their abilities (S. Caprio, personal communication, February 10, 2020).

¹³Recall that the signals satisfy the monotone likelihood ratio property, see Duggan and Martinelli (2001) for how this translates into a voting rule characterized by a threshold crossing condition.

¹⁴Note that the cut-off signal from (1) is outside its support when $\tau < 1$ and $\mu_0^j \notin (\frac{1-\tau}{\frac{1}{\pi} + \frac{\tau}{\pi} - 2\tau}, \frac{\tau+1}{\frac{1}{\pi} - \frac{\tau}{\pi} + 2\tau})$. If $\mu_0^j \leq \frac{1-\tau}{\frac{1}{\pi} + \frac{\tau}{\pi} - 2\tau}$ it is optimal for member $i_{t=E}^j$ to vote no for any signal realization and when $\mu_0^j \geq \frac{\tau+1}{\frac{1}{\pi} - \frac{\tau}{\pi} + 2\tau}$ it is optimal for member $i_{t=E}^j$ to vote yes for any signal realization.


The herd type uses the history of votes to update his/her beliefs about θ^j . Let $\mu_i^j \equiv P(\theta^j = 1|h_i^j)$ indicate member i^j 's beliefs about the state after observing the preceding votes (not yet taking his/her own signal into account) and updating using Bayes' rule (or naïve updating). Optimal voting behavior for the herd type can be characterized by the cut-off signal, $\bar{s}_{i,t=H}^j$:

$$P(\theta^j = 1|\bar{s}_{i,t=H}^j, h_i^j) = \pi \iff \bar{s}_{i,t=H}^j = \frac{\mu_i^j - \pi}{2\tau(2\mu_i^j\pi - \pi - \mu_i^j)} + \frac{1}{2} \quad (2.2)$$

Member $i_{t=H}^j$ votes yes if $s_i^j > \bar{s}_{i,t=H}^j$ and no if $s_i^j < \bar{s}_{i,t=H}^j$.¹⁵ Under simultaneous voting $\bar{s}_{i,t=E}^j = \bar{s}_{i,t=H}^j$.

Note that increases in τ makes the value of the cut-off signal move towards one half (from below when $\mu > \pi$ and from above when $\mu < \pi$). Moreover, the cut-off signal becomes less sensitive to changes in μ and π when τ increases. As illustrated in Figure 2.2 the cut-off is decreasing in μ . In particular, the cut-off value is decreasing at an increasing (decreasing) rate when $\pi > 0.5$ ($\pi < 0.5$). In comparison to a committee member following the preponderance of evidence ($\pi = 0.5$), there is a larger range of signals for which an incautious expert will vote yes, and a smaller range of signals for which a cautious expert votes yes. This difference in the signal required to vote yes is most pronounced when the common prior is close to 0.5.

Figure 2.2: Cut-off signals (illustration with $\tau = 1$)



¹⁵The cut-off signal from (2) is outside its support when $\tau < 1$ and $\mu_i^j \notin (\frac{1-\tau}{\frac{1}{\pi} + \frac{\tau}{\pi} - 2\tau}, \frac{\tau+1}{\frac{1}{\pi} - \frac{\tau}{\pi} + 2\tau})$. When $\mu_i^j \leq \frac{1-\tau}{\frac{1}{\pi} + \frac{\tau}{\pi} - 2\tau}$ it is optimal for member $i_{t=H}^j$ to vote no for any signal realization and when $\mu_i^j \geq \frac{\tau+1}{\frac{1}{\pi} - \frac{\tau}{\pi} + 2\tau}$ it is optimal for member $i_{t=H}^j$ to vote yes for any signal realization.

2.2.1 Herd voting

We refer to the notion of local *herd voting* whenever a herd type is swayed by the history of votes.¹⁶ That is, when the herd type, following the cutoff rule, votes yes (no), whereas had he/she ignored the vote-history and followed the cutoff rule in (1), like the expressive type, he/she would have voted no (yes). In addition, we term it local *cascade voting* if it is optimal for committee member $i_{t=H}^j$ to vote, say yes, for *any* possible private signal realization and, at the same time, had member $i_{t=H}^j$ ignored the vote-history he/she would have voted no after observing $s_{i,t=H}^j$. Cascade voting is only possible when $\tau < 1$. The formal definitions are:

DEFINITION 1 (herd-voting): Provided that $\tau \geq 1$ or $\tau < 1$ and $\mu_i^j \in (\frac{1-\tau}{\frac{1}{\pi} + \frac{\tau}{\pi} - 2\tau}, \frac{\tau+1}{\frac{1}{\pi} - \frac{\tau}{\pi} + 2\tau})$, committee member $i_{t=H}^j$ engages in *herd-yes-voting* if $\bar{s}_{i,t=E}^j > s_{i,t=H}^j > \bar{s}_{i,t=H}^j$ and *herd-no-voting* if $\bar{s}_{i,t=E}^j < s_{i,t=H}^j < \bar{s}_{i,t=H}^j$.

DEFINITION 2 (cascade-voting): Given $\tau < 1$, committee member $i_{t=H}^j$ engages in *cascade-yes-voting* if $\mu_i^j \geq \frac{\tau+1}{\frac{1}{\pi} - \frac{\tau}{\pi} + 2\tau}$ and $s_{i,t=H}^j < \bar{s}_{i,t=E}^j$ and *cascade-no-voting* if $\mu_i^j \leq \frac{1-\tau}{\frac{1}{\pi} + \frac{\tau}{\pi} - 2\tau}$ and $s_{i,t=H}^j > \bar{s}_{i,t=E}^j$.

Unless $\tau < 1$ and $\mu_0^j \notin (\frac{1-\tau}{\frac{1}{\pi} + \frac{\tau}{\pi} - 2\tau}, \frac{\tau+1}{\frac{1}{\pi} - \frac{\tau}{\pi} + 2\tau})$ or $\mu_0^j = \mu_i^j$ (which is the case for simultaneous voting and the first voter under sequential voting) there will always be some private signal realizations whereby member $i_{t=H}^j$ engages in herd- or cascade voting.¹⁷ Cascade voting is not possible when $\tau \geq 1$, as there will always be some possible private signal realization that shifts the beliefs across the herd type's threshold and thus herd types never ignore their signal when $\tau \geq 1$. Even if τ and N are large the updated prior cannot hit 1 or 0. When $\mu_0^j < \mu_i^j$ the probability that member $i_{t=H}^j$ engages in herd-yes-voting can be computed as the probability that $s_{i,t=H}^j$ lands in the interval $[\bar{s}_{i,t=H}^j, \bar{s}_{i,t=E}^j]$ (NB the linearity of the state contingent signals allow for analytical solutions). Similarly when $\mu_0^j > \mu_i^j$ and herd-no-voting is possible. Thus, a larger discrepancy between μ_0^j and μ_i^j implies that the probability of a herd vote increases. However, for $\tau < 1$ and a sufficiently extreme μ_i^j the herd type will ignore his/her signal and cascade vote. This does not imply that all subsequent herd-types will cascade, as votes from expressive types may reverse the beliefs about θ^j . Even for $\tau < 1$ a positive measure of expressive types ensures that

¹⁶This definition of herd voting is similar to herd-buying and herd-selling in Cipriani and Guarino (2014).

¹⁷Note that when $\tau < 1$ and $\mu_0^j \notin (\frac{1-\tau}{\frac{1}{\pi} + \frac{\tau}{\pi} - 2\tau}, \frac{\tau+1}{\frac{1}{\pi} - \frac{\tau}{\pi} + 2\tau})$ there is no private signal realization that can overcome the prior and expressive- and herd types always vote the same. In this case, learning is stuck from the beginning.

learning never ceases and is unbounded.

2.2.2 The likelihood function

The likelihood function under Bayesian updating. To estimate our voting model we have to specify its likelihood function. Recall that the state is independently drawn across the voting questions $j = 1, \dots, J$ and the private signals $\{s_1^j, \dots, s_{N^j}^j\}$ are independent and identically distributed conditional on θ^j . Therefore, the events $v^j = v_1^j, \dots, v_{N^j}^j$ and $v^{k \neq j} = v_1^{k \neq j}, \dots, v_{N^{k \neq j}}^{k \neq j}$ are independent and the likelihood of a sequence of votes over the set of voting questions can be written as

$$P(\{v^j\}_{j=1}^J | \Phi) = \prod_{j=1}^J P(v^j | \Phi) \quad (2.3)$$

Where Φ is the vector of parameters $\{\mu_0, \tau, \lambda, \pi, \xi\}$. To demonstrate how to derive $P(v^j | \Phi)$ we consider sequential voting and any voting sequence $v^j = v_1^j, \dots, v_{N^j}^j$.

$$\begin{aligned} P(v^j | \Phi) &= \mu_0^j \prod_{i=1}^{N^j} P(v_i^j = 1 | \Phi, h_i^j, \theta^j = 1)^{v_i^j} P(v_i^j = 0 | \Phi, h_i^j, \theta^j = 1)^{1-v_i^j} \\ &+ (1 - \mu_0^j) \prod_{i=1}^{N^j} P(v_i^j = 1 | \Phi, h_i^j, \theta^j = 0)^{v_i^j} P(v_i^j = 0 | \Phi, h_i^j, \theta^j = 0)^{1-v_i^j} \end{aligned} \quad (2.4)$$

Conditional on the state and the voting history h_i^j , the individual votes are independent across the members. Thus, the vector of votes follows a mixture distribution, with mixing probability μ_0^j . The state specific voting probabilities are calculated as follows:¹⁸

$$\begin{aligned} P_{1,i} \equiv P(v_i^j = 1 | \Phi, h_i^j, \theta = 1) &= \lambda P(s_i^j > \bar{s}_{i,t=H}^j | \Phi, h_i^j, \theta = 1) \\ &+ (1 - \lambda) P(s_i^j > \bar{s}_{i,t=E}^j | \Phi, \theta = 1) \end{aligned}$$

¹⁸By the monotone likelihood ratio property of the signals it is ensured that $P_{1,i} \geq P_{0,i}$ and we can identify the state-specific voting probabilities. Identification in this setting is proven in a number of papers dealing with identification of mixture models such as Allman, Matias, and Rhodes (2009).

$$P_{0,i} \equiv P(v_i^j = 1 | \Phi, h_i^j, \theta = 0) = \lambda P(s_i^j > \bar{s}_{i,t=H}^j | \Phi, h_i^j, \theta = 0) \\ + (1 - \lambda) P(s_i^j > \bar{s}_{i,t=E}^j | \Phi, \theta = 0)$$

To solve for $\bar{s}_{i,t=H}^j$ we need member i 's updated probability that the state is good, given the history of votes. We do this recursively, whereby using Bayes rule

$$\mu_i^j \equiv P(\theta^j = 1 | h_i^j) = \frac{\mu_{i-1}^j P_{1,i-1}^{v_{i-1}} (1 - P_{1,i-1})^{1-v_{i-1}}}{\mu_{i-1}^j P_{1,i-1}^{v_{i-1}} (1 - P_{1,i-1})^{1-v_{i-1}} + (1 - \mu_{i-1}^j) P_{0,i-1}^{v_{i-1}} (1 - P_{0,i-1})^{1-v_{i-1}}}$$

for $i \geq 2$ and our base is $\mu_1^j = \mu_0^j$. If voting is simultaneous we can compute $P(v^j | \Phi)$ in the same manner, with the important difference that all the committee members vote with an empty vote-history.

The likelihood function under naïve updating. If herd types are “naïve”, as in Eyster and Rabin (2010), they (incorrectly) believe that each previous person’s action reflects solely that person’s private information. In our model this translates into herd types believing that everyone before them is an expressive type. This assumption requires an adjustment to the way we calculate the updated prior. We denote the updated prior under the assumption of naïve updating as $\mu_{N,i}^j$.

$$\mu_{N,i}^j \equiv P(\theta^j = 1 | h_i^j) = \frac{\mu_{N,i-1}^j P_{N,1,i-1}^{v_{i-1}} (1 - P_{N,1,i-1})^{1-v_{i-1}}}{\mu_{N,i-1}^j P_{N,1,i-1}^{v_{i-1}} (1 - P_{N,1,i-1})^{1-v_{i-1}} + (1 - \mu_{N,i-1}^j) P_{N,0,i-1}^{v_{i-1}} (1 - P_{N,0,i-1})^{1-v_{i-1}}}$$

where,

$$P_{N,1,i} \equiv P(v_i^j = 1 | \Phi, h_i^j, \theta = 1) = P(s_i^j > \bar{s}_{i,t=E}^j | \Phi, \theta = 1)$$

$$P_{N,0,i} \equiv P(v_i^j = 1 | \Phi, h_i^j, \theta = 0) = P(s_i^j > \bar{s}_{i,t=E}^j | \Phi, \theta = 0)$$

2.3 Data

2.3.1 FDA advisory committees

In the United States the producers of new drugs are required to win approval from the FDA in order to market their products. The review process gives the FDA the option to refer a matter of drug approval to one of its advisory committees. Around half of the drugs that the FDA reviews goes to a committee - typically those where the available data renders decision making particularly difficult or the drug or disease involved is controversial (Moffitt, 2010). Advisory committees are intended to provide the FDA with independent opinions and recommendations from outside

experts. Although the expert committees provide recommendations to the FDA, the FDA makes the final decisions and usually obtains additional clinical data and has discussions internally and with the sponsor company after the committee meetings are held.¹⁹ At the beginning of a meeting the FDA and the sponsor company present data from clinical trials and results regarding the risks and benefits of the drug or product under consideration. After the presentations the committee members deliberate and usually, after lengthy discussions, vote on one or more questions put forth by the FDA. These questions are generally scientific in nature and can involve a range of subjects, including the assessment of a drug or biological product's efficacy, safety, or overall approvability.

There are currently 18 different advisory committees under the Center for Drug Evaluation and Research. The committees are specialized on a particular disease or topic e.g. the Cardiovascular and Renal Drugs Advisory Committee or the Oncologic Drugs Advisory Committee. Each committee typically meets 1 to 4 times per year at the request of the FDA. On average a committee comprises around a dozen members. Each committee has a chair, who leads the meetings, several regular scientific members (serving 4-year terms), plus a qualified consumer and sometimes a patient representative.²⁰ Additional experts are usually added as temporary voting members. Like regular members, the invited temporary members have recognized expertise in the relevant field.²¹ Temporary members can be invited outside experts, members of the center's consultancy pool, or members of other advisory committees. Before each meeting all the committee members receive briefing material for preparation.

Following the Draft Guidance for FDA Advisory Committees of 2007 the voting procedure changed from sequential to simultaneous (electronic) voting.²² Under sequential voting the polling starts at one end of the committee table, at the chair's discretion, and continues in a stepwise fashion according to the seating plan of the

¹⁹As explained by the clinical FDA team leader at the meeting of the approval of Olodaterol: *"Before I close, I just wanted to mention the legal framework that gives the FDA the ability to hold advisory committees to ask for scientific advice and recommendations from experts in the field. As I noted previously, the FDA takes very seriously the advice of the committee. However, the Commissioner has sole discretion on actions taken with regard to drug approval, especially since there may be other issues, such as manufacturing, not discussed at the meeting, that impact approval decisions."*

²⁰As stated by the FDA, the role of the consumer representative is to represent the consumer perspective and serve as a liaison between the committee and interested consumers and consumer organizations. The consumer representatives are usually experts in the field like the regular committee members. Patient representatives have experience with the disease either as a patient or primary caregiver.

²¹Guidance for Industry, Advisory Committees: Implementing Section 120 of the Food and Drug Administration Modernization Act of 1997.

²²In the transition from sequential to electronic voting some committee meetings used voting by a show of hands. We exclude these meetings from our analysis.

meeting. The seating plan is jointly decided by the committee’s executive secretary and the chair. In Appendix 2.9.1 we provide evidence that, based on observable member characteristics, there are no clear patterns in the way that the committee members are seated (except for the chairperson). When a meeting has several voting questions, which is often the case, the chair usually alternates so that voting starts at each end of the committee table at every other voting question (on few occasions also starting from the middle and going clockwise or counter-clockwise around the table). The chair also votes and is seated in a central position. Often members are allowed to accompany their votes with comments, motivation, or provisos. Under simultaneous voting the members place their votes with electronic voting pads and after the votes have been locked in they sequentially go on the record and state what they voted and give comments. The order of announcements follows the seating plan of the meeting and the chair decides at which end of the table to begin. As with sequential voting, the chair usually alternates so that the announcements starts at each end of the committee table at every other voting question.

Around the same time the voting procedure changed, the FDA Amendments Act of 2007 was passed by Congress. Notably, the law established a limit on the number of committee members with financial conflicts.²³ Under USC Section 208, the FDA has authority to grant waivers to committee members who have potential financial conflicts when it is determined that the need for a particular individual’s services outweighs his or her potential financial conflict of interest.²⁴ One concern is whether the decline in conflict of interest (COI) waivers occurring around the time of the shift in voting procedure could have caused the changes in voting patterns that we see in the data. In our analysis we control for COI waivers and we are able to rule out that this change is driving our results.

2.3.2 Data collection and variables

Our data source is the full set of meeting transcripts that can be downloaded via www.fda.gov. The public records start in 1996 and we have data until June 2014. We consider committee meetings with one or more binary voting questions on the agenda and where the overall topic concerns approval of a drug or biological product. For sequential voting, this gives us 138 committee meetings and 375 binary voting questions with the full sequence of votes and for simultaneous voting it is 189 com-

²³The law also extended the authority to levy fees on companies applying for drug approvals, expanded clinical trial guidelines for pediatric drugs and enhanced the authorities to require post-approval studies.

²⁴Potential financial conflicts include investments, consulting, expert witness testimony, grants, patents and royalties, and primary employment in the sponsor company or its competitors.

mittee meetings and 438 voting questions. In total the data consists of 10,466 yes or no votes.²⁵ A “yes” vote is always associated with a favorable assessment of the relevant drug or biological product. This means that on a few occasions (negated questions) we reverse the votes. On average a committee comprises of 13 members, with a minimum of 5 members and a maximum of 28 members.

For each committee meeting, we indicate whether voting is simultaneous or sequential, the name of the advisory committee, type of application, and proposed trade name.²⁶ Our data covers 15 different topical committees. Applications can be a New Drug Application, a Biologic License Application, a supplemental New Drug Application, or a supplemental Biologic License Application.²⁷ Drug applications can be under “priority review”; a mechanism which seeks to expedite the review process for drugs that are expected to have a large impact on the treatment of a disease. Information on which applications are under priority review is obtained from the Drugs@FDA database.

We record the wording of the voting question and classify it depending on whether the question is about efficacy, safety, approval, or other (e.g. questions about methodology, dose, or labeling). We also report a score on the FDA reviewer(s) assessment of efficacy, safety, and approval. The FDA reviewer score is based on the FDA presentations and introductory remarks. Before each committee meeting the FDA’s review team analyses the efficacy and safety studies in question and prepares presentations to be held in front of the committee. These presentations take place before the voting stage and the reviewers are not members of the committee and do not vote. The FDA reviewer score on efficacy (1, 0, -1) reflects the review team’s conclusions regarding efficacy of the proposed drug or biological product. This revolves around the primary endpoints of the efficacy studies.²⁸ If the FDA reviewer(s) state that all the primary endpoints were met (usually with respect to p-values less than 0.05) in all the efficacy studies we code the efficacy score as 1. If the efficacy results are mixed or the FDA reviewer has major methodology concerns we code the efficacy score as 0. If the FDA reviewer concludes that the drug or product has no effect we code the efficacy score as -1.

²⁵The voting members also have the option to abstain, although they rarely do so. In our data, 1.2 percent of the sequential votes are abstentions and 2.7 percent of the simultaneous votes are abstentions. For simplicity we ignore abstentions in our analysis (the individual abstentions drop out as if they had not been placed).

²⁶In some cases, multiple drugs or products are considered on the same day and meetings are then split between morning and afternoon sessions.

²⁷Companies are allowed to make changes to drugs and biological products or their labels after they have been approved. To change a label, market a new dosage or strength, or change the way the treatment is manufacturing, a company must submit a supplemental application.

²⁸These studies are often placebo controlled trials, but can also be non-inferiority or superiority studies with respect to an already approved comparator.

The FDA reviewer score on safety (1, 0, -1) reflects the review team's conclusions regarding safety issues and adverse events. If the reviewer concludes that the safety profile is not worse than already approved products, or if the reviewer states that the safety studies reveal no significant safety concerns, we record a score of +1. On the other hand, if the FDA reviewer expresses serious safety concerns, also with respect to what is the standard for the relevant indication, we code the safety score as -1. If neither explicit positive or negative conclusions are drawn from the safety review we code the safety score as 0. Finally, the approval score simply adds the efficacy score to the safety score. Where the total score is 1 or higher, the FDA reviewer score for approval questions will be 1, similarly when it is -1 or lower, the score will be -1. In cases where there is no FDA reviewer assessment, we assign a score of 0.²⁹ Table 2.1 indicates the frequency of voting questions in each group.

At the voting question level we record the individual votes. For meetings with simultaneous voting, we order the votes in the way that the committee members announce them as they go on the record after the votes have been locked in. For sequential voting, the votes are entered according the order in which they are cast.³⁰ For each vote, we register the name of the voter, gender, educational background, whether he/she was granted a COI waiver for the given committee meeting, and whether the voter is a consumer representative, patient representative, regular or temporary member. All this information, except for gender, appears in the meeting transcripts or the summary minutes. In total, we observe the votes made by 1,378 unique voters. 12% of these voters (which account for 25% of votes) are present under both voting procedures. Table 2.2 indicates the frequency of voter characteristics, split by voting regime. A clear difference in committee composition after 2007 is the reduction in the number of members with a conflict of interest owing to the FDA Amendments Act. Another difference is the increased use of temporary committee members, as opposed to regular members, under simultaneous voting.

²⁹Questions that are classified as 'other' receive a score of 0. Meetings at the Cardiovascular and Renal Drugs Advisory Committee do not use FDA presentations until the end of 2005. For these meetings we insert reviewer scores of 0.

³⁰Under sequential voting, we have 20 voting questions with low data quality in the sense that the voting order was interrupted during voting (11), the voting question was modified after the first vote was given (5), or members did not precisely specify their vote (4). In the latter case, we include a yes or no vote depending on whether the member expressed himself/herself positively or negatively on the voting question. These voting questions are all included in our empirical analysis. None of our results or estimates significantly change depending on whether we include these observations or not.

Table 2.1: Voting question characteristics

	Sequential	Simultaneous	All
<i>FDA reviewer score</i>			
-1	37	26	63
0	287	293	580
1	51	119	170
<i>Priority review</i>			
Yes	123	71	194
No	252	367	619
<i>Question type</i>			
Efficacy	75	83	158
Other	153	120	273
Risk v. Benefit	126	179	305
Safety	21	56	77
<i>Application type</i>			
Non-supplementary	298	325	623
Supplementary	77	113	190
<i>Product category</i>			
Drug	349	336	685
Biologic	26	102	128
<i>Committee</i>			
Anesthetic and Life Support	1	13	14
Anti Infective	25	37	62
Anti Viral	16	16	32
Arthritis	13	31	44
Cardiovascular and Renal	76	31	107
Dermatologic and Ophthalmic	11	17	28
Endocrinologic and Metabolic	61	43	104
Gastrointestinal	24	44	68
Medical Imaging	2	3	5
Nonprescription	5	7	12
Oncologic	70	36	106
Peripheral and Central Nervous System	20	42	62
Psychopharmacologic	14	42	56
Pulmonary Allergy	16	49	65
Reproductive Health	21	27	48
Total	375	438	813

Table 2.2: Voter characteristics

	Sequential	Simultaneous	All
<i>Type</i>			
Regular member	2839 (66%)	2878 (47%)	5717 (54%)
Temporary member	1080 (25%)	2549 (41%)	3629 (35%)
Consumer representative	251 (6%)	372 (6%)	623 (6%)
Patient representative	128 (3%)	369 (6%)	497 (5%)
<i>Degree</i>			
Medical	2579 (60%)	3581 (58%)	6160 (59%)
PhD	1352 (31%)	1808 (29%)	3160 (30%)
Other degree	259 (6%)	395 (7%)	654 (6%)
No degree	108 (3%)	384 (6%)	492 (5%)
<i>Conflict of Interest</i>			
Yes	1011 (24%)	82 (1%)	1093 (10%)
No	3287 (76%)	6086 (99%)	9373 (90%)
<i>Gender</i>			
Male	2782 (65%)	4081 (66%)	6863 (66%)
Female	1516 (35%)	2087 (34%)	3603 (34%)
Total	4298	6168	10466

2.4 Descriptive analysis

In this section, we construct descriptive statistics and use reduced-form techniques to investigate changes in vote outcomes and voting behavior following the switch to simultaneous voting. Our findings are consistent with herd behavior: The main insights are that the probability of a unanimous vote outcome and the probability of a committee member voting the same as the person seated before him/her are significantly higher under sequential voting. We also find significant differences in voting behavior across committee members with differing observable characteristics. This further warrants the estimation of our model which allows for sequential learning and heterogeneity in the behavior of committee members.

2.4.1 Vote outcomes

To obtain precursory insights into voting outcomes under sequential and simultaneous voting, we construct the following variables at the voting-question level; 1) an indicator variable that takes the value 1 if the outcome of the vote is unanimous, 2) the size of the majority (as a percentage) and 3) the percentage of yes votes. De-


scriptive statistics for these variables are presented in Table 2.3. Under sequential voting 48% of vote outcomes are unanimous, this figure is 29% under simultaneous voting. Figure 2.3 shows the distribution of the percentage of yes votes for a given voting question. Clearly, there appears to be more agreement in votes under the sequential procedure.

Table 2.3: Descriptive statistics for vote outcomes

Variable	Obs.	Mean	Std. Dev.	Min	Max	Mean by Voting Rule		
						Sequential	Simultaneous	Difference
Unanimous	813	0.375	0.484	0	1	0.477	0.288	-0.19***
Majority size	813	0.853	0.154	0.5	1	0.871	0.838	-0.033***
Percent yes	813	0.637	0.36	0	1	0.664	0.614	-0.05*

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure 2.3: Agreement of votes



2.4.2 Voting behavior

In this subsection we run several regressions that model voting behavior at the individual level. The aim of this analysis is to convince the reader that there is a change in voting behavior following the switch to simultaneous voting, and that this change is in line with herd behavior under sequential voting. Furthermore, we use regression analysis to explore differences in voting behavior across different types of committee members.

In our analysis we focus on three outcome variables; 1) an indicator variable taking the value 1 if member i 's vote matches the preceding vote: $I(v_i^j = v_{i-1}^j)$ ³¹, 2) an indicator variable taking the value 1 if a member votes in line with the majority up until that point: $I(\text{WithMajority})$, and 3) an indicator variable taking the value 1 if a member votes "yes": $I(v_i^j = \text{yes})$. The first two measures aim to (crudely) reflect how individuals are influenced by previous votes. The regression specifications take the following form:

$$Pr[I_{ij} = 1] = \gamma_0 + \gamma_1 \text{Sequential}_j + X_{ij}\delta + \epsilon_{ij},$$

where I_{ij} is one of three binary dependent variables outlined above, Sequential_j is an indicator variable for a sequential voting procedure and X_{ij} is a vector of individual-level and question-level covariates.

At the question-level we control for committee size, FDA reviewer score, whether or not the drug is under priority review, if the application pertains to a biological product, if the application is a supplementary application, and the share of members in the committee with a conflict of interest waiver (as a percentage). We include fixed effects for the 15 topical committee categories and four question types (efficacy, risk vs. benefit, safety, other). At the individual-level we control for the member's position in the voting order ("seat"), voter type (i.e. temporary, regular, consumer representative and patient representative), level of education, gender and whether or not the committee member has a conflict of interest waiver. In order to provide an initial indication of heterogeneity in behavior across different types of committee members, we estimate specifications with interaction terms between individual characteristics and sequential voting.

Table 2.4 provides summary statistics for the three outcome variables and the non-categorical explanatory variables. There is substantial variation in the share of members with a COI waiver, both under sequential and simultaneous voting, which we use to separate the effect of a change in voting procedure from the effect of new regulations introduced around the same time which limited the number of members with a COI waiver.³²

³¹For simultaneous and sequential voting the order we use follows the order in which votes are announced.

³²The share of members with a COI waiver under sequential voting ranges from a minimum of 0 to a maximum of 1, with a standard deviation of 0.19 and a mean of 0.23. The share of members with a COI waiver under sequential voting ranges from a minimum of 0 to a maximum of 0.3, with a standard deviation of 0.04 and a mean of 0.013.

Table 2.4: Summary statistics

Variable	Frequency	Mean	Std. dev.	Min	Max
$I(v_i^j = v_{i-1}^j)$	9653	0.787	0.41	0	1
$I(WithMajority)$	9231	0.83	0.376	0	1
$I(v_i^j = yes)$	10466	0.633	0.482	0	1
Seat	10466	7.54	4.766	1	28
Size	10466	14.08	4.34	5	28
Share COI	10466	0.104	0.171	0	1

Table 2.5 presents the regression results using ordinary least squares.³³ Note that we exclude the first vote when using $I(v_i^j = v_{i-1}^j)$ and $I(WithMajority)$ as the dependent variable. We also exclude votes that follow a 50/50 outcome from regressions where $I(WithMajority)$ is the dependent variable. Controlling for observable vote and voter characteristics, on average, the probability that a member’s vote is the same as the previous vote is 6.3 percentage points higher under sequential voting (column 1). Further, a committee member is almost 5 percentage points more likely to vote with the present majority under sequential voting (column 3). The probability to vote yes is higher for sequential voting (column 5). These findings are consistent with earlier votes influencing later votes under sequential voting.

A number of additional results are worth pointing out. We find that regular members are significantly less likely to vote with the present majority under sequential voting. Patient representatives are also less likely to vote with present majority under sequential voting (column 4). Furthermore, we find that both variables which control for the change in COI waivers (*Share COI* and *COI waiver*) do not significantly affect the probability that a member’s vote matches the preceding vote or the probability of voting in line with the present majority. The positive sign on the interaction between *Seat* and *Sequential* in columns 2 and 4 provides some (weak) evidence that a committee member is more likely to be influenced by previous votes the later on in the sequence they vote.

In columns 5 and 6 we report which variables are correlated with a member’s probability to vote “yes”. An FDA reviewer score of -1 is negatively and significantly correlated with voting yes, whereas a score of 1 is positively and significantly correlated with voting yes. Priority drugs are also more likely to receive yes votes on average, as well as biological medications.³⁴ As might be expected, we find that consumer representatives are less likely to vote in favor of a drug all else constant, whereas patient representatives are more likely to vote “yes”.

³³Our results are robust to probit and logit specifications.

³⁴Later, we will use these features of vote questions to characterize the prior in our structural model.

The share of members in the committee with a conflict of interest waiver does not have a significant impact on the probability to vote in favor of the drug. There is some (weak) evidence that members with a conflict of interest may be more likely to vote in favor of a drug under simultaneous voting, but are relatively less likely to do so under sequential voting. It is not surprising that we do not find a strong effect of COI waivers on voting behavior. Previous studies using voting data from the FDA’s advisory committees have produced mixed results concerning the connection between voting and industry ties.³⁵

Overall, a descriptive analysis of the data suggests that there is a change in voting behavior following the switch to simultaneous voting and provides evidence consistent with earlier votes influencing later votes under sequential voting. However, we cannot rule out the possibility that voting outcomes may be driven by more extreme priors for the voting questions under the sequential procedure, which would lead to more agreement (regardless of the voting procedure). This is something that our reduced form model cannot directly control for. Moreover, even if we believe that, on average, the priors are similar for the sequential and simultaneous voting questions, a reduced-form model cannot speak to the mechanism that causes more agreement in the case of sequential voting or quantify the extent of “herd votes”. A structural approach is invaluable in this situation to explicitly incorporate the unobserved prior for each voting question, experts’ private information and differences in the cautiousness of experts in order to provide estimates that have a clear interpretation in the context of a model of herd behavior. Ultimately, using our model and the estimated parameters, we can also say something about which voting procedure leads to more efficient information aggregation.

³⁵Lurie and Zieve (2006) find a weak positive relationship between members’ votes for approval and financial ties. Ackerley et al. (2009) expand the data-set used in Lurie and Zieve (2006) and show a tendency for committee members to vote against their financial interests. Pham-Kanter (2014) finds that individuals with financial interests solely in the sponsoring firm are more likely to vote in favor of the sponsor than members with no financial ties. Cooper and Golec (2017) find that conflicts of interests are not significantly related to votes in FDA committees. Using a structural model, Camara and Kyle (2016) estimate each member’s skill and bias associated with financial ties to a drug’s sponsor or its competitors. Their results suggest that members with financial ties are more likely to vote in favor of both “good” and “bad” drugs. However, members with financial ties also have somewhat higher estimated ability, and hence are more favorable towards good drugs. Notably, these studies do not distinguish between sequential and simultaneous voting.

Table 2.5: Reduced-form results

	Dependent variable					
	(1) $I(v_t^j = v_{t-1}^j)$	(2) $I(v_t^j = v_{t-1}^j)$	(3) $I(WithMajority)$	(4) $I(WithMajority)$	(5) $I(v_t^j = yes)$	(6) $I(v_t^j = yes)$
Sequential	0.0633*** (0.0121)	0.0492* (0.0276)	0.0451*** (0.0114)	0.0652** (0.0258)	0.0572*** (0.0138)	0.0459 (0.0296)
Size/10	0.00657 (0.0134)	0.00436 (0.0134)	0.00198 (0.0125)	-0.000915 (0.0125)	-0.00444 (0.0144)	-0.00330 (0.0144)
FDA Reviewer Score = -1	-0.0971*** (0.0190)	-0.0976*** (0.0190)	-0.0649*** (0.0182)	-0.0671*** (0.0182)	-0.144*** (0.0199)	-0.144*** (0.0199)
FDA Reviewer Score = 1	0.0710*** (0.0113)	0.0711*** (0.0113)	0.0717*** (0.0104)	0.0718*** (0.0104)	0.183*** (0.0122)	0.183*** (0.0122)
Priority	-0.00421 (0.0107)	-0.00534 (0.0107)	-0.00926 (0.0101)	-0.0109 (0.0101)	0.119*** (0.0114)	0.121*** (0.0114)
Share COI	0.0302 (0.0334)	0.0318 (0.0337)	0.0564* (0.0308)	0.0535* (0.0311)	0.0633 (0.0419)	0.0655 (0.0420)
Supplementary	0.0144 (0.0107)	0.0138 (0.0108)	0.0144 (0.0101)	0.0152 (0.0101)	-0.00811 (0.0120)	-0.00929 (0.0120)
Biologic	0.0192 (0.0137)	0.0186 (0.0137)	0.0188 (0.0130)	0.0179 (0.0130)	0.0432*** (0.0156)	0.0434*** (0.0156)
Seat	-0.000200 (0.00104)	-0.00146 (0.00129)	0.000738 (0.000991)	-0.000155 (0.00124)	-0.000447 (0.00107)	0.000198 (0.00128)
Regular	-0.00792 (0.00942)	-0.00158 (0.0121)	0.00651 (0.00886)	0.0305*** (0.0115)	-0.0251** (0.0104)	-0.0267** (0.0129)
Patient Rep.	-0.0572*** (0.0220)	-0.0487* (0.0261)	-0.0312 (0.0204)	-0.0113 (0.0242)	0.0547** (0.0221)	0.0476* (0.0264)
Consumer Rep.	-0.0333* (0.0190)	-0.0481* (0.0250)	-0.0489*** (0.0186)	-0.0506** (0.0246)	-0.0859*** (0.0211)	-0.129*** (0.0268)
PhD	-0.00736 (0.00929)	-0.0156 (0.0125)	-0.0129 (0.00878)	-0.0193 (0.0118)	-0.0146 (0.0102)	-0.0232* (0.0133)
Male	0.00266 (0.00911)	0.0108 (0.0123)	-0.0107 (0.00850)	-0.00795 (0.0115)	-0.00286 (0.0100)	-0.00802 (0.0131)
COI Waiver	0.0220 (0.0155)	0.0293 (0.0436)	0.0128 (0.0144)	0.0166 (0.0419)	0.00230 (0.0175)	0.0849* (0.0498)
Seat X Seq.		0.00367* (0.00192)		0.00285 (0.00180)		-0.00200 (0.00208)
Regular X Seq.		-0.0160 (0.0189)		-0.0675*** (0.0174)		0.0105 (0.0214)
Patient Rep. X Seq.		-0.0328 (0.0489)		-0.0763* (0.0451)		0.0222 (0.0490)
Consumer Rep. X Seq.		0.0314 (0.0381)		-0.0106 (0.0372)		0.115*** (0.0433)
PhD X Seq.		0.0220 (0.0185)		0.0186 (0.0175)		0.0175 (0.0208)
Male X Seq.		-0.0195 (0.0182)		-0.00704 (0.0170)		0.0137 (0.0203)
COI X Seq.		-0.00590 (0.0459)		-0.00007 (0.0439)		-0.0920* (0.0524)
Question Type	yes	yes	yes	yes	yes	yes
Topical Committee	yes	yes	yes	yes	yes	yes
Constant	0.634*** (0.0418)	0.642*** (0.0431)	0.713*** (0.0394)	0.713*** (0.0408)	0.341*** (0.0420)	0.346*** (0.0432)
Observations	9,653	9,653	9,231	9,231	10,466	10,466
R-squared	0.032	0.032	0.026	0.028	0.082	0.083

Notes: OLS regression. Standard errors in parentheses are robust. ** * $p < 0.01$, * * $p < 0.05$, * $p < 0.1$.

2.5 Estimation and identification

In this section, we describe the specification of the prior and the estimation of our model. We then provide some intuition on which variation in the data allows us to identify the parameters of the model.

To make the model tractable, we place a parametric restriction on the prior.³⁶ We

³⁶This is in line with previous literature e.g. Iaryczower and Shum (2012) and Camara and

allow the prior μ_0^j , the common belief that the correct answer to the voting question is yes, to depend parametrically on characteristics of the voting question captured by X_j (e.g. FDA reviewer score) via the following logit formulation:

$$\mu_0^j(X_j; \beta) = \frac{\exp(X_j' \beta)}{1 + \exp(X_j' \beta)} \in (0, 1) \quad (2.5)$$

This specification implies that there are certain observable vote characteristics that have an effect on the probability of the state being “yes” or “no”. The state for each voting question j is determined by these characteristics and a question-specific unobserved shock term drawn from a standard logistic distribution. In Appendix 2.9.2 we elaborate on how this specification can accommodate correlation in the true state for voting questions that are part of the same meeting.

To incorporate heterogeneity in herd behavior, caution and expertise we specify λ , π and τ as a function of categorical voter characteristics including category of committee member (Regular, Temporary, Consumer Representative or Patient Representative), gender, whether or not the committee member has a PhD and whether or not the member has a conflict of interest. Specifically,

$$\lambda_i = \gamma_{reg.} + \gamma_{temp.} + \gamma_{cons.} + \gamma_{pat.} + \gamma_{phd} + \gamma_{COI} + \gamma_{male} \quad (2.6)$$

$$\pi_i = \alpha_{reg.} + \alpha_{temp.} + \alpha_{cons.} + \alpha_{pat.} + \alpha_{phd} + \alpha_{COI} + \alpha_{male} \quad (2.7)$$

$$\tau_i = \eta_{reg.} + \eta_{temp.} + \eta_{cons.} + \eta_{pat.} + \eta_{phd} + \eta_{COI} + \eta_{male} \quad (2.8)$$

Note that this specification defines four different intercepts for regular members, temporary members, patient and consumer representatives which can shift depending on gender, education and COI status. The parameters to be estimated are the vectors β , γ , α and η . To recover the parameter estimates, we maximize the likelihood function (3) directly using the full dataset of 10,466 individual votes. To find the parameters that minimize the negative log-likelihood function, we used both the quasi-Newton algorithm for unconstrained optimization and the Nelder-Mead simplex direct search algorithm.³⁷ Standard errors are calculated by taking the square root of the diagonal elements of the inverse estimated Hessian of the likelihood function at the solution.

Dupius (2014).

³⁷Both methods converge to the same parameter estimates. We do not need to impose constraints on our parameters to obtain reasonable estimates. Results are robust to different starting values. We conducted a Monte Carlo exercise using a simulated dataset and verified that our procedure yields reasonably precise, unbiased estimates of the parameters of the model.

Regarding the model’s identification, the degree of agreement in votes, the proportion of yes votes, the exact order of votes under sequential voting, and differences in how certain voters behave across questions are important. The prior, μ_0 , is identified by the proportion of yes votes at the vote question level. All experts tend to receive higher private signals when the state is “yes” and thereby a high prior will induce many yes votes. The level of cautiousness, π , is identified by variation in how many “yes” vs. “no” votes are cast across questions with different priors. A voter’s π follows them across questions with different priors. Intuitively, if a committee member (or group of committee members) is particularly cautious there will be less variation in their votes and they will vote “no” more often.

The identification of τ is characterized by the degree of agreement in votes. An increase in τ is distinguishable from an increase in μ_0 , or a decrease in π , as the increase in τ may cause a higher concentration of both yes and no votes across multiple voting questions. At the individual level, a committee member with high expertise will be more likely to receive a signal that squares with the true state. Thus, members whose votes are typically in line with the majority, particularly under simultaneous voting, will be estimated as having a high expertise. The probability that a committee member is a herd type, λ , is identified by the sequence of votes and differences in how voters with certain characteristics react to the history of votes that they observe. Given μ_0 , π , and τ we can identify how likely it is that a voter is a herd type based on how they vote in response to previous voters. Intuitively, if we see a long sequence of yes votes, and thereafter a voter who votes no, this voter is more likely to be an expressive type.

While in theory, all parameters can be identified on the basis of sequential data alone, in practice, with a limited sample of voting questions under sequential voting and a limited number of committee members voting on each question, access to simultaneous data is crucial. Intuitively, both an increase in the precision of information and the share of herd types will create more agreement in votes, hence without access to simultaneous data identification of λ hinges on the exact order of votes under non-unanimous outcomes. Simultaneous data is used to get a grip on τ (as well as μ_0 and π) when there are no herd effects at play, which allows us to better separate the effect of these parameters from λ .³⁸ Effectively, the results are similar to what would be obtained if we applied a two-step procedure whereby first simultaneous data is used to estimate τ , π and the determinants of the common prior, and then plugging these estimates into the model, sequential data is used to

³⁸This is confirmed by simulations where, with a limited number of voting questions, the accuracy of all parameter estimates is improved by using both simulated sequential and simultaneous data. Increasing committee size (the length of the sequence) also improves the estimates the parameters.

recover λ .

2.6 Estimation results

In this section, we present our results for the voting model introduced in Section 2.2. We first present the estimates of the model parameters and then discuss the frequency of herd voting.

2.6.1 Estimates

Table 2.6 presents the estimates and the standard deviations for the parameters of the model. In the baseline model we do not include voter characteristics. We find that on average the probability that any given committee member is a herd type is close to one half (48% for Bayesian version and 52% for naïve version). Put differently, on average half of committee members take into consideration the vote history when placing their vote under sequential voting. We should bear in mind that this does not mean that half of the committee members actually herd, that is, change their vote from what it would have been if ignoring the vote history. Herd types will only change their vote if the information inferred from the previous votes is sufficiently strong and opposite to their private information. It's not hard to find examples from the FDA transcripts where panelists are open about paying attention to the previous votes, without necessarily being swayed by them.³⁹ We discuss our approach to quantifying herd votes in the following subsection.

We estimate the model allowing for heterogeneity across voters with different characteristics. We present the estimates for both the model which assumes standard Bayesian updating and the model which assumes naïve updating. The results are qualitatively similar, and hence we will focus our discussion on the model which assumes herd types are fully Bayesian.

The average proportion of herd types masks differences in λ across voters with certain characteristics. The results indicate that temporary committee members are the most susceptible to herd behavior. Regular committee members and patients representatives are less likely to be herd types. Members with a conflict of interest are more likely to be herd types. Finally, the results suggest that gender and whether or not an expert has a PhD has little impact on λ .

The low proportion of herd types among regular members compared to temporary members could be due to the fact that these members regularly participate

³⁹For example, voting after one "yes" and four "no" votes it's Dr. Martino's turn: "Having struggled and heard all of you struggles, my answer is going to be no." From the meeting of the Oncologic Drugs Advisory Committee in March 2006 on the drug Gemzar.

in advisory meetings. This process would make their self-esteem and esteem as expert panelists fairly settled and potentially turn some regular “herd” members into expressive voters (see Brennan and Pettit (2000) on the economics of esteem). Another explanation may be career concerns which are typically thought to be higher for less experienced members (see Hansen, McMahon, and Prat, 2018). Scharfstein and Stein (1990) show that agents with career concerns unsure of their expertise tend to herd on the same action, thereby avoiding being the only one to take an incorrect decision. Hong, Kubik, Solomon (2000) compare the behavior of inexperienced and experienced equity analysts and find that inexperienced analysts deviate less from consensus forecasts. They interpret this finding as being consistent with career-concern-motivated herding theories. With respect to social conformity,⁴⁰ some experiments demonstrate that subjects are more likely to conform when grouped with strangers as opposed to friends (McKelvey and Kerr, 1988). In Appendix 2.9.4 we expand on our finding in Section 2.4.2 that regular members are less likely to vote with the present majority under sequential voting. We find that members who are attending a meeting for the first time are significantly more likely to vote with the majority under sequential voting. We are also able to rule out that a committee member’s age is driving the result. This lends support to the conjecture that more frequent attendance of meetings can reduce the extent to which committee members are influenced by previous votes.

Our results indicate that gender does not impact the likelihood to be influenced by previous votes in the context of FDA Advisory Committees. This result is in contrast to previous findings from the social conformity literature. The results are mixed, however, Eagly and Carli (1981) performed a meta-analysis of 148 studies of influenceability and find that women are more persuadable and more conforming than men in group pressure situations that involve surveillance.

On average, we find that committee members are cautious and would prefer to incorrectly vote to reject a good drug than incorrectly vote to approve a bad drug. On average, members vote yes if they believe that the probability that the true state is “yes”, given all information, is at least 58%. We also find heterogeneity in the cautiousness of members. Consumer representatives are the most cautious as one might expect. On the other hand, patient representatives vote yes if the probability that the state is 1 is greater than 49%. Unlike with the tendency to

⁴⁰In a social conformity framework, individuals are influenced by observing others’ actions, not because of information revealed about an underlying state, but due to social dynamics, see Asch (1951). For discussions on the distinction between informational and social influence see Deutsch and Gerard (1955), Shiller (1995), and Bernheim and Exley (2015). In a field experiment conducted with a financial brokerage, Bursztyn et al. (2014) implement a novel design to separately identify these two channels of influence.

herd, the standard of proof required to vote in favor of a drug does not differ between temporary and regular committee members.

On average the precision of information for the FDA advisory committees is quite high ($\tau = 1.26$). This implies that the probability that a member gets an incorrect signal (i.e. a signal < 0.5 when the state is 1, or a signal > 0.5 when the state is 0) is 20%. Precision of information varies across committee members. Regular members have the most accurate private information whereas the signals of consumer and patient representatives are less likely to align with the true state.

We use information on the FDA reviewer score, whether the medication is a drug or biological product and whether the drug is under priority review to characterize the prior. The inclusion of these characteristics is motivated by which variables are significant in the reduced-form analysis. We find that our estimates are robust to the inclusion of more vote characteristics in the specification of the prior, e.g. inclusion of voting question type. We select this specification as it is parsimonious, while at same time providing a good sense of the range of μ_0 . In Appendix 2.9.3 results are presented for a model where we estimate a common prior for each committee meeting.

Estimates of the average common prior μ_0 for categories of voting questions can be computed using β . For example, the average prior for voting questions relating to a drug under priority review with an FDA reviewer score of 1 can be computed as $\frac{\exp(1.67)}{1+\exp(1.67)} = 0.84$. Estimated average priors range from 0.44 to 0.84.

Table 2.6: Estimation results

Parameter	Baseline				Parameter	Heterogeneity			
	Bayesian		Naïve			Bayesian		Naïve	
	Estimate	SE	Estimate	SE		Estimate	SE	Estimate	SE
λ	0.48	0.04	0.52	0.04	$\gamma_{reg.}$	0.38	0.06	0.46	0.06
					$\gamma_{temp.}$	0.55	0.07	0.67	0.07
					$\gamma_{cons.}$	0.45	0.10	0.51	0.09
					$\gamma_{pat.}$	0.32	0.15	0.35	0.14
					γ_{phd}	0.06	0.06	-0.01	0.06
					γ_{COI}	0.16	0.06	0.11	0.06
					γ_{male}	0.00	0.06	-0.04	0.05
π	0.58	0.02	0.59	0.03	$\alpha_{reg.}$	0.58	0.03	0.58	0.03
					$\alpha_{temp.}$	0.56	0.03	0.57	0.03
					$\alpha_{cons.}$	0.67	0.03	0.69	0.03
					$\alpha_{pat.}$	0.49	0.03	0.50	0.03
					α_{phd}	0.04	0.01	0.03	0.01
					α_{COI}	-0.01	0.02	-0.01	0.02
					α_{male}	0.00	0.01	-0.01	0.01
τ	1.26	0.03	1.28	0.03	$\eta_{reg.}$	1.38	0.07	1.36	0.07
					$\eta_{temp.}$	1.24	0.07	1.24	0.07
					$\eta_{cons.}$	1.09	0.10	1.14	0.11
					$\eta_{pat.}$	1.12	0.10	1.12	0.10
					η_{phd}	-0.09	0.06	-0.08	0.06
					η_{COI}	-0.10	0.11	-0.06	0.10
					η_{male}	0.03	0.06	0.06	0.06
$\beta_{score=1,priority,drug}$	1.64	0.17	1.61	0.17	$\beta_{score=1,priority,drug}$	1.67	0.17	1.58	0.17
$\beta_{score=0,priority,drug}$	0.75	0.11	0.65	0.13	$\beta_{score=0,priority,drug}$	0.73	0.11	0.63	0.12
$\beta_{score=-1,priority,drug}$	-0.26	0.22	-0.24	0.22	$\beta_{score=-1,priority,drug}$	-0.25	0.22	-0.24	0.23
$\beta_{score=1,drug}$	1.09	0.12	1.12	0.12	$\beta_{score=1,drug}$	1.10	0.12	1.13	0.12
$\beta_{score=0,drug}$	0.42	0.09	0.48	0.09	$\beta_{score=0,drug}$	0.43	0.09	0.48	0.09
$\beta_{score=-1,drug}$	0.00	0.15	0.05	0.15	$\beta_{score=-1,drug}$	0.06	0.15	0.11	0.15
$\beta_{score=1,biologic}$	0.22	0.17	0.24	0.17	$\beta_{score=1,biologic}$	0.24	0.17	0.26	0.17
$\beta_{score=0,biologic}$	0.46	0.11	0.51	0.11	$\beta_{score=0,biologic}$	0.47	0.11	0.52	0.11
$\beta_{score=-1,biologic}$	0.33	0.25	0.35	0.25	$\beta_{score=-1,biologic}$	0.34	0.25	0.35	0.25

2.6.2 Herd votes

We now investigate the frequency of actual herd votes. Using our structural model and the estimated parameters we are able to construct a simulated dataset of votes under sequential and simultaneous voting. By comparing an individual's simulated vote under sequential and simultaneous voting we can directly observe which votes are herd votes.

We simulate a dataset of 1,000 vote questions which are voted on by a committee of 13 members (the average committee size) under both a sequential and simultaneous

procedure. We use parameter values in line with our baseline estimates and repeat the procedure for three different values of the common prior; 0.5, 0.65 and 0.8.⁴¹ Specifically, the procedure is as follows:

1. For each voting question draw an error term ϵ_j from the standard logistic distribution
2. Assign each voting question a state as follows: $\theta^j = 1$ if $y_j^* > 0$ and 0 otherwise, where $y_j^* = \beta + \epsilon_j$
3. Given θ^j and τ , draw private signals for each voter, for each voting question
4. With probability λ assign each voter to be a herd type
5. Given the assigned signals, μ_0 , τ and π simulate voting under the simultaneous voting rule
6. Given the assigned signals, μ_0 , τ , π , types and λ simulate voting under the sequential voting rule

Thus, we have the same voters voting on the same voting question, once under simultaneous rule, and once under the sequential rule. If an individual votes differently under sequential voting, this is counted as a herd vote. We also compare the proportion of unanimous outcomes and the average size of the majority across the simulated datasets.

Table 2.7 presents key statistics calculated using the simulated data based on the Bayesian model. For a balanced prior, we find that 18.4% of herd types actually herd and thus 8.9% of all sequential votes are “herd votes”. Across all three common priors, the proportion of herd votes is fairly stable at around 8.5%. In each case the proportion of unanimous vote outcomes increases under sequential voting. The number of unanimous outcomes also increases as the common prior moves away from 0.5 and there is less uncertainty. With a prior of 0.8, we find that 42.6% of vote outcome are unanimous under sequential voting, in comparison to 28.5% under simultaneous voting. Thus, voting data simulated according to our model is able to re-produce patterns similar to those appearing in the real data.⁴²

Table 2.8 presents the same statistics for the case of naïve updating. Here, we find that roughly 9% of all sequential votes are herd votes and that unanimous outcomes occur more frequently. This is in part driven by a higher estimate of τ in the naïve model, which explains why the share of unanimous outcomes is also higher under simultaneous voting. In general, the patterns are similar to those observed under the assumption of Bayesian updating.

⁴¹The corresponding β 's are 0, 0.62 and 1.39.

⁴²Although not directly comparable to the simulated dataset, recall that in the actual data, 47.7% of vote outcomes under sequential voting are unanimous and 28.8% under simultaneous voting.

Table 2.7: Simulated outcomes with Bayesian updating

	$\mu_0=0.5$		$\mu_0=0.65$		$\mu_0=0.8$	
	Seq.	Sim.	Seq.	Sim.	Seq.	Sim.
Proportion of herd votes	8.9%	NA	8.2%	NA	8.4%	NA
Proportion of unanimous outcomes	25%	9%	28.7%	10.1%	42.6%	28.5%
Average size of the majority	0.88	0.8	0.89	0.82	0.88	0.82

Notes: $\tau = 1.26$, $\pi = 0.58$, $\lambda = 0.48$, $N=13$ **Table 2.8: Simulated outcomes with naïve updating**

	$\mu_0=0.5$		$\mu_0=0.65$		$\mu_0=0.8$	
	Seq.	Sim.	Seq.	Sim.	Seq.	Sim.
Proportion of herd votes	9.8%	NA	9.4%	NA	6.7%	NA
Proportion of unanimous outcomes	27.5%	10.6%	31.3%	10.1%	54.6%	37.4%
Average size of the majority	0.88	0.8	0.89	0.82	0.91	0.87

Notes: $\tau = 1.28$, $\pi = 0.59$, $\lambda = 0.52$, $N=13$

2.7 Information aggregation

In this section, we consider information aggregation through voting and how a switch to simultaneous voting affects the quality of the committee’s overall assessment. We first do this by calculating the probability that the committee makes the correct assessment about a medication under each voting procedure. Thereafter, we match our sample of voting questions with final FDA approval decisions. On the basis of our model and estimates we find that the committee is more likely to make the correct assessment under simultaneous voting. Furthermore, we find that the FDA’s final approval decisions are more likely to be in alignment with the committee’s assessment under simultaneous voting.

2.7.1 Probability of a correct assessment

In this subsection, we calculate the probability that a committee of a given size makes the correct assessment about a medication under each voting procedure. We define the committee’s overall assessment as being favorable when the updated beliefs about the state being “yes” after everyone has voted is greater than one half. Beliefs about the state being “yes” after everyone has voted are computed in the same way that committee members update beliefs about the state (taking all the parameter values of the model into consideration). We consider both Bayesian and naïve updating.

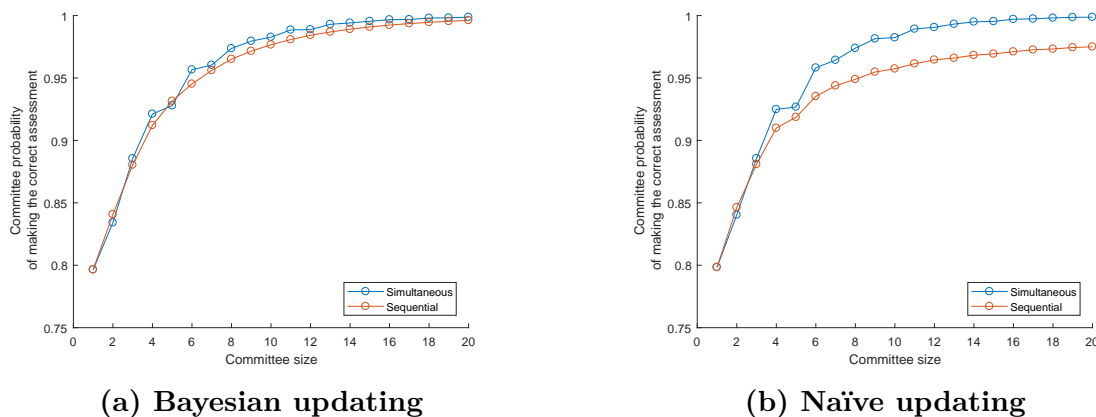
Let $A^j \in \{0, 1\}$ represent the committee's overall assessment for voting question j where $A^j = 1$ represents a favorable assessment of the drug. For each committee size N we can calculate the probability that the assessment is correct (i.e. a favorable assessment of a good drug or a negative assessment of a bad drug) by calculating the number of instances where $A^j = \theta^j$ across all possible voting profiles and weighting appropriately by the probability of the state and the voting profile conditional on the state.

Let v_N^j be a specific sequence of votes with N voters. The set V^j comprises of all possible voting sequences with N voters. We denote the updated beliefs about the state after N experts have voted by $\mu_{N+1}^j(v_N^j) \equiv P(\theta^j = 1|v_N^j)$. We assume $A^j(v_N^j) = 1 \iff \mu_{N+1}^j(v_N^j) > 0.5$. Let $I(\mu_{N+1}^j(v_N^j) > 0.5)$ be an indicator variable that takes on the value one if $\mu_{N+1}^j(v_N^j) > 0.5$. The probability that the committee's overall assessment is correct can be computed as follows:

$$\begin{aligned} Pr(A^j(v_N^j) = \theta^j) &= \mu_0^j \sum_{v^j \in V^j} Pr(v_N^j | \Phi, \theta = 1) \times I(\mu_{N+1}^j(v_N^j) > 0.5) \\ &\quad + (1 - \mu_0^j) \sum_{v^j \in V^j} Pr(v_N^j | \Phi, \theta = 0) \times (1 - I(\mu_{N+1}^j(v_N^j) > 0.5)) \end{aligned}$$

Figure 2.4 illustrates how the total probability of making the correct assessment changes with committee size.⁴³ For both sequential and simultaneous voting, the committee is more likely to make the correct assessment as more members are added. The informational gain of adding more committee members displays diminishing returns. After around 12 committee members, there are only small gains from expanding committee size.

Figure 2.4: Information aggregation



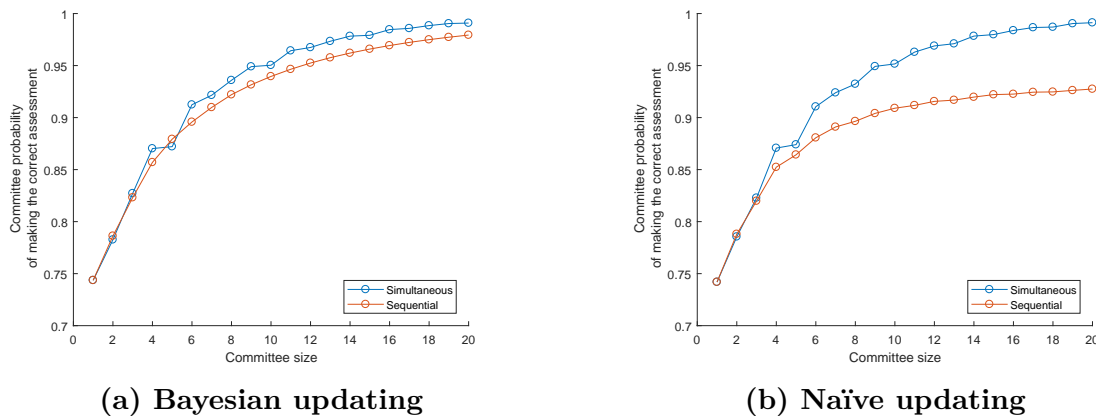
⁴³As in the previous section, we set the parameters to their average values and use a common prior of 0.5.

In all but one case (where $N=5$ and updating is Bayesian), simultaneous voting outperforms sequential. Committee members can herd in both the correct and incorrect direction. Occasional incorrect local herds under sequential voting drive our result that simultaneous voting outperforms sequential voting on average. Incorrect local herds tend to occur in the event that voters with incorrect signals vote early on, causing subsequent herd types to get the wrong idea about the state. This can take the updated beliefs about the state being “yes” above (below) 0.5 when the true state is actually “no” (“yes”). Clearly, local herds also form in the correct direction. In such cases updated beliefs can exceed (fall below) 0.5 faster under sequential voting when the true state is “yes” (“no”). When a local herd forms in the correct direction, as committee size increases, updated beliefs after everyone has voted under simultaneous voting “catch up” with those under sequential voting. However, updated beliefs formed by incorrect local herds, tend to remain divergent from beliefs under simultaneous voting for a wider range of committee sizes and may not be overturned even for very large committees, leading to incorrect assessments.

The deleterious informational consequences of herding are more prominent if belief updating is naïve. In the naïve version, herd types take the preceding votes at face value, thus belief updating may accelerate faster and it is also harder to overturn beliefs that get off on the wrong foot. In the Bayesian version herd types take into account that a preceding vote which breaks a trend (e.g. a no vote after four yes votes) may come from another herd type, which signals that this voter had a very strong signal against the trend. This positive effect is absent in the naïve version where the herd type believes that the “breaking” vote comes from an expressive type with an “average” signal against the trend. In Appendix 2.9.5 we present the results of a counterfactual scenario in which we assume that all members are herd types.

To assess the effect of a counterfactual decrease in information precision we set τ equal to 1 and hold all other parameters constant. This corresponds to a situation in which experts receive incorrect signals 25% of the time. As illustrated in Figure 2.5, the effect of a decrease in information precision is to exacerbate the negative consequences of herd behavior regardless of which type of updating is applied. Finally, we should mention that there are possible parameter values of our model for which sequential voting outperforms simultaneous voting. This stands in contrast to most of the herding literature building on a binary signal structure.⁴⁴

⁴⁴An exception is Wiedman (2014) who shows that sequential voting may increase information transmission compared to simultaneous voting in a model with binary signals and competent versus incompetent experts.

Figure 2.5: Information aggregation with a lower precision of information

2.7.2 Votes and FDA approvals

In this subsection we match our sample of voting questions with FDA approval decisions. The observations are matched on the basis of application number. The sample is restricted to risk vs. benefit questions relating to non-supplementary applications as this is most reflective of the committee’s opinion on whether or not a new medication should be approved. Information on drug approvals is taken from the FDA Orange Book and information on biological product approvals comes from the FDA CDER List of Licensed Biological Products, both of which are publicly available online. In total our sample comprises of 241 voting questions. As previously mentioned, the FDA committees do not operate under a majority rule and one should bear in mind that the agency often conducts additional investigations after the committee meetings have taken place and before a final decision is made.

We show simple but powerful descriptive statistics on the link between the committee vote and final approval decisions under the different voting regimes. Table 2.9 illustrates how often FDA approval decisions are in line with the committee’s recommendation when the committee’s vote is unanimous. We find that the FDA is more likely to go against the committee’s unanimous recommendation when voting is sequential. In 71% of cases where the committee unanimously votes yes, the FDA approves the medication under sequential voting. This is in comparison to 94% of cases where the committee unanimously votes yes under simultaneous voting. The difference is significant. There is also suggestive evidence that the FDA is more likely to reject a drug conditional on a unanimous no under simultaneous voting. However owing to the small number of observations, this difference is not significant.

The pattern we observe could be driven by the FDA placing more weight on committee votes under simultaneous voting, or individual votes being more aligned

with the true state, which the FDA is more accurately able to assess at the point in time when the final decision is made, or a combination. Both mechanisms are in line with our main finding that the assessment of the committee is more likely to match the true state under simultaneous voting.

Table 2.9: Alignment of FDA decisions and unanimous voting outcomes

Variable	Obs.	Mean	Std. Dev	Mean by Voting Rule		
				Sequential	Simultaneous	Difference
Pr(Approval Unanimous yes)	69	0.83	0.05	0.71	0.94	-0.23***
Pr(Rejection Unanimous no)	24	0.79	0.08	0.77	0.82	-0.05

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

2.8 Conclusion

In this paper, we use data from FDA committees to estimate the extent and importance of herd behavior under sequential polling of expert recommendations. We find that around half of the committee members are susceptible to herd behavior and might go against their private assessment if the votes from previous experts indicate otherwise. On average, around 9 percent of the sequential votes are actual herd-votes. Temporary committee members invited on an ad hoc basis are more apt towards herding than regular (standing) members.

Considering the consequences of herding on information aggregation we find that simultaneous voting lowers the probability that the committee’s overall assessment is misaligned with the state. Further, we find that simultaneous voting performs substantially better than sequential voting if information is imprecise and/or when members do not take into consideration that preceding members may have herded (i.e. members are “naïve”). We bolster these findings with descriptive evidence that indicates that committee votes are more aligned with the FDA’s final approval decisions under simultaneous voting.

We believe that our analysis and results are relevant beyond scientific advisory committees: In situations where people are assembled to give their advice or vote on certain questions; from corporate advisory boards and hiring committees to roll call voting by elected representatives in community councils and commissions. The main policy implication of this study is to follow the example of the FDA and substitute sequential voting with simultaneous (electronic) voting.

In future work we intend to explore the mechanisms behind our main findings in more depth. In this regard, diving into the comments and pre-vote discussions from the FDA meeting transcripts using sentiment analysis and text mining could be

fruitful (see e.g. the method and techniques applied in Hansen, McMahon, and Prat (2018) concerning policy makers' deliberations in the FOMC). Directly incorporating and testing for reputational or career concerns would be another way in which to extend the present research.

Another array of future research that we plan to undertake is to study the FDA's decision making process, taking into account both the advisory committees and the clinical study by the sponsor company, and the quality of drug approval decisions. To this end, Li and Agha (2015) study the success of peer-review NIH panels and are able to track the merits of NIH funded research (likewise for unfunded NIH applications in Li (2017)). The quality of FDA decisions may be gauged by e.g. considering drug withdrawals, market reactions, or decisions made by other agencies like the European Medicines Agency.

2.9 Appendix

2.9.1 Seating order

In this appendix we present descriptive statistics concerning the order in which votes are announced (“seating order”) under sequential voting. In Table 2.10 we split the sample of votes by voter characteristics and calculate summary statistics on seat number for each category. It is evident that different types of committee members, e.g. those with a PhD or male members, are very spread out in terms of where they sit in the sequence. The average seat position for all groups is between 6 and 7 with a standard deviation of around 4.5. There are no clear clusters of certain types of committee members at start or end of the voting sequence. In Table 2.11, we focus on the ten most frequent voters. It is evident that even individual committee members are spread out in terms of where they sit and do not always get placed in the same position.

Table 2.10: Summary statistics for seat no. by voter characteristics

Voter Characteristic	Frequency	Mean Seat	Std. dev.	Min	Max
Regular	2839	6.78	4.38	1	28
Temporary	1080	6.91	4.80	1	21
Consumer rep.	251	7.11	4.25	1	22
Patient rep.	128	7.40	4.73	1	21
PhD	1352	6.37	4.10	1	25
Male	2782	6.63	4.47	1	28
COI waiver	1011	6.48	3.69	1	24
All Sequential Votes	4298	6.85	4.50	1	28

Table 2.11: Summary statistics for seat no. for 10 most frequent voters

Voter ID	Frequency	Mean Seat	Std. dev.	Min	Max
570	62	5.10	2.29	1	11
985	46	8.91	3.44	2	15
1284	46	5.67	3.10	1	11
539	44	5.93	3.01	1	11
1141	43	7.28	5.01	1	18
1051	42	5.79	3.90	1	14
791	40	6.00	3.75	2	15
813	40	7.30	2.74	3	13
981	40	6.95	4.01	2	17
848	38	4.87	3.08	1	11

2.9.2 Correlation in vote questions part of the same meeting

In our empirical implementation, we assume that there is an underlying (latent) variable y_j^* that determines the true state for a voting question j . We specify $y_j^* = X_j'\beta + \epsilon_j$ such that:

$$\theta^j = \begin{cases} 1 & \text{for } y_j^* > 0 \\ 0 & \text{for } y_j^* \leq 0 \end{cases}$$

We further assume ϵ_j has a standard logistic distribution which implies $Pr(\theta^j = 1|X) = \frac{\exp(X_j'\beta)}{1+\exp(X_j'\beta)}$.

Given that we can have multiple voting questions taking place on the same day and relating to the same drug (on average 2-3 voting questions per meeting), we may expect correlation in the true state across voting questions that are part of the same meeting.

We can allow for such correlation by introducing an error term at the meeting level ϵ_m . We now specify the underlying (latent) variable as $y_j^* = X_j'\beta + ((1-\sigma)\epsilon_j + \sigma\epsilon_m)$. If σ is 0 this collapses to the previous specification, as σ increases there is more correlation in the true state for voting questions that are part of the same meeting. Given that private signals are state dependent, there will also be more correlation in private signals within a meeting. Assuming ϵ_m follows the standard logistic distribution, the new error term $((1-\sigma)\epsilon_j + \sigma\epsilon_m)$ also follows the standard logistic distribution. Hence we still have $Pr(\theta^j = 1|X) = \frac{\exp(X_j'\beta)}{1+\exp(X_j'\beta)}$.

In our simulations we implement such an error structure by grouping questions into sets of four and drawing the same ϵ_m for the set. We use a σ of 0.7. We find that our estimates of β , λ , and τ are unbiased and very similar to what they were before introducing the correlation.

2.9.3 Alternative specification of the prior

In order to illustrate the robustness of our results to different specifications of the prior, we implement a model where we estimate a prior for each committee meeting i.e. we estimate 327 priors corresponding to the 327 committee meetings in our dataset. In order to ease estimation of a model with so many parameters, we use constrained minimization. We constrain λ , π and all elements in the vector of common priors μ_0 to be between 0 and 1. We constrain τ to be between 0 and 5.

Table 2.12 provides the main parameter estimates of this model for both the Bayesian and naïve version. As one might expect, estimating a prior for each meetings allows the prior to vary much more across voting questions and so reduces our estimates of τ and λ , but not by too much. We now find that 32-39% of members are herd types.

Table 2.12: Estimation results

Parameter	Baseline			
	Bayesian		Naïve	
	Estimate	SE	Estimate	SE
λ	0.32	0.04	0.39	0.05
π	0.6	0.02	0.6	0.02
τ	1.17	0.04	1.15	0.04

2.9.4 Additional regressions

In this appendix we perform additional regressions to shed light on the finding that regular committee members are less likely to vote with present majority under sequential voting. In column 1 we include an indicator variable that takes the value 1 when a committee member is attending a meeting for the first time. We find that first timers are 3 percentage points more likely to vote with the present majority under sequential voting. In column 2 we include the count of meetings that a member has attended and its interaction with *Sequential*, the effect has the expected sign, the more meetings a member attends the less likely he/she is to vote with majority under sequential voting, however the coefficient is not significant. We collect information on the age of members from Healthgrades.com. We are not able to find the age for all committee members in our sample. In column 3, we re-run the regression in Table 2.13 column 3 controlling for age where possible. Our result that regular members are less likely to vote with present majority persists when we control for age.


Table 2.13: Additional reduced-form results

	<i>I(WithMajority)</i>	<i>I(WithMajority)</i>	<i>I(WithMajority)</i>
Sequential	0.0312** (0.0130)	0.0559*** (0.0138)	0.0690*** (0.0193)
First Meeting	-0.0116 (0.0122)		
First Meeting X Seq.	0.0328** (0.0162)		
No. of Meetings		0.00111 (0.00144)	
No. of Meetings X Seq.		-0.00391 (0.00270)	
Regular	0.00662 (0.00955)	0.00670 (0.00906)	0.0149 (0.0139)
Regular X Seq.			-0.0446** (0.0203)
Age			0.00003 (0.000572)
Consumer Rep.	-0.0495*** (0.0188)	-0.0491*** (0.0186)	-0.144*** (0.0434)
Patient Rep.	-0.0319 (0.0204)	-0.0313 (0.0204)	-0.319*** (0.106)
Seat	0.000738 (0.000991)	0.000750 (0.000991)	0.00135 (0.00121)
PhD	-0.0128 (0.00879)	-0.0129 (0.00888)	-0.00368 (0.0120)
Male	-0.0112 (0.00851)	-0.0117 (0.00856)	-0.00250 (0.0105)
COI Waiver	0.0135 (0.0144)	0.0143 (0.0144)	0.0228 (0.0170)
All question-level controls	yes	yes	yes
Constant	0.719*** (0.0399)	0.711*** (0.0395)	0.751*** (0.0600)
Observations	9,231	9,231	6,046
R-squared	0.027	0.027	0.030


Notes: OLS regression. Standard errors in parentheses are robust. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

2.9.5 Counterfactual with 100% herd types


In this appendix we present the results of a counterfactual scenario in which we assume that all committee members are herd types. Interestingly when all members are Bayesian updaters and information precision is high, see Figure 2.6 (a), sequential voting marginally outperforms simultaneous voting for committees that are smaller than 12. On the other hand, if all members are naïve updaters, sequential voting performs very poorly in comparison to simultaneous voting (Figure 2.6 (b)). If information is less precise, a higher proportion of herd types is detrimental for information aggregation as shown by Figure 2.7.




(a) Bayesian updating



(b) Naïve updating

Figure 2.6: Information aggregation with $\tau = 1.26$ and $\lambda=1$ 

(a) Bayesian updating



(b) Naïve updating

Figure 2.7: Information aggregation with $\tau = 1$ and $\lambda=1$

Chapter 3

Common Ownership and Market Entry: Evidence from the Pharmaceutical Industry¹

Chapter Abstract

Common ownership - where two firms are partially owned by the same investor - and its impact on product markets has recently drawn attention. This chapter focuses on implications for entry. We consider the entry decisions of generic pharmaceutical firms into drug markets opened up by the end of regulatory protection in the US. We find robust evidence that common ownership affects entry: a one-standard-deviation increase in common ownership decreases the probability of individual entry by 15-18%, whereas a one-standard-deviation increase in market-level common ownership decreases the total number of entrants by as much as 15% in that market.

3.1 Introduction

Johnson & Johnson, Pfizer, Abbott Laboratories, Perrigo and Allergan, some of the largest brand and generic companies in US pharmaceutical markets, had the same top two shareholders in 2005: BlackRock and Vanguard (Thomson Reuters Global Ownership Database, 2015). BlackRock and Vanguard are amongst the world's largest institutional investors.² Investors' holdings in multiple firms give

¹An earlier version of this chapter is published in the DIW Discussion Paper Series as: Newham, M., Seldeslachts, J. and Banal-Estañol, A. Common Ownership and Market Entry: Evidence from the Pharmaceutical Industry. *DIW Berlin Discussion Paper No. 1738*. We thank Hendrik Meder and Manuel Gigena for sharing data with us; and Anna Sama and especially Julian Hidalgo for their excellent research assistance.

²Institutional investors such as Blackrock and Vanguard manage other people's money by buying and controlling equity in companies.

rise to what is known as “common ownership.” A controversial question is if, and if so in which way, firms’ strategic decisions are altered by the presence of common ownership.³

This article investigates the effect of common ownership on one of the most important strategic decisions firms make: market entry. Specifically, we analyze generic firms’ entry decisions into pharmaceutical markets opened up by the end of regulatory protection. Monopolized markets are a vital source of revenue for brand firms. Brand revenues can decline by as much as 90% following generic entry (Branstetter et al., 2016). Moreover, losses to the brand and gains to the generics are highly asymmetric. According to one estimate, brand firms value deterring entry at about \$4.3 billion on average (Jacobo-Rubio et al., 2020). In contrast, generic firms value the right to enter at about \$204.3 million. This asymmetry in brand losses and generic gains is confirmed in our sample too, as we will later show: with the event of generic entry, not only brand revenues but also total market revenues decrease as compared to the case of no generic entry. Thus, generic entry decisions may crucially depend on whether the owners of generic firms also have an interest in the brand firms.

We investigate whether a higher level of common ownership between potential generic entrants and the market’s incumbent brand reduces the likelihood of entry, both at the level of an individual potential entrant and at the market level for all potential entrants. To do so we combine patent and drug approval data from the US Food and Drug Administration’s (FDA) Orange Book with ownership data of publicly listed pharmaceutical companies from the Thomson Reuters Global Ownership Database. The US pharmaceutical industry is an attractive industry for studying entry because; (i) pharmaceutical markets are well defined, (ii) one can identify clear entry windows and (iii) US health care expenditure as a percentage of GDP is among the highest in the world and generic medicines are crucial to keeping down healthcare costs. Indeed, promoting generic entry has become an important goal for the FDA in recent years, and there are still several hundred off-patent branded drugs which do not face any generic competition (FDA, 2019).

We first present a theoretical framework to understand the effects of common ownership between an incumbent and the potential entrants. We model, in particular, the simultaneous entry decisions of a set of generic firms, where we take into consideration both individual profits and levels of common ownership with the incumbent brand firm. We find that, in response to a higher level of common own-

³Rather than maximizing their own value, commonly-owned firms may maximize shareholders’ *portfolio* values. See Backus et al. (2019) and Schmalz (2018) for reviews of the available academic evidence.

ership with the brand, an individual generic should find entry less profitable, for any belief concerning the entry decisions of the other generics. We then solve for equilibrium and show that there will be fewer entrants in markets characterized by higher levels of common ownership with the brand.

Thereafter we empirically test and corroborate the proposition that higher common ownership reduces individual generic entry. This result is robust to several measures of common ownership, different econometric methods, different definitions of the set of potential entrants, different time-horizons for the decision-making process and different fixed effects. Our regressions include the controls used in previous literature including pre-entry market size, molecular substitutes, entrant experience and the presence of an authorized generic. The average effect is large: a one-standard-deviation increase in common ownership decreases the probability of entry by that generic firm by 15-18%. Furthermore, our results indicate a non-linear impact of common ownership on entry, where high levels have a much stronger impact than low levels. Our results hold if we instrument common ownership with stock market index membership.

Going to the market level, we find that a one-standard-deviation increase in overall common ownership between the brand and all potential entrants decreases the total number of generics in that market by 13-15%. We find that common ownership also delays generic entry and increases the probability that the brand will face zero competition from generic entrants.

Common ownership is a pervasive feature not only of pharmaceutical companies, but of many industries in the US as well as in Europe (Fichtner et al., 2017; Seldeslachts et al., 2017). Although large institutional investors may own 5-8% of a single company, this is often enough to position them as a top investor with privileged access to the firms' management (Malenko and Shen, 2016). There is growing evidence that institutional investors engage in active discussions with companies' board and management with a view to influence the companies' strategies (e.g., McCahery, 2016; Fichtner and Garcia-Bernardo, 2017).⁴ However, institutional investors need not actively influence companies to have an impact on firm strategies. They may employ "selective omission"; encouraging actions that increase both firm value and portfolio profits and remaining silent when this is not the case (Hemphill and Kahan, 2019). They may have an effect by crowding out and occasionally voting against other investors (Antón et al., 2020). Moreover, firms that are largely owned by shareholders who also have sizeable stakes in competitors might just simply act in these shareholders' interest, which leads them –rather than maximizing their own

⁴We present some anecdotal evidence in Appendix 3.8.1 that investors confirm this view, both in general and for pharmaceutical markets.

profits— to maximize the return of their shareholders’ portfolios (Azar, 2017). In our theoretical framework, we present different measures of common ownership that reflect these different channels on how common ownership might influence firms’ behavior.

The ongoing concentration of ownership in the hands of a few large investors and the corresponding escalation in common ownership is unprecedented. Dubbed “an economic blockbuster” and “the major new antitrust challenge of our time,” common ownership is undoubtedly an important, new topic in competition economics (Elhauge, 2016; Posner et al., 2017).⁵ But empirical research on the topic is still in its infancy. For a large sample of US public firms, He and Huang (2017) find that common ownership by institutional investors facilitates explicit forms of product market coordination which in turn improves innovation productivity and operating profitability. Azar et al. (2018) provide empirical evidence that common ownership in the airline industry is linked to higher prices. The results of these studies have been subject to ongoing debate (see e.g., O’Brien and Waehrer, 2017). There is, however, a resounding agreement that more research is required to understand the implications of common ownership (Patel, 2017; OECD, 2017).

This article is the first to directly consider the influence of common ownership on market entry. Whereas pricing decisions are typically made on a regular basis by specialized pricing teams, market entry is a one-off decision with substantial consequences for the firm. Another advantage of the current article over previous empirical studies is the fact that we do not only look at market-level common ownership, but also at ownership links at the pair level; between individual generics and the incumbent brand.

The rest of the article is organized as follows. Section 3.2 provides a literature overview of entry in pharmaceutical markets and common ownership. Section 3.3 introduces the theoretical framework. Section 3.4 presents the data. Section 3.5 shows the empirical analysis and results of the effect of common ownership on individual entry. Section 3.6 deals with the effect of common ownership on market outcomes. Section 3.7 concludes. We include appendices on (i) anecdotal evidence on how institutional investors influence firms’ decisions, (ii) data construction, (iii) empirical robustness checks and (iv) mathematical proofs.

⁵The issue has also received significant media attention and instigated public debate; see e.g. The Economist (2015), The New York Times (2016), Handelsblatt Global (2016) and OECD (2017).

3.2 Literature

We separately discuss the most relevant articles on the entry decisions of generic firms in pharmaceutical markets and common ownership.

Generic entry. Several articles have considered the determinants of generic entry decisions in off-patent drug markets, i.e., markets where the patent of the brand company has expired. A common finding from this literature is that generic entry increases with the size of the branded drug's market prior to the loss of patent protection, where market size is commonly measured as brand-generated revenues (Scott Morton, 1999, 2000; Hudson, 2000; Saha et al., 2006; Moreno-Torres et al., 2008; Appelt, 2015).

Scott Morton (1999) considers other aspects of generic entry decisions in US pharmaceutical markets. She finds that generic firms are more likely to enter markets in which they have previous experience in drug form, therapy class or ingredient. Kyle (2006) and Appelt (2015) similarly confirm the importance of generic firm characteristics. Scott Morton (1999, 2000) also highlights the role of the characteristics of the drugs. Appelt (2015) examines the impact of authorized generics, i.e., the distribution and marketing of the brand product under a generic label through an authorized generic distributor (typically just before the loss of the patent). She finds that authorized generic entry has no significant effect on the likelihood of 'independent' generic entry.

Scott Morton (2002) reviews how direct ownership links between the brand firm and a generic firm influences the likelihood of generic entry. She finds that generics owned by the original innovator (i.e., the brand company) are less likely to enter the market. Helland and Seabury (2016) investigate the link between Paragraph IV challenges, settlements and entry. They find that a Paragraph IV challenge increases generic entry, although a settlement effectively reverses the effect. Hovenkamp and Lemus, finally, (2017) confirm that settlements after Paragraph IV challenges cause generics to stay out of the market.

Common ownership. In terms of theoretical work, beginning with Rubinstein and Yaari (1983) and Rotemberg (1984), a number of authors have remarked that shareholder diversification can lead firms to internalize the externalities they impose on rivals; see Schmalz (2018) for a full overview. These models show that common ownership of competitors reduces incentives to compete as the gains of aggressive competition to one firm come at the expense of other firms in the investors' portfolio. Consequently, common ownership is predicted to lead to higher prices and boost

industry profits. On the other hand, Lopez and Vives (2019) find that cost-reducing R&D investment with spillovers in a Cournot oligopoly may lead to higher welfare when there is higher common ownership.

Previous empirical studies on common ownership have mainly focused on price effects. In an empirical study focusing on the US airline industry, Azar et al. (2018) use the modified Herfindahl-Hirschman index (MHHI), developed by O'Brien and Salop (2000), which provides a measure of the extent of common ownership at the market level. They find that ticket prices are about 3-12% higher than would be the case under separate ownership. Azar et al. (2016) focus on the US banking industry, extending the MHHI to take into account cross-ownership –the degree of which banks own shares in each other– and find that common and cross-ownership are positively correlated with banking fees. Further studies that look at the effect of common ownership on prices in airlines (Kennedy et al, 2017) and banking (Gramlich and Grundl, 2017), using different methodologies, measures and samples, find mixed effects. Scott Morton and Boller (2020) study the effect of common ownership on future expected profits as captured by stock prices. They find that increases in common ownership cause increases in stock returns, consistent with a hypothesis that common ownership raises profit.

Xie and Genakos (2019) find that institutional investors' common holdings between US generic and brand companies increase the likelihood of settlement agreements after generic companies have disputed the brand's patent validity through a Paragraph IV challenge, which is the section of the Hatch-Waxman act under which generic entrants dispute pharmaceutical patents. Additionally, through positive brands' abnormal stock market returns around the settlement date, they conclude that these settlements have facilitated collusion between brand and generics. Their study, thus, is complementary to this article as it showcases a plausible channel of how entry can be deterred.

Some recent empirical studies highlight the positive effects that common ownership can have on innovation and vertical relations. Antón et al. (2017) examine how common ownership affects R&D investments and innovation output. Geng et al. (2017) find that vertical common-ownership links can mitigate hold-up problems arising from patent complementarities, which in turn is correlated with more innovation. Cici et al. (2015) and Freeman (2019) find that common ownership between vertically connected firms can help strengthen business relationships.

Finally, there is a small but growing body of literature in corporate finance that investigates channels through which institutional investors might have an impact on governance, policies and strategic decisions of firms (e.g., Aghion et al., 2013; Brav et al., 2018). Appel et al. (2016) find that passive mutual funds have a significant

and positive impact on several aspects of corporate governance (board composition, anti-takeover provisions and unequal voting rights). Their evidence suggests that a key mechanism by which theseF investors exert their influence is through their large voting blocks.

Furthermore, institutional investors state that they have a fiduciary duty to weigh on firms' decisions and do so through informal meetings with management and through voting at annual general meetings by the employment, for example, of proxy voters such as Institutional Shareholder Services (ISS) (Malenko and Shen, 2016). Boone and White (2015) examine the effects of institutional ownership on firm transparency and information production. They find that higher institutional ownership is associated with greater management disclosure; resulting in lower informational asymmetries. In line with the findings of Appel et al. (2016), they discover that indexing investors have the highest influence on information production.

3.3 Theoretical framework

We now present a simple framework to understand the effects of common ownership on market entry. We model, in particular, the decisions of a set of symmetric generic firms that have the possibility to produce a drug and simultaneously enter a market currently dominated by the product of a brand firm. We first analyse how an increase in the levels of common ownership between a brand and a focal generic affects this individual generic's entry decision, taking as given the decisions of other generics. We also determine the overall number of market entrants in equilibrium, as a function of the level of common ownership of all the generics with the brand. Finally, we propose several measures of common ownership between brand and generic firms.

Common ownership and market entry. Consider N (≥ 1) symmetric generic firms that can simultaneously enter the market of a brand firm b .⁶ We focus first on the decision of a *focal* generic g as a function of its beliefs about the entry decisions of the other generics (i.e. the best-reply function).

Denote by p_k the probability, assigned by this (risk-neutral) focal generic, to the event that a number k of the *other* generic firms enter the market, where $k = 0, \dots, N - 1$ and $\sum_{k=0}^{N-1} p_k = 1$. Denote by $\pi_g^k (> 0)$ the focal generic's profits in a market that also includes k other generic firms (and thus the market contains in

⁶Our main empirical specification specifies an entry window of 6 quarters, where entry decisions should be considered as simultaneous. This is because the entire application process for generic drugs takes about 6 quarters on average, during which period information on ANDA's received by the FDA is kept secret until approval.

total $k + 2$ firms, when also counting the brand firm). Profits π_g^k may include fixed costs of entry, and are thus net of these entry costs. Denote by $\Delta\pi_b^k (< 0)$ the loss in profits of the brand firm b due to an increase from k to $k + 1$ in the total number of generic entrants in the market.

Let us make the following assumptions. Naturally, we assume that generic competition reduces individual generic profits, i.e. π_g^k is decreasing in k , and that the change in the brand firm's profit loss decreases with the number of entrants, i.e. $|\Delta\pi_b^k|$ is decreasing in k . We also posit that the gains of the generic are lower than the losses of the brand, as generic competition reduces a brand firm's profits significantly (Branstetter et al., 2015). In other words, although generic firm profits increase, $\pi_g^k > 0$, joint profits decrease with entry, $\pi_g^k + \Delta\pi_b^k < 0$, independently of the number of entrants.⁷

Common ownership between the generic and the brand makes the entry decision non-trivial. Indeed, shareholders of the generic that also own shares in the brand may care about the reduction of joint profits. As a result, the decision-makers of g may consider the “net” gains from entry when deciding whether to enter, thereby taking also into account the reduction of joint profits. Formally, denoting by δ the weight the decision-makers of g place on the joint profits, rather than on individual generic firm profits, g shall enter the market if $\Pi_g \geq 0$, i.e., the expected net gains from entry are positive, where

$$\Pi_g(p_0, \dots, p_{N-1}, \delta) \equiv \sum_{k=0}^{N-1} p_k [(1 - \delta)\pi_g^k + \delta(\pi_g^k + \Delta\pi_b^k)]. \quad (3.1)$$

An increase in common ownership between g and b will naturally increase δ .⁸ In the absence of common ownership between g and b , $\delta = 0$, and therefore generic g should place no weight on joint profits and entry will occur, as $\pi_g^k > 0$ for any number k of other generic entrants. At the other extreme, in the case where common ownership is so high that joint profits are as important as individual generic profits, $\delta = 1$, entry will not occur, as $\pi_g^k + \Delta\pi_b^k < 0$ for any k . More in general, the gains from entry of a generic g decrease in its level of common ownership with the brand, as

$$\partial\Pi_g(p_0, \dots, p_{N-1}, \delta)/\partial\delta = \sum_{k=0}^{N-1} p_k \Delta\pi_b^k < 0 \text{ for any } p_0, \dots, p_{N-1}.$$

⁷Both assumptions are consistent with the evidence we provide in Figure 3 on the relationship between number of entrants on the one hand, and brand and total market revenues on the other hand. That means that the business stealing effects caused by generic entry on the brand firm are larger than any market expansion effect. This should hold true for markets with low demand elasticity of which pharmaceutical markets are a primary example (Duggan and Scott Morton, 2010).

⁸Thus, δ can be viewed as our “measure of common ownership.” We discuss possible common ownership measures at the end of the theory section.

The next proposition summarizes this result, on the individual decisions of a focal generic, as well as the pure-strategy equilibrium entry decisions of all the N generics. The characterization of the equilibrium entry decisions, as a function of symmetric level δ of common ownership of each of the generics and the brand, can be found in Appendix 3.8.3.

Proposition 1. *An increase in the (bilateral) level of common ownership between an individual generic and the brand reduces entry by this generic. The number of entrants in equilibrium decreases as the symmetric (market) level of common ownership between all the generics and the brand increases.*

Common ownership measures. We now propose several measures of common ownership that aim to capture how common investors' interests in the two firms affect the weight that the generic firm places on joint rather than on individual firm profits. We posit that shareholdings in the brand provide common investors with *incentives* to steer decisions towards joint profits and shareholdings in the generic provide investors with the *ability* to influence such decisions (Posner et al., 2017). The main difference between our various measures is how incentives and ability to influence decisions are taken into account.

We propose three approaches that to some extent cover different channels of investor influence. In broad terms, the first approach has some flavour of investors actively engaging with decision-making, as it parametrizes the effect of shareholders' interests into an index of decision-making influence. The second approach assumes that the generic firm's decision-makers are aware of and take shareholders' portfolio interests into account, and hence investors do not need to explicitly engage. The third approach posits that the top common investors in generic and brand firms (according to their rank) have the strongest incentives and ability to influence the firm's decisions. Whereas the first and second approach make use of the *size* of investors' shareholding, the third approach constructs measures of common ownership that rely on the *ranking* of common investors in terms of their holdings.

PRODUCTION FUNCTION APPROACH. This approach assumes that there exists a "production function" that transforms each common investor's shareholdings in the brand and generic firms (inputs) into a "joint profit steering index" (output). This index increases with the size of the investor's shareholdings in the brand because this increases her concerns about the reduction of joint profits (incentives). The index also increases with the size of the investor's shareholdings in the generic because larger shareholdings naturally imply a greater capacity to influence the generic firm's decisions (ability). For simplification, assuming perfect coordination among common

investors, the weight that the generic firm places on joint, rather than on individual, profits is the sum of joint profit steering indices across common investors.⁹ In formal terms, there exists a function f such that

$$\delta(g, b) = \sum_j f(\gamma_{jg}, \gamma_{jb}), \quad (3.2)$$

where γ_{jg} and γ_{jb} are the shareholdings of a common shareholder j that owns shares in the generic and brand, respectively. The marginal effect of each of the two arguments of f should be positive, but there could additionally be some degree of complementarity between the two. In other words, the marginal effect of incentives may be larger if the ability is higher, and vice versa. We apply two extreme production function examples (Gilje et al., 2019). First, the two shareholdings can be “perfect substitutes,” i.e., $f(\gamma_{jg}, \gamma_{jb}) = (\gamma_{jg} + \gamma_{jb})/2$, and thus:

$$\delta_S(g, b) \equiv \sum_j (\gamma_{jg} + \gamma_{jb})/2. \quad (3.3)$$

Second, the two shareholdings can be “perfect complements,” i.e., $f(\gamma_{jg}, \gamma_{jb}) = \min\{\gamma_{jg}, \gamma_{jb}\}$, and thus:

$$\delta_C(g, b) \equiv \sum_j \min\{\gamma_{jg}, \gamma_{jb}\}. \quad (3.4)$$

Note that both functions are assumed to be symmetric with respect to the two inputs. Moreover, the scale is such that both measures range between zero and one. In both cases, the generic firm will place no weight on joint profits ($\delta(g, b) = 0$) if there are no common shareholders, and a necessary condition for full-weight on joint profits ($\delta(g, b) = 1$) is that all shareholders are common.

In terms of interpretation, perfect substitutes (equation (3.3)) assumes that the marginal effect of an increase in incentives does not depend on ability, and vice versa. On the other hand, perfect complements (equation (3.4)) assumes that incentives require ability, and vice versa. This means that the perfect substitutes measure does not penalize unequal shareholdings in the two firms whereas the perfect complements measure does. For instance, a shareholder that owns 5% of the shares of one firm and 15% of the other would have the same contribution to $\delta(g, b)$ as someone that owns 10% in both firms when applying the perfect substitutes measure but only half of it when applying the perfect complements measure. Of course, both measures are similar if the relative holdings of all common investors in the brand and generic are similar.

⁹We assume thus that common investors coordinate their collective decision making. This assumption makes sense if common owners have similar interests. For example, a case study of a shareholder vote at the company DuPont indicates how common investors can group together and use the power of their large voting block to implement their objectives (Schmalz, 2015).

WEIGHTED SUM OF INTERESTS APPROACH. This approach, following O’Brien and Salop (2000), assumes that the decision makers of the generic firm maximize a weighted sum of the interests of all investors in the firm, where (i) the interests of an investor are given by her shareholdings in the two firms and (ii) the weights are given by the investor’s degree of control of the firm. The interests of any (common or non-common) shareholder i who has holdings γ_{ig} and γ_{ib} are given by $\gamma_{ig}\pi_g + \gamma_{ib}\pi_b$. Assuming that control is proportional to financial interest, the degree of control of the generic firm is given by γ_{ig} . Decision-makers of the generic firm should maximize

$$\sum_i \gamma_{ig} [\gamma_{ig}\pi_g + \gamma_{ib}\pi_b],$$

where γ_{ig} and γ_{ib} are the shareholdings of any shareholder i that owns shares in either or both of the two firms. Straightforward algebra shows that maximizing this function is equivalent to maximizing

$$\pi_g + \frac{\sum_i \gamma_{ig}\gamma_{ib}}{\sum_i \gamma_{ig}^2} \pi_b,$$

and thus

$$\delta_L(g, b) \equiv \frac{\sum_i \gamma_{ig}\gamma_{ib}}{\sum_i \gamma_{ig}^2}$$

can be thought of a measure of common ownership. This measure captures the importance of the shareholdings in the generic (ability) and shareholdings in the brand (incentives) taking into account the ownership concentration of the generic. See O’Brien and Waehrer (2017) and Backus et al. (2019) for a thorough discussion of this measure, often called “lambda.”

RANK-BASED APPROACH. Another class of measures is based on investors’ *rankings* within a company in terms of holdings. We construct two measures based on investors’ rank. In particular, we construct counts based on the number investors that are ranked in the top 5 or top 10, respectively, in both the brand and generic companies.

$\delta_{top5}(g, b)$ = Number of investors in the top 5 of both the brand and the generic

$\delta_{top10}(g, b)$ = Number of investors in the top 10 of both the brand and the generic

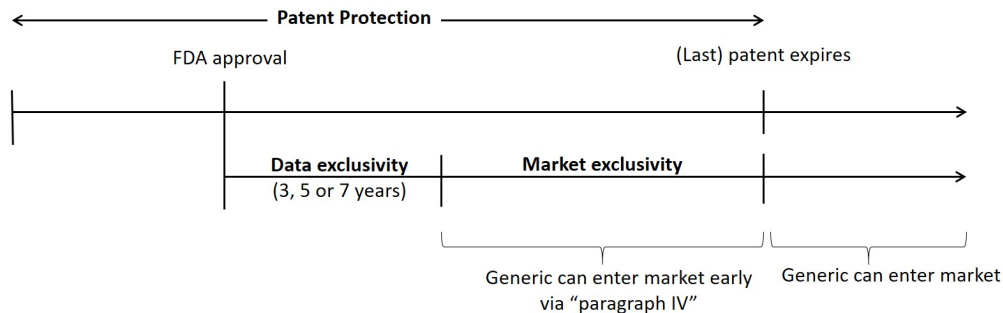
3.4 Data

We explain both the pharmaceutical and common ownership data in this section. More details on the data and construction of the dataset can be found in Appendix 3.8.2.

Entry in the pharmaceutical industry. Broadly speaking, pharmaceutical firms can be categorized as brand firms or generic firms.¹⁰ Brand firms undertake costly research and development to discover new medications and bring them to market, and must apply for FDA approval through the new drug application (NDA) procedure. Once a brand has received FDA approval, it is awarded “data exclusivity” for a period of three, five or seven years, depending on the drug type. Data exclusivity protects the underlying clinical data and runs concurrently with patent protection. The period that spans between the end of data exclusivity and the expiration of the last patent, if any, is commonly referred to as “market exclusivity.”

Generic firms produce bioequivalent replications of brand drugs at a much lower cost, after they have already been marketed as brand-name products. Generic firms are able to enter a particular drug market once the regulatory protections afforded to the brand product have expired. During the market exclusivity period, generics can challenge the monopoly rights of the brand in court, for instance through Paragraph IV certification. Generic companies can also apply for FDA approval once all patents are expired. In both instances, an abbreviated new drug application (ANDA) must be submitted to the FDA. The protection conferred to new drugs is illustrated in Figure 3.1.

Figure 3.1: Exclusivities and patent protection in pharmaceuticals



Notes: This figure illustrates the two types of protection awarded to new drugs. Data exclusivity protects the underlying clinical data and runs concurrently with patent protection. At the end of data exclusivity, a drug is protected only by its patents until they expire, a period termed “market exclusivity.”


We use FDA approval as an indicator of generic entry, in line with several articles on the topic (e.g., Helland and Seabury, 2016; Hovenkamp and Lemus, 2018; Scott

¹⁰Note that we define firms as being a “brand” or a “generic” on a market basis. It is possible that the same firm is a potential generic entrant for one market and the brand company in another market. This can occur because some companies produce both branded drugs and generic drugs.

Morton, 1999, 2000). We consider a market to be open for generic entry at the earlier of either the date of first generic entry or the end of the market exclusivity period.¹¹ If we observe FDA approval of the first generic entrant before the end of the market exclusivity period, then a generic successfully challenged the brand’s patent through a Paragraph IV procedure.¹² We term this point in time the “end of exclusivity.”

We focus on entry that occurs within 6 quarters after the end of exclusivity, as generics prefer to enter a market as early as possible (Wang et al., 2018, Scott Morton, 1999) and it indeed captures most of the actual generic entries in our sample (see Figure 3.2); see further below on the details of our sample. However, given the potential sensitivity of results to our time window, we will show that results are robust to other entry period definitions.

Figure 3.2: Histogram of entry



Notes: This figure illustrates the entry patterns in our data after the “end of exclusivity.” The dark gray area shows the probability that entry occurs within 6 quarters after the end of exclusivity.

Pharma data sources and variables. We obtain NDA and ANDA information

¹¹We consider a market to be open for generic entry at the earlier of either the date of first generic entry or the date at which the *last* patent listed in the Orange Book for the drug expires. To check the robustness of our results to an earlier date for the end of market exclusivity we use information (where available in the Orange Book) on the type of patent. As a robustness check, we take the end of market exclusivity to be the earlier of either the date at which the *substance* patent expires and the date of first generic entry. We re-run our main specification with this adjustment, and find that our results hardly change. We repeat this process by taking the earlier of either the date at which the *product* patent expires and the date of first generic entry. We re-run our main specification with this adjustment, and find that our results hardly change.

¹²Other generics can then enter too, although possibly with a delay of 2 quarters due to temporary monopoly rights conferred to the first paragraph IV filer (see e.g., Hovenkamp and Lemus, 2018).

from the FDA Orange Book. The FDA Orange Book provides data on all launched pharmaceutical products in the United States since 1982. The data includes information on the launching company, type of drug (NDA or ANDA), associated patents, list of ingredients, dosage form, strength, approval date and status (prescription, over-the-counter, or discontinued). Information on the submission class of the brand product is merged in from the “Drugs@ FDA” database using the FDA application number; see also Helland and Seabury (2016) and Hovenkamp and Lemus (2018) for more details on this data source. Products are linked to their relevant therapeutic field using the ATC/DDD Index 2015 and applying exact text matching, based on compound-name.¹³

We define drug markets at the ingredient-form level. For example, the drug with the brand-name Zyrtec in syrup form with the ingredient Cetirizine Hydrochloride 5mg/5ml is considered to be in the same drug market as Zyrtec in syrup form with the ingredient Cetirizine Hydrochloride 10mg/10ml. However, the product Zyrtec Allergy with the ingredient Cetirizine Hydrochloride 10mg in the form of a tablet constitutes a different market. The therapeutic field in which Zyrtec falls, at the ATC-2 level, is “Antihistamines for systemic use.”

We match the brand product (NDA) with the full sample of potential generic entrants to form a brand product-generic observation. The sample of potential generic entrants includes all pharmaceutical companies that launched at least one generic product in our drug markets and have previous experience in launching generic drugs of the same form (i.e. oral, injection etc.) as the relevant brand drug. Results are robust to a set of different definitions of the entrant set, as we will show when discussing the results.

Following prior literature, we construct variables used to control for relevant drug market and generic firm characteristics (Hurwitz and Caves, 1988; Scott Morton, 1999; Kyle, 2000; Hudson, 2000; Saha et al., 2006; Regan, 2008; Glowicka et al., 2009; Moreno-Torres et al., 2009; Appelt, 2015). *Pediatric Drug* is an indicator variable which takes on the value 1 if a drug can be used in children. *Orphan Drug* is an indicator variable that takes the value 1 if a drug treats a rare disease.¹⁴ The indicator variable *Authorized Generic* takes on the value 1 if the brand firm has launched an authorized generic in that particular market.¹⁵

¹³The ATC/DDD Index 2015 categorizes all chemical compounds used in any therapeutic field according to a five-level hierarchical system, called the Anatomical Therapeutic Chemical (ATC) Classification System.

¹⁴This information is obtained by looking at the exclusivities afforded to the drug in the FDA Orange Book. There are special exclusivity provisions for pediatric drugs and orphan drugs.

¹⁵Note that our left-hand variable is independent generic entry. Authorized generics can be launched without FDA approval and at any point in time (typically shortly before patent expiry). An authorized generic may be launched by a partially-owned generic or subsidiary of the brand,


We proxy drug market size using a measure of the brand's pre-generic-entry revenues obtained from Medicaid reimbursements (available publicly from the Medicaid website). We match the drugs in our sample with Medicaid reimbursement data using National Drugs Codes (NDC) which are unique product identifiers for drugs in the US. The *Market Size* variable used in the analysis is the dollar value in billions of total national Medicaid reimbursements for the brand drug in the two years before the end of exclusivity.

We also use the Medicaid reimbursement data matched with our final sample of drug markets to investigate the decline in brand revenues and total drug market revenues after the end of exclusivity, and how this varies with the number of generic entrants. For our sample of drugs, we find that average annual revenues for brand *before* entry are \$46 million in comparison to average revenues of \$2 million for a generic upon entry. Figure 3.3 shows the average percentage decline in revenues two years after the end of exclusivity relative to two years before. Relative to the case where no generics enter the market, on average, brand revenues decline steadily with the number of entrants. Total revenues decline for any positive number of generic entrants.¹⁶ Our data, thus, confirm the key assumptions of our framework: brand revenues decline after entry, and generics' increased revenues do not compensate for this decline, such that total market revenues decline after entry.

and hence would not enter as an *independent* generic.

¹⁶Note that two outliers where total market revenues increased by over 1500% after the end of exclusivity have been removed from the sample when creating this figure.

Figure 3.3: Decline in brand and total revenues with number of generic entrants



Notes: This figure illustrates the average decline in brand revenues and total market revenues two years after the end of exclusivity relative to two years before by the number of generic entrants in the market. The declines reported are relative to the average revenue change when there are no generic entrants.

We also take into account the intensity of inter-molecular competition in the therapeutic field (Appelt, 2015; Regan, 2008). *Substitutes on Patent* provides a count of the number of on-patent substitutive active ingredients listed in the same therapeutic field at the ATC-2 level in the quarter prior to the end of exclusivity. Similarly, *Substitutes off Patent* measures the number of off-patent substitutive active ingredients. Further market characteristics include the therapeutic field of the drug (at the ATC-2 level), submission class of the brand product, drug dosage form/route and the year of the end of exclusivity.^{17,18}

Generic firm characteristics aim to capture the prior experience of the generic in the relevant market. Controlling for generic firm characteristics has shown to be crucial in previous studies (Scott Morton, 1999; Scott Morton, 2002; Kyle, 2006). *Experience Route* serves as a proxy for the potential entrant's experience in the brand drug form/route by counting the number of products with identical route of administration previously launched by the generic one quarter prior to the end of exclusivity. Similarly, *Experience ATC2* serves as a proxy of the entrant's experience in the relevant therapeutic field at the ATC2 level. *Experience New Drug* is constructed as a count of the entrant's previously launched new drugs. Generic

¹⁷Submission classes include Type 1 New Molecular Entity, Type 2 New Active Ingredient, Type 3 New Dosage Form, Type 4 New Combination, Type 5 New Formulation or Other Differences.

¹⁸We recode the FDA form/route variable to construct five form/route classes namely oral, injection, topical, ophthalmic and inhalation.

entrants that are also active in producing new drugs may hold some patents that ease entry. *Breadth of Experience* accounts for the breadth of the generic entrant's portfolio by counting the number of distinct therapeutic fields in which the generic has been active in one quarter prior to the end of exclusivity. The variables concerning generic firm experience and substitutes are calculated using the full FDA Orange Book. Counts start in 1994, 10 years before the start of the sample; results are robust to other starting points.

Common ownership data. We use the Thomson Reuters Global Ownership Database, which includes holdings by each shareholder in each publicly listed firm for every year-quarter. For US-listed firms Thomson Reuters collects ownership information from 13F, 13D and 13G filings, and forms 3, 4, and 5. For companies outside the US, information is sourced from stock exchange filings, trade announcements, company websites, company annual reports and financial newspapers.

The advantages with regard to datasets used by other articles on common ownership are considerable. Most recent articles on common ownership use Thomson's Spectrum database (e.g., Azar et al., 2017; He and Huang, 2017; Xie and Gerakos, 2019). This database is limited to 13F filings, which contains only large investors in US companies, whereas some pharmaceutical companies are not listed on a US stock market. Moreover, the Thomson's Spectrum database shows holdings assigned to the owner that filed the 13F. This is what is commonly referred to as an "as-filed view." Our database utilizes a "money-manager view." With this view, the database combines together one or more filings to link the holdings to the actual firm that manages the investments. In other instances, it might break apart a single filing in order to accomplish the same. The holdings would then be assigned to one or more of the managers listed on the file.

For each firm for each quarter in the period 2003-2014 we extracted data on the shareholders that own at least 1% of the shares, and computed yearly ownership averages. Table 3.1 gives an example of the top 5 investors for the brand-generic pair Johnson & Johnson-Mylan in 2013. As shown, in this pair common shareholders account for the lion's share of the ownership of the top 5 investors.

Table 3.1: Top 5 largest investors (2013)

Brand		Generic	
Johnson & Johnson		Mylan	
State Street Global	6%	Vanguard Group	7%
BlackRock	6%	BlackRock	6%
Vanguard Group	5%	State Street Global	4%
Royal Bank of Canada	2%	Wellington Mgmt.	4%
Wellington Mgmt.	2%	John Paulson	4%

Source: Thomson Global Ownership Database

Our data indicates that top shareholders in the generic have a substantial interest in brand profits. For pairs where both the brand and generic are publicly listed we find that on average, the top 10 shareholders in the *generic* collectively own 51%, valued at \$3.7 billion, in the *generic*. They collectively own 6.5%, valued at \$7.6 billion, in the *brand*. For 75% of brand-generic pairs the value held in the brand by the top 10 shareholders of the generic exceeds the value held in the generic. Given large losses to brand profits upon generic entry, even small stakes in the brand would incentivize common owners to influence generic entry.

3.5 Individual entry

In this section, we empirically investigate the impact of pairwise common ownership linkages between a brand and a generic firm on that particular generic's entry decision for a variety of different empirical specifications.

Common ownership variables. We construct empirical counterparts of the five measures introduced in the theory section: δ_S , δ_C , δ_L , δ_{top5} and δ_{top10} . When constructing the empirical measure of δ_S , for each brand-generic pair, the denominator comprises of the sum all shareholdings in the brand and the generic in our database. As our database includes only investors with at least 1% ownership stake, the denominator may be smaller than the theoretical 2.¹⁹

For private firms, i.e. not listed on a stock-exchange, we assume that they do not have common investors with any other firm. For firms with a presence in the UK, we verified that this assumption holds true using annual return filings with full

¹⁹For clarity, the formula is $\delta_S = \frac{\sum_j (\gamma_{jg} + \gamma_{jb})}{\sum_i (\gamma_{ig} + \gamma_{ib})}$ where the denominator runs over all i investors in our database.

shareholder lists that are also available for private firms from the company registry (Companies House).

We pay particular attention to the case in which the potential generic entrant is a subsidiary of the brand firm. We create an indicator variable that takes on the value 1 if the potential generic entrant is a subsidiary of the brand and 0 if it is not. In the former, the common ownership variables are set to zero.²⁰

We report results using common ownership measured in the year prior to the end of exclusivity, as entry requires time to acquire an approved source of materials and suitable production facilities. About one to two years before filing an ANDA application, the generic firm starts preparing to enter (Reiffen and Ward, 2005). However, as it is unclear at exactly what point time the final entry decision of the generic firm is made, we also check that our results are robust to the use of common ownership measured two and zero years prior to the end of exclusivity.

Figure 3.4 and Figure 3.5 show the evolution of the common ownership measures over time.²¹ It is evident that common ownership has increased significantly from 2003 to 2014. The growth of common ownership was relatively small until the beginning of 2010. The average level of common ownership almost doubled in the last four years of the sample.

²⁰We consider a firm X to be subsidiary of a firm Y if firm Y has a direct ownership stake of more than 50% in firm X . We can also identify minority shareholdings, i.e., when one firm has an ownership stake of less than 50% in another firm. However there are only three pairs in the dataset where the brand has a stake-holding in the potential generic entrant and one pair where the potential generic entrant has a stake-holding in the brand. As this ends up being too few observations to draw meaningful statistical conclusions, we do not consider these links in the analysis.

²¹We only include the company-pairs that are observed for the entire period, as this provides a robust overview of how the degree of connectedness between brand and generic pairs has changed over time.

Figure 3.4: Evolution of common ownership


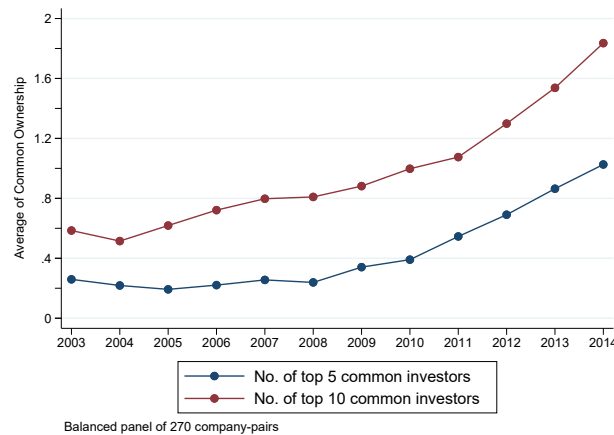


Figure 3.5: Evolution of common ownership - Rank measures



Sample and descriptive statistics. Our final sample consists of 395 drug product markets and 34,144 drug product-brand-generic observations. We consider only drug products that faced generic entry or patent expiry between 2004 and 2014, as this is the range for which we have data on all relevant variables. In total there are 93 unique brand companies. Companies may enter (by incorporation) or exit the sample (by acquisition or bankruptcy). There are 10,453 unique generic-brand pairs. On average there are 86 potential generic entrants per market.

Table 3.2 gives an example of the structure of our data in terms of drug market, brand firm, potential generic entrants, entry and common ownership measures. The example relates to the drug Natrecor which is used for the treatment of heart failure and is produced by Johnson & Johnson. The relevant market is defined by

the ingredients (nestiritide recombinant) and dosage form (solution; intravenous). The patent associated with Natrecor expired in 2014q2. Entry is defined within 6 quarters of the end of market exclusivity, in this case between 2014q2 and 2015q4. According to this definition no generics have entered the market. Indeed, the drug is currently on the FDA List of off-patent, off-exclusivity drugs without an approved generic.²² The common ownership measures correspond to those of the year 2013.

Table 3.2: Example data structure

obs.	trade name	ingredients	dosage form	brand	generic entrant	entry	δ_S	δ_C	δ_L
1	natrecor	nestiritide recombinant	solution; intravenous	JOHNSON & JOHNSON	MYLAN	0	0.67	0.23	0.90
2	natrecor	nestiritide recombinant	solution; intravenous	JOHNSON & JOHNSON	BARR	0	0.51	0.02	0.25
3	natrecor	nestiritide recombinant	solution; intravenous	JOHNSON & JOHNSON	RANBAXY	0	0.05	0.01	0.00
4	natrecor	nestiritide recombinant	solution; intravenous	JOHNSON & JOHNSON	SANDOZ	0	0.45	0.09	0.33
5	natrecor	nestiritide recombinant	solution; intravenous	JOHNSON & JOHNSON	AMNEAL	0	0	0	0
6	natrecor	nestiritide recombinant	solution; intravenous	JOHNSON & JOHNSON	APOTEX	0	0	0	0
.
.

Table 3.3 outlines the key characteristics for the 395 drug markets. The unconditional probability of entry is 2.8%.²³ In 28% of the markets the brand has launched a generic itself, i.e. started selling an authorized generic. In terms of market size, pre-entry brand revenues through Medicaid reimbursements average 100 million USD. The average potential generic entrant has launched 21 generic products of the same route/form as the brand and is active in 14 therapeutic fields.

²²<https://www.fda.gov/downloads/Drugs/ResourcesForYou/Consumers/BuyingUsingMedicineSafely/UnderstandingGenericDrugs/UCM564441.pdf>

²³Both number of entrants and realized entry opportunities are comparable with previous studies: in Scott Morton (1999) there are 123 potential generic entrants per drug market and in Appelt (2015) there are 100 potential entrants per drug market. Furthermore, in Scott Morton (1999) 2-7% of entry opportunities are realized, in Kyle (2006) 2.5% of entry opportunities are realized, and in Appelt (2015) 10% of entry opportunities are realized.

Table 3.3: Summary statistics

VARIABLES	(1) N	(2) mean	(3) sd	(4) min	(5) max
Entry (0/1)	34,144	0.0278	0.164	0	1
δ_S	34,144	0.0851	0.160	0	0.946
δ_C	34,144	0.0249	0.0552	0	0.363
δ_L	34,144	0.0712	0.172	0	1.365
δ_{top5}	34,144	0.290	0.731	0	5
δ_{top10}	34,144	0.674	1.375	0	10
Subsidiary (0/1)	34,144	0.00246	0.0495	0	1
Market Size	34,144	0.101	0.245	0	2.143
Authorized Generic (0/1)	34,144	0.280	0.449	0	1
Orphan Drug (0/1)	34,144	0.0917	0.289	0	1
Pediatric Drug (0/1)	34,144	0.318	0.466	0	1
Substitutes on Patent (ATC2) \div 10	34,144	2.498	1.732	0	7.300
Substitutes off Patent (ATC2) \div 10	34,144	1.714	1.268	0	6.100
Experience Route \div 10	34,144	2.124	3.633	0.100	29.90
Experience ATC2 \div 10	34,144	0.0969	0.261	0	3.200
Experience New Drug \div 10	34,144	0.220	0.466	0	2.800
Breadth (ATC2) \div 10	34,144	1.373	1.243	0.100	6.100

Empirical implementation. We determine which individual generic firms are more likely to enter a given drug market. As our main variable of interest –common ownership between a potential generic entrant and the brand– is firm-specific, our regressions in this section are based on the individual probability of entering, rather than on the market-level number of entrants. However, it is important to remember that –as in our theoretical analysis of the individual entry decision– other potential generic entrants are part of the analysis through their inclusion in the set of potential entrants.

The binary dependent variable thus contains the entry decision of generic firm g in the market m of the brand b . The resulting equation to be estimated is:

$$Pr[Entry_{gm} = 1] = \beta_0 + \beta\delta_{gm}(g, b) + \eta Z_m + \gamma X_{gm} + A_m + \mu_t + \epsilon_{gm}.$$

$Entry_{gm}$ takes on the value 1 when generic g enters market m within 6 quarters after the end of exclusivity. δ_{gm} is one of the measures of common ownership between the generic firm and the brand of the market, i.e., δ_{gm} can be δ_S , δ_C , δ_L , δ_{top5} or δ_{top10} . Z_m is a vector of market characteristics, including market size as measured by pre-generic-entry sales, an indicator for the presence of an authorized generic,

an indicator for pediatric drugs, an indicator for orphan drugs, and the number of on- and off-patent inter-molecular substitutes in same therapeutic field. X_{gm} is a vector of generic-market characteristics, the generic's previous experience with drug from/route, generic's previous experience with the therapeutic class, generic's previous experience with new drugs, number of therapeutic fields in which the generic has experience and region of generic's company headquarters. A vector of fixed effects A_m is included for drug dosage form, submission class and therapeutic field (ATC-2 level), as well as a fixed effect μ_t for the year of the end of exclusivity.

We estimate a linear probability model. Coefficients for the probit and logit models are reported in Appendix 3.8.4.²⁴ The coefficient β measures the impact of common ownership between the brand and the generic on the generic's entry decision.

Results. Table 3.4 presents our results. The coefficient on δ_{gm} across all measures is negative and significant. Thus we find that common ownership between the brand and generic indeed reduces the likelihood of generic entry. The coefficient on common ownership should be interpreted bearing in mind the unconditional probability of entry for the sample. The unconditional probability of entry for the sample of firms and markets is 2.8%. An increase of one standard deviation as measured by δ_S implies a $0.16 \times 0.027 = 0.0043$ decrease in the probability of entry *ceteris paribus*. This is therefore a $0.0043/0.028 = 15\%$ reduction in the unconditional probability of entry. Similarly, an increase of one standard deviation in δ_C and δ_L imply a 15% and 18% decrease, respectively, in the probability of entry.

Column 4 shows that one additional top 5 common investor leads to a 0.6 percentage point decrease in the probability of entry. Therefore, an additional top 5 common investor leads to a $0.006/0.028 = 21\%$ decline in the probability of entry. The effect of having an additional top 10 common investor is also highly significant and negative (see column 5), although the size of the effect is about half. These findings, therefore, are consistent with the idea that higher ranked investors have more power, and effectively use this power to reduce entry.

²⁴There are several therapeutic fields at the ATC2 level which do not experience any entry in our sample, thus the dummy indicators for these ATC2 fields become perfect predictors for a zero outcome. These observations are thus dropped in the logit and probit models.

Table 3.4: Main specification

VARIABLES	(1)	(2)	(3)	(4)	(5)
δ_S	-0.0273*** (0.00676)				
δ_C		-0.0776*** (0.0183)			
δ_L			-0.0287*** (0.00556)		
δ_{top5}				-0.00584*** (0.00142)	
δ_{top10}					-0.00298*** (0.000811)
Subsidiary (0/1)	-0.0555*** (0.0145)	-0.0540*** (0.0145)	-0.0546*** (0.0145)	-0.0275* (0.0148)	-0.0284* (0.0150)
Market Size	0.0344*** (0.00867)	0.0343*** (0.00867)	0.0344*** (0.00865)	0.0343*** (0.00866)	0.0344*** (0.00869)
Authorized Generic (0/1)	0.00647 (0.00409)	0.00654 (0.00410)	0.00652 (0.00409)	0.00645 (0.00410)	0.00646 (0.00410)
Orphan Drug (0/1)	-0.00104 (0.00706)	-0.000929 (0.00706)	-0.00104 (0.00706)	-0.000961 (0.00706)	-0.000963 (0.00706)
Pediatric Drug (0/1)	0.0126*** (0.00477)	0.0125*** (0.00477)	0.0126*** (0.00477)	0.0126*** (0.00479)	0.0126*** (0.00479)
Substitutes on Patent (ATC2) \div 10	-0.00416 (0.00638)	-0.00407 (0.00637)	-0.00408 (0.00637)	-0.00394 (0.00638)	-0.00406 (0.00639)
Substitutes off Patent (ATC2) \div 10	-0.00620 (0.00497)	-0.00617 (0.00498)	-0.00622 (0.00496)	-0.00632 (0.00498)	-0.00624 (0.00498)
Experience Route \div 10	0.00837*** (0.000855)	0.00838*** (0.000856)	0.00841*** (0.000857)	0.00833*** (0.000854)	0.00837*** (0.000856)
Experience ATC2 \div 10	0.0609*** (0.0104)	0.0611*** (0.0104)	0.0609*** (0.0104)	0.0611*** (0.0104)	0.0611*** (0.0104)
Experience New Drug \div 10	0.00453 (0.00286)	0.00408 (0.00283)	0.00467 (0.00285)	0.00385 (0.00280)	0.00407 (0.00285)
Breadth (ATC2) \div 10	0.00101 (0.00231)	0.00114 (0.00231)	0.000952 (0.00229)	0.000670 (0.00230)	0.000891 (0.00231)
Observations	34144	34144	34144	34144	34144
R-squared	0.0855	0.0855	0.0857	0.0855	0.0854
Therapeutic field	Yes	Yes	Yes	Yes	Yes
Drug form	Yes	Yes	Yes	Yes	Yes
Submission type	Yes	Yes	Yes	Yes	Yes
Generic region of origin	Yes	Yes	Yes	Yes	Yes
Year end of exclusivity	Yes	Yes	Yes	Yes	Yes
Drug markets	395	395	395	395	395

Notes: OLS regression. Standard errors in parentheses are clustered at the drug market level. The dependent variable is entry within 6 quarters. The constant term is estimated but not reported. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The control variables carry the expected signs; higher pre-entry brand revenues and greater entrant experience all significantly increase the likelihood of entry. Pe-

diatric drugs are also more likely to experience entry. On the other hand, we find that the number of molecular substitutes on and off-patent do not have a significant impact on generic entry.

Directly relevant for the topic of the study, the effect of common ownership is smaller than the effect of being a subsidiary of the brand. For example, if we set δ_S to 1 – that is the brand and generic share all the same common owners – then the probability of entry falls by 2.7 percentage points. On the other hand, if the relationship is parent-subsidiary then the probability of entry falls by 5.5 percentage points *ceteris paribus*.

The fact that we find a significant effect across all measures of common ownership, and similar effects in terms of economic magnitude implies that we cannot say much about which measure of common ownership best captures the *manner* in which common investors' incentives and ability translate into the weight that the generic firm places on joint profits. This is in fact no surprise as empirically we find that the five measures of common ownership are highly correlated with each other (see Table 3.5). Thus, although in theory our measures capture quite different mechanisms of influence, the empirical counterparts are quite similar and the variation across brand-generic pairs is small.

Table 3.5: Cross-correlations between common ownership measures

Variables	δ_S	δ_C	δ_L	δ_{top5}	δ_{top10}
δ_S	1.000				
δ_C	0.890	1.000			
δ_L	0.894	0.890	1.000		
δ_{top5}	0.793	0.763	0.792	1.000	
δ_{top10}	0.874	0.828	0.805	0.862	1.000

In Table 3.6 we present results where common ownership is specified as a categorical variable in order to investigate whether greater levels of common ownership have a larger impact; i.e., whether the relationship between common ownership and entry is non-linear. We focus on the measure δ_S . This measure can be interpreted as the fraction of total ownership in the pair held by common investors, and hence presents natural thresholds. We construct three categorical variables based on the value of δ_S : $\delta_S(0 < \delta \leq 0.3)$ takes on the value 1 if $\delta_S \in (0; 0.3]$, $\delta_S(0.3 < \delta \leq 0.5)$ takes on the value 1 if $\delta_S \in (0.3; 0.5]$, and $\delta_S(0.5 < \delta \leq 1)$ takes on the value 1 if $\delta_S \in (0.5; 1]$.

The results in Table 3.6 indicate that the effect of common ownership increases the greater the level of common ownership. The coefficients on each categorical variable increase in magnitude (become more negative) with higher common ownership.

Furthermore, once δ_S is greater than 0.5 the coefficient is significant at the 1% level. A change from zero common ownership to common ownership of greater than 0.5 reduces the entry probability of a generic by 1.5 percentage points on average. This is approximately a 50% decline in the unconditional probability of entry. In our sample, there are 552 unique brand-generic pairs with a δ_S of greater than 0.5 at some point in time. This is 5% of all brand-generic pairs.

In sum, these results indicate that common ownership levels have a non-linear impact on entry, where high levels have a much stronger impact than low levels. In particular, common ownership has its strongest and most significant effect when more than half of the total ownership in the pair is in the hands of common investors. We will use this finding later on in our market-level analysis when constructing market-level measures of common ownership.

Table 3.6: Categorical specification

VARIABLES	(1)
δ_S ($0 < \delta \leq 0.3$)	0.00186 (0.00315)
δ_S ($0.3 < \delta \leq 0.5$)	-0.00785* (0.00410)
δ_S ($\delta > 0.5$)	-0.0149*** (0.00464)
All controls	Yes
Therapeutic field	Yes
Drug form	Yes
Submission type	Yes
Generic region of origin	Yes
Year end of exclusivity	Yes
Drug markets	395
Observations	34144
R-squared	0.0854

Notes: OLS regression. Standard errors in parentheses are clustered at the drug market level. The dependent variable is entry within 6 quarters. The constant term is estimated but not reported. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.7: Instrumental variables regression

VARIABLES	(1)	(2)	(3)	(4)	(5)
δ_S	-0.0285* (0.0147)				
δ_C		-0.0636* (0.0328)			
δ_L			-0.0239* (0.0124)		
δ_{top5}				-0.00648* (0.00338)	
δ_{top10}					-0.00324* (0.00167)
Subsidiary (0/1)	-0.0557*** (0.0146)	-0.0536*** (0.0144)	-0.0541*** (0.0144)	-0.0249 (0.0198)	-0.0264 (0.0191)
Market Size	0.0345*** (0.00864)	0.0343*** (0.00864)	0.0344*** (0.00862)	0.0344*** (0.00863)	0.0344*** (0.00867)
Authorized Generic (0/1)	0.00647 (0.00408)	0.00654 (0.00409)	0.00653 (0.00409)	0.00644 (0.00409)	0.00645 (0.00409)
Orphan Drug (0/1)	-0.00103 (0.00706)	-0.000991 (0.00705)	-0.00108 (0.00705)	-0.000927 (0.00707)	-0.000936 (0.00706)
Pediatric Drug (0/1)	0.0126*** (0.00474)	0.0125*** (0.00475)	0.0125*** (0.00475)	0.0126*** (0.00475)	0.0127*** (0.00476)
Substitutes on Patent (ATC2) \div 10	-0.00417 (0.00636)	-0.00407 (0.00636)	-0.00408 (0.00635)	-0.00392 (0.00637)	-0.00406 (0.00637)
Substitutes off Patent (ATC2) \div 10	-0.00619 (0.00495)	-0.00621 (0.00496)	-0.00625 (0.00495)	-0.00631 (0.00497)	-0.00622 (0.00496)
Experience Route \div 10	0.00837*** (0.000854)	0.00837*** (0.000853)	0.00840*** (0.000855)	0.00833*** (0.000852)	0.00837*** (0.000855)
Experience ATC2 \div 10	0.0609*** (0.0104)	0.0611*** (0.0104)	0.0609*** (0.0104)	0.0610*** (0.0104)	0.0611*** (0.0104)
Experience New Drug \div 10	0.00465 (0.00308)	0.00363 (0.00287)	0.00414 (0.00296)	0.00410 (0.00296)	0.00429 (0.00300)
Breadth (ATC2) \div 10	0.00105 (0.00235)	0.000931 (0.00234)	0.000786 (0.00232)	0.000747 (0.00232)	0.000971 (0.00234)
Observations	34144	34144	34144	34144	34144
Therapeutic field	Yes	Yes	Yes	Yes	Yes
Drug form	Yes	Yes	Yes	Yes	Yes
Submission type	Yes	Yes	Yes	Yes	Yes
Generic region of origin	Yes	Yes	Yes	Yes	Yes
Year end of exclusivity	Yes	Yes	Yes	Yes	Yes
Drug markets	395	395	395	395	395

Notes: 2SLS regression. Standard errors in parentheses are clustered at the drug market level. The dependent variable is entry within 6 quarters. The constant term is estimated but not reported. The instrument is an indicator equal to 1 if both firms are listed on the Dow Jones US Select Pharmaceutical Index. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Instrumental variables. If investors adjust their holdings in response to entry opportunities, common ownership might be endogenous. The direction of the bias

is not clear a priori. For example, if investors in the brand increase investment in generics with entry plans, common ownership between the brand and generic will increase before entry, causing β to be biased upwards. Alternatively, if investors with shareholdings in generic firms reduce the size of their stakes in brand firms with drugs that face impending generic entry, then common ownership will decrease before entry, causing β to be biased downwards.

To address the concern of endogeneity, we therefore also perform IV estimations and instrument for common ownership with financial index membership at the pair level.²⁵ We use information on the holdings included in the Dow Jones US Select Pharmaceutical Index during the 2006-2014 period. Our data on the composition of the Dow Jones US Select Pharmaceutical Index comes from historical data on the composition of BlackRock's iShares US Pharmaceutical exchange-traded fund (ETF) which tracks the Dow Jones US Select Pharmaceutical Index.

Figure 3.6 in Appendix 3.8.1 provides a snapshot of the top 10 investments of the fund as of November 2013. As can be seen, both brand and generic firms are present in the fund; e.g. Johnson & Johnson is a brand company, whereas Mylan primarily produces generic drugs. On average, the fund has been comprised of 39 holdings over time, each allocated a specific weight that changes over time; relative weightings are computed according to their total market capitalization.²⁶ Since May 2006, each listed company has been included in the ETF for an average of 4 years. This evidences the pattern of entry and exit of the fund that has been marked by various periods of high entrance and exit – for instance, more than 6 companies dropped out and entered the fund in the last quarter of 2013 and the third quarter of 2015, respectively – and periods of no change.

Our instrument, *Index Presence* is an indicator equal to 1 if both firms are listed on the Dow Jones US Select Pharmaceutical Index at the point in time when common ownership is measured. We expect that if both companies in the pair appear in the Index, common ownership will increase by virtue of the fact that investors who track the index will hold shares in both companies. The identifying assumption is that inclusion in the pharmaceutical index, is exogenous to a particular market entry,

²⁵A similar approach has been applied by several other articles in the literature. For example, Aghion (2013) use the inclusion of a firm in the S&P 500 as an instrument for institutional ownership. Bena et al. (2017) instrument foreign institutional ownership with stock additions and deletions to the MSCI all country world index. Schmidt and Fahlenbach (2017) instrument passive institutional ownership with switches between the Russel 1000 and Russel 2000 indexes. Scott Morton and Boller (2020) use instances of a stock entering the S&P 500 index to test if an increase in common ownership changes future expected profits of the entering firm and its product market rivals.

²⁶A detailed description of how the Dow Jones US Select Pharmaceutical Index is constructed can be found at: <https://www.spglobal.com/spdji/en/documents/methodologies/methodology-dj-us-select-sector-specialty.pdf>

except through its effect on common ownership. This is the case provided that the index is not created with potential entry opportunities in mind and that, controlling for other factors, addition to the index does not directly affect entry decisions except through common ownership.

The results of the IV regressions are presented in Table 3.7. The IV results are very similar to the OLS results. The first-stage results, reported in Table 3.8, indicate that the instrument is highly relevant and positively correlated with δ (see F-stat of excluded instruments). However, the Durbin-Wu-Hausman test shows that we cannot reject the hypothesis that δ is exogenous for all measures of δ . This suggests that, in this context, the endogeneity of common ownership is not a large concern. One possible explanation for this finding is that a large share of common investors in our dataset are passive investors and are therefore unlikely to actively adjust their holdings in generic or brand firms due to potential entry opportunities.

Table 3.8: First-stage IV regressions

VARIABLES	(1) δ_S	(2) δ_C	(3) δ_L	(4) δ_{top5}	(5) δ_{top10}
Index Presence	0.337*** (0.00885)	0.151*** (0.00444)	0.402*** (0.0152)	1.481*** (0.0637)	2.968*** (0.0752)
Constant	0.00990 (0.0328)	-0.00263 (0.0101)	0.00824 (0.0295)	-0.136 (0.137)	-0.217 (0.253)
Observations	34,144	34,144	34,144	34,144	34,144
R-squared	0.400	0.487	0.404	0.421	0.480
Fixed Effects	Yes	Yes	Yes	Yes	Yes
Drug markets	395	395	395	395	395
F-stat excl. instruments	1454	1157	704.7	540.6	1556
Endogeneity test (p-val)	0.929	0.593	0.649	0.833	0.857

Notes: For simplicity only the coefficient associated with the excluded instrument is reported. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Robustness checks. Our results are robust to a series of different specifications, as can be seen from the tables in Appendix 3.8.4.

Table 3.12 shows results where we add drug product fixed effects. The significance and magnitude of the coefficients stays virtually the same. Table 3.13 and Table 3.14 present probit and logit regressions for our main specification respectively. Results show that our five common ownership measures negatively impact entry.

In our main specification we use Medicare reimbursements as a proxy for market size. To check the robustness of results to this measure, in Tables 3.15 and 3.16 we

use a different proxy for market size based on the total sales volume for the drug in the US. Data on the sales of pharmaceutical drugs are available commercially but are expensive to acquire. We thus use publicly available sales data from drugs.com. Drugs.com provides the annual US sales figures for the top 200 drugs for the years 2003 - 2010 (source: Verispan/VONA) and the top 100 drugs for the years 2011 - 2013 (source: IMS Health/Midas). In Table 3.15 we substitute Medicare reimbursements with an indicator variable for whether or not the drug is in the top 100 in terms of sales. In Table 3.16 we limit the sample to drugs where we have information on the annual US sales in the year before the market becomes open for entry. Our results are robust to both of these alternative measures of market size.

Another issue may be the set of potential entrants. In our main specification, the set of potential entrants is quite narrowly defined. We exclude generic firms from the potential entrant set that have not previously launched a generic drug of the same form as the relevant brand drug. Doing so, however, means that we drop 51 *actual* entry observations (5% of all actual entry observations). To check the robustness of our results to a broader potential entrant set, we expand the set to include generics without experience in the relevant drug form. Results in Table 3.17 show that the effects are qualitatively identical to our main results: for all common ownership measures, the effect is negative and significant at the 1% level.

We also test the robustness of our results to different entry time windows, as entry may be slower or faster than our chosen 6 quarter window. In Table 3.18 we alter the dependent variable such that we consider entry within any time period. We also consider alternative time windows such as entry within a 2 year time period. These results are available on request and findings are qualitatively the same as in our main specification, i.e., entry is significantly negatively influenced by common ownership and this holds for different time windows.

3.6 Market outcomes

Up until this point, we have empirically established that a higher level of common ownership between a brand and a particular generic reduces the entry probability of that specific generic firm. However, it is relevant to consider if, as a result of this, common ownership has an impact on outcomes at the market level such as the total number of generics in a market.

Common ownership measures. To assess the effects of common ownership at the market level, we need to construct market-level measures of common ownership. Our first measure of common ownership at the market level is a simple average of δ_S taken over all potential generic entrants for the relevant market. Our second measure

is motivated by the fact that common ownership has non-linear effects (see results in Table 3.6). To calculate the second measure we first define a common ownership link between a brand and potential generic entrant as occurring when $\delta_S > 0.5$. For each market, we then count the number of common ownership links and divide this sum by the total number of potential generic entrants. Whereas the first measure provides an idea of the average level of common ownership in the market, the second measure focuses on establishing the extent of “strong” common ownership links.

$$\delta_S \text{ (market)} \equiv \frac{\sum_{g \in S_m} \delta_S}{S_m} \quad (3.5)$$

and

$$\text{Share links } \delta_S > 0.5 \equiv \frac{\sum_{g \in S_m} I(\delta_S > 0.5)}{S_m}, \quad (3.6)$$

where $I(x)$ is an indicator function that takes value 1 if the event x occurs, and 0 otherwise and S_m is the set of potential entrants in market m .

Empirical implementation. We consider several outcomes at the market level; 1) the number of generic entrants within 6 quarters, 2) the number of generic entrants ever, 3) the share of generic entrants within 6 quarters²⁷, 4) the share of generic entrants ever²⁸, 5) the duration of time (in quarters) until the first generic enters²⁹ and finally 6) the probability that there is no generic entry at all.

For each outcome variable we estimate the following linear model:

$$Y_m = \beta_0 + \beta \delta_m + \eta Z_m + A_m + \mu_t + \epsilon_{gm}.$$

where Y_m is one of the outcome variables mentioned above, δ_m is one of the two market-level measures of common ownership, Z_m is a vector of market-level control variables, A_m is a vector of fixed effects for drug dosage form, submission class and therapeutic field (ATC-2 level), and lastly μ_t is a fixed effect for the year of the end of exclusivity.

Table 3.9 presents summary statistics for the market-level measures of common ownership and the market outcomes. The average level of common ownership as measured by δ_S at the market level is 0.09. On average the brand firm is connected

²⁷The share is calculated as the number of actual entrants within 6 quarters divided by the total number of potential entrants.

²⁸The share is calculated as the number of actual entrants to ever enter the market divided by the total number of potential entrants.

²⁹The duration is calculated as the number of quarters from when the market becomes open for entry until the first generic entrant. If there is no generic entry within the sample, the duration is calculated as the time between when the market becomes open for entry and the end date of the dataset.

with 4% of potential generic entrants by a common ownership link where $\delta_S > 0.5$. This measure varies from a minimum of 0 to a maximum of 32%. The average number of generic entrants within 6 quarters for our sample of drug markets is 2.4. This figure increases to 3.8 if we do not limit ourselves to a specific time window and consider all occurrences of entry in the data. On average the share of generic entrants who actually enter, out of the set of potential entrants, is 3.7%. On average, the first generic enters 6 quarters after the market becomes open for entry. The final row of Table 3.9 indicates that for 20% of our markets, there is no generic entry at all.

Table 3.9: Summary statistics at the market-level

VARIABLES	(1) N	(2) mean	(3) sd	(4) min	(5) max
δ_S (market)	395	0.0913	0.0572	0	0.283
Share links $\delta_S > 0.5$	395	0.0434	0.0558	0	0.318
No. entrants within 6 quarters	395	2.400	2.978	0	17
No. entrants	395	3.838	4.070	0	23
Entry share within 6 quarters	395	0.0365	0.0603	0	0.500
Entry share	395	0.0564	0.0757	0	0.625
Time until first entry (in quarters)	395	5.954	11.93	0	55
No entry (0/1)	395	0.197	0.399	0	1

Results. Table 3.10 presents the coefficient estimates for the first measure of common ownership and Table 3.11 presents the results for the second measure. In both tables, the coefficients on common ownership carry the same signs: Common ownership has a negative effect on the total number of entrants within 6 quarters (column 1), a negative effect on the number of entrants ever (column 2), a negative effect on the share of entrants within 6 quarters (column 3), a negative effect on the share of entrants ever (column 4), a positive impact on the time until the first generic entry (column 5), and a positive impact on the probability that there is no entry whatsoever for the market (column 6).

In Table 3.11 all the estimates of β are significant, whereas in Table 3.10 we find that the coefficients on the share of entry ever and the duration until the first generic entry are not significant. This result is in line with our previous finding that common ownership displays non-linear effects, where higher levels of common ownership have a larger and more significant impact. Thus, we find more significant effects when we focus on “strong” common ownership links.

We now consider the economic magnitude of these effects. We first describe

how common ownership affects the number of generic entrants in the market. We find that a one standard deviation increase in δ_S (*market*) leads to a decrease of 0.33 (0.057×-5.821) entrants (see Table 3.10 column 1). This is a 14% decline in the unconditional average number of generic entrants within 6 quarters. A one standard deviation increase in *Share links* $\delta_S > 0.5$ leads to a decrease of 0.36 (0.056×-6.389) entrants (see Table 3.11 column 1). This is a 15% decline in the unconditional average number of generic entrants within 6 quarters. The size of the effect is similar when considering the number of entrants within any time frame (column 2).

In columns 3 and 4 in both tables, the outcome variable of interest is what *proportion* of potential entrants actually enter the market. We find that a one standard deviation increase in *Share links* $\delta_S > 0.5$ leads to a 18% decrease in the average share of entrants within 6 quarters.

The results in column 5 indicate that common ownership has a positive impact on the duration of time until generic entry. A one standard deviation increase in *Share links* $\delta_S > 0.5$ extends the time to generic entry by 1.4 quarters which is a 24% increase in the average time until generic entry. Thus, not only does common ownership result in fewer generic entrants, but it also delays the onset of generic competition.

Finally, an increase in common ownership also makes it more likely that a brand firm will face zero competition from generic entrants: a one standard deviation increase in *Share links* $\delta_S > 0.5$ increases the probability of no generic entry by $0.056 \times -1.147 = 6.4$ percentage points.

We can use the estimated coefficients to predict the total number of entrants for different values of *Share links* $\delta_S > 0.5$ using the results of column (2) in Table 3.11. We find that when going from the minimum level of *Share links* $\delta_S > 0.5$, i.e., having no major common ownership links at all, to the maximum market level of 0.318, the average number of entrants in a market would go down from about 4.2 to 1.7, keeping all else constant. Thus, we find that common ownership has an economically significant effect on total generic entry as it reduces the average number of total entrants by 50% at its maximum.

Table 3.10: Market outcomes - δ_S (market)

VARIABLES	(1) N-6q	(2) N	(3) Share-6q	(4) Share	(5) Time	(6) No entry
δ_S (market)	-5.821* (3.006)	-8.231** (3.956)	-0.0818 (0.0605)	-0.136* (0.0773)	20.57 (14.47)	0.957** (0.470)
Market Size	3.622*** (0.914)	4.762*** (1.089)	0.0320** (0.0129)	0.0488*** (0.0163)	-4.779** (2.197)	-0.114* (0.0655)
Authorized Generic (0/1)	0.471 (0.380)	0.789* (0.460)	0.0101 (0.00800)	0.0141 (0.00920)	-2.223 (1.380)	-0.0958** (0.0426)
Orphan Drug (0/1)	-0.403 (0.611)	-0.193 (0.808)	0.00649 (0.0168)	0.0116 (0.0207)	1.772 (2.277)	0.0799 (0.0838)
Pediatric Drug (0/1)	0.990** (0.442)	1.545*** (0.548)	0.0216** (0.0106)	0.0308** (0.0122)	-2.879* (1.588)	-0.0821 (0.0523)
Substitutes on Patent (ATC2) \div 10	-0.414 (0.486)	-0.200 (0.631)	-0.00543 (0.00900)	-0.00994 (0.0120)	0.827 (1.756)	7.57e-05 (0.0577)
Substitutes off Patent (ATC2) \div 10	-0.106 (0.336)	-0.428 (0.419)	-0.0112 (0.0105)	-0.0151 (0.0123)	-0.403 (2.086)	-0.0172 (0.0596)
Observations	395	395	395	395	395	395
R-squared	0.391	0.467	0.381	0.446	0.326	0.348
Therapeutic field	Yes	Yes	Yes	Yes	Yes	Yes
Drug form	Yes	Yes	Yes	Yes	Yes	Yes
Submission type	Yes	Yes	Yes	Yes	Yes	Yes
Year end of exclusivity	Yes	Yes	Yes	Yes	Yes	Yes

Notes: OLS. Standard errors in parentheses are robust. The constant term is estimated but not reported.
 *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.11: Market outcomes - Share links $\delta_S > 0.5$

VARIABLES	(1) N-6q	(2) N	(3) Share-6q	(4) Share	(5) Time	(6) No entry
Share links $\delta_S > 0.5$	-6.389** (2.641)	-8.157** (3.636)	-0.117* (0.0624)	-0.164** (0.0772)	25.48* (13.80)	1.147** (0.506)
Market Size	3.643*** (0.901)	4.781*** (1.088)	0.0326** (0.0129)	0.0495*** (0.0164)	-4.890** (2.232)	-0.118* (0.0657)
Authorized Generic (0/1)	0.543 (0.386)	0.887* (0.467)	0.0112 (0.00805)	0.0158* (0.00925)	-2.486* (1.376)	-0.108** (0.0426)
Orphan Drug (0/1)	-0.408 (0.615)	-0.199 (0.810)	0.00639 (0.0167)	0.0114 (0.0204)	1.792 (2.287)	0.0808 (0.0845)
Pediatric Drug (0/1)	0.932** (0.440)	1.458*** (0.548)	0.0209** (0.0105)	0.0295** (0.0121)	-2.687* (1.540)	-0.0730 (0.0504)
Substitutes on Patent (ATC2) $\div 10$	-0.334 (0.488)	-0.0923 (0.635)	-0.00410 (0.00915)	-0.00794 (0.0123)	0.520 (1.786)	-0.0139 (0.0572)
Substitutes off Patent (ATC2) $\div 10$	-0.152 (0.333)	-0.490 (0.414)	-0.0119 (0.0105)	-0.0162 (0.0123)	-0.229 (2.082)	-0.00921 (0.0593)
Observations	395	395	395	395	395	395
R-squared	0.393	0.467	0.385	0.449	0.330	0.353
Therapeutic field	Yes	Yes	Yes	Yes	Yes	Yes
Drug form	Yes	Yes	Yes	Yes	Yes	Yes
Submission type	Yes	Yes	Yes	Yes	Yes	Yes
Year end of exclusivity	Yes	Yes	Yes	Yes	Yes	Yes

Notes: OLS. Standard errors in parentheses are robust. The constant term is estimated but not reported.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

3.7 Conclusion

Ownership linkages between firms, which typically arise due to large investors that invest in multiple firms in an industry, are a defining feature of firm ownership structures in the present day. Consequently the question of whether these investors influence firm strategies and correspondingly whether common ownership between rival firms has an effect on product markets outcomes has recently attracted significant attention.

In this article we consider the effect of common ownership on market entry decisions in the pharmaceutical industry. Given that generic entry results in substantial revenue losses for the brand firm that can be much higher than the generic's gains from entry, a simple theory model shows that higher common ownership reduces generic entry as common owners have both the incentive and ability to push back

entry. Empirical results lend robust support to this proposition. We show that higher common ownership between a potential generic entrant and the brand firm (incumbent) in a specific drug market has a significant negative effect on the likelihood that the generic firm will enter the market. Based on a linear probability model that relates generic entry to several measures of common ownership with the brand, we find that a one-standard-deviation increase in common ownership decreases the probability of generic entry by 15-18%. Moreover, we show that common ownership has an effect on the overall number of generic firms in a market. A one-standard-deviation increase in common ownership at the market level decreases total entry by as much as 15%. Still, it is perhaps important to stress that, as compared to the effect of being fully owned by the brand, the effect of any level of common ownership between the generic and the brand is smaller.

This research contributes to the literature on the product markets effects of common ownership and informs the current debate. We provide evidence that is consistent with the hypothesis that common shareholders indeed influence strategic decisions of companies. Given the importance of generic entry in terms of reducing drug prices and therefore overall healthcare costs, common ownership in the pharmaceutical industry may have the potential to raise the costs to consumers and healthcare payers.

There is room for future work on the topic in several dimensions. First, to make a clear welfare assessment on the link between common ownership and welfare, a more structural empirical model is needed where entry, pricing and innovation decisions are explicitly modelled. Further, much still needs to be done to understand the corporate governance of common ownership, both how holdings translate into incentive, and ability to influence on the one hand, and how preferences of diverse investors are aggregated into firm's decisions on the other hand. Finally, US pharmaceutical markets are a clear example where common ownership can impact entry. Indeed, given the large asymmetries between brand and generic profits, incentives are high. It would be interesting to identify other markets where both incentives and abilities are high, and to investigate whether common owners exert influence there on entry too.

3.8 Appendix

3.8.1 Common ownership

Anecdotal Evidence

We provide some anecdotal evidence that institutional investors are interested in influencing governance, policies and strategic decisions of firms. Evidence in Appel et al. (2016) suggests that informal discussions between institutions and managers, backed with the threat of voice (i.e., voting in shareholding meetings), are often used to exert influence. Glenn Booraem, controller of Vanguard funds, notes that engagement with directors and management of companies is a key component and that Vanguard has “found through hundreds of discussions every year” that it is “frequently able to accomplish as much -or much more through dialogue” as through voting (Booraem, 2014).

Furthermore, Vanguard’s chairman recently stated that Vanguard seeks active interactions with firms they invest in: “In the past, some have mistakenly assumed that our predominantly passive management style suggests a passive attitude with respect to corporate governance. Nothing could be further from the truth.”³⁰ A similar message emerges from BlackRock’s chairman Larry Fink, “We are an active voice, we work with companies, we need to work for the long-term interest.”³¹

Specifically in pharmaceutical markets, institutional investors can be seen to take an active interest in the strategic decisions of companies. In 2016, a group of representatives of major US mutual funds (Fidelity Investments, T. Rowe Price Group Inc., Wellington Management Co., among others) met up with top biotechnology and pharmaceutical executives and lobbyists to discuss the pricing conditions of the market and the possible steps that could be taken in order to avoid future regulations. This example also illustrates that investor interactions need not be addressed to a particular company but can be extended to a specific industry.³²

³⁰Letter sent by F. William McNabb III, Vanguard’s Chairman and CEO, to the independent leaders of the boards of directors of the Vanguard funds’ largest portfolio holdings, dated 27 February 2015, available at [https://about.vanguard.com/vanguard-proxy-voting/CEO Letter 03 02 ext.pdf](https://about.vanguard.com/vanguard-proxy-voting/CEO%20Letter%2003%20ext.pdf).

³¹Wall Street Journal, ‘BlackRock’s Larry Fink: typical activists are too short-term’, dated 16 January 2014, available at <http://blogs.wsj.com/moneybeat/2014/01/16/blackRocks-larry-fink-typical-activists-are-too-short-term/>

³²Chen, C. (2016). Mutual fund industry to drug makers: stand up and defend yourself. Bloomberg News. Retrieved from <https://www.bostonglobe.com/business/2016/05/10/mutual-fund-industry-drugmakers-stand-and-defend-yourself/REKxLITGDeQR2oVmUZaTIP/story.html>

Figure 3.6: iShares U.S. Pharmaceutical ETF (IHE) - Snapshot of Holdings

iShares U.S. Pharmaceuticals ETF [Fact Sheet](#) [Prospectus](#) [Download](#)

Overview Performance Key Facts Characteristics Fees **Portfolio** Literature

Top 10 All

as of Nov 29, 2013 [Custom Columns](#)

Ticker	Name	Sector	Weight (%)	Notional Value
JNJ	JOHNSON & JOHNSON	Pharmaceuticals	10.43	-
PFE	PFIZER INC	Pharmaceuticals	9.59	-
MRK	MERCK & CO INC	Pharmaceuticals	7.85	-
BMY	BRISTOL MYERS SQUIBB	Pharmaceuticals	6.84	-
ABT	ABBOTT LABORATORIES	Pharmaceuticals	5.59	-
A60	ACTAVIS INC	Pharmaceuticals	5.06	-
LLY	ELI LILLY	Pharmaceuticals	4.76	-
AG4	ALLERGAN	Pharmaceuticals	4.19	-
MYL	MYLAN INC	Pharmaceuticals	3.38	-
PRGO	PERRIGO COMPANY	Pharmaceuticals	3.32	-

3.8.2 Dataset construction

This Appendix contains a detailed description of how the data used for the analysis in this article was constructed. The Orange Book has been downloaded from the FDA website for each year (2001q4, 2002q4,..., 2017q4) using Internet Archive. In the current version of the Orange Book online the names of companies have been partially back-dated to display the current manufacturer of a drug. To establish the company name and drug status at the time of approval, we merged information from multiple versions of the FDA Orange Book.

Duplicate applications in the FDA Orange Book were identified and removed. Where duplicate applications had different approval dates, the earlier date was taken. Thereafter the products in the dataset were merged with historical patent data from the FDA based on the FDA drug application number and product number. The patent data provides a complete list of which patents are associated with the product and their corresponding expiration dates.

In the FDA Orange Book, a drug product can be identified as a unique ingredient-form-strength combination. For example, Cetirizine Hydrochloride in syrup form with a strength of 5mg/5ml. Initially, the FDA Orange Book reports 3964 products at the ingredient-form-strength level that were launched from 1982q1 until 2017q2. For our purposes we restricted the data in multiple ways. First, we consider only drug products that faced generic entry or patent expiry in the time frame 2004q1 to 2014q4 (this is the range where we have data on all variables). This results in a sample of 1080 unique drug products. We then drop drug products which are not linked to any patent (as this study focuses on market entry in markets that are initially protected by patents). This results in 666 unique drug products. Thereafter we drop OTC drugs, keeping only prescription drugs. This results in 640 unique drug products.

On the basis of information contained in the Orange Book we seek to remove drug products where the original brand drug was withdrawn for safety reasons. We identify these products as cases where the original brand has been discontinued, and there is no note in the Orange Book that the discontinuation was not for safety reasons. Dropping these brand products results in 554 unique drug products. We drop two further drug products where generic applications (ANDAs) were approved before the NDA application for the same ingredient-form-strength. This results in 552 drug products.

We then aggregate these drug products to the ingredient-form level. We take the first strength that was approved by the FDA at the ingredient-form level as the relevant brand product. We then identify subsequent ANDAs that were approved at the same ingredient-form level. In cases where a generic enters with multiple strengths, we keep only the earliest entry. This results in 457 unique drug product markets, or brand products, at the ingredient-form level.

A variable is constructed that takes the earlier of either generic entry or the date of the last expiring patent for the relevant product market at the ingredient-form level; called “end of exclusivity.”

Each product is linked through exact text matching, based on compound-name, with the ATC/DDD Index 2015.³³ The ATC/DDD Index 2015 is used to identify relevant therapeutic

³³The ATC/DDD Index 2015 categorizes all chemical compounds used in any therapeutic field according to a five-level hierarchical system, called the Anatomical Therapeutic Chemical (ATC) Classification System. The highest level (ATC1) consist of 14 anatomical main groups (e.g. Alimentary Tract and Metabolism (A) or Cardiovascular System (C)). The next lower level (ATC2) describes 88 therapeutic main groups (e.g. Drugs used in Diabetes (A10) or Diuretics (C03)). Lower levels make even finer distinctions between products. The lowest level (ATC5) indicates 3709 chemical substances.

markets and chemical classes for different levels of the ATC classification system. Whereas the ATC3 level is most in line with market definition in M&A approval procedures in Europe and the United States, through the matching process one drug may be linked with numerous therapeutic classes at the ATC3 level. To ensure that we obtain a unique therapeutic class for each drug, we use the broader market definition of ATC2.

For each drug product market, we identify if the brand firm has launched its own generic in the market (an “authorized generic”) using the FDA list of authorized generics. The merge was conducted on the basis of trade name and form. Additional information, such as submission class, is merged in using the FDA application number.³⁴ We recode the FDA form/route variable to construct five form/route classes namely oral, injection, topical, ophthalmic and inhalation.

The data on firms and their product launches from the FDA Orange book is then matched with the Thomson Reuters ownership dataset based on the name of the pharmaceutical company. We correct for the fact that firms may change their name over the course of the sample period and undergo mergers, on the basis of public information. We record the year-quarters in which each firm is either publicly listed or not. For example, some companies in the sample start out being publicly listed, and then are taken off the stock exchange (e.g., if they experience a leveraged buyout) and then are later made public again. It can occur that a company that is known to have been public in a specific year-quarter, has no ownership information in this year-quarter in the Thomson Reuters dataset. Where we have a public firm in the pair that has missing ownership data we remove this pair from the analysis. A total 6 markets are dropped due to missing ownership data, resulting in 451 drug markets.

We then match the brand drug products in our sample with Medicaid reimbursement data, publicly available from medicaid.gov, at the national level using National Drugs Codes (NDC) which are unique product identifiers for drugs in the US. A drug product in our sample may be matched with multiple NDC codes due to the fact we define drug products at the ingredient-form level, whereas NDC codes are defined at the finest level taking drug strength and package size into account. We aggregate information on the total amount reimbursed per year by summing over NDC codes for a drug product. Due to that fact that some drugs cannot be matched with Medicaid reimbursements, we are left with 395 unique drug product markets.

Subsidiary firms are assigned the ownership structure of the parent firm under the assumption that they are fully controlled by the parent. However in recognition of the fact that the subsidiary is a separate entity from the parent with its own previous experience, we determine all experience variables at the subsidiary level. That is, we do not assign the experience of the parent to the subsidiary. In the final dataset, there are 93 unique brand companies and 189 unique generic companies operating within the relevant markets and time period. Given that the focus of the article is on links between brand and generic companies, we then make our dataset pairwise; creating brand-generic pairs. There are 10,453 unique pairs.

The common ownership measures are constructed at the pair level using data from Thomson Reuters Global Ownership Database from 2003 to 2014. We calculate common ownership measures in the year of the end of exclusivity (lag 0), one year prior (lag 1) and two years prior (lag 2). When constructing measures of common ownership, we restrict ourselves to the investor holdings that represent at least one percent in the equity of the firms. Investor acquisitions during this period and ultimate owners are identified on the basis of public sources.

³⁴The main submission classes include Type 1 New Molecular Entity, Type 2 New Active Ingredient, Type 3 New Dosage Form, Type 4 New Combination, Type 5 New Formulation or Other Differences (e.g., new indication, new applicant, new manufacturer).

3.8.3 Formal model and proofs

In this Appendix, we characterize the equilibrium entry decisions of the N potential entrants, as a function of their symmetric “market-level” common ownership with the brand, δ . To this end, we first analyse the strategic interaction between generics’ entry decisions. All the proofs can be found in the last subsection of this Appendix.

Strategic effects: complements or substitutes?

For ease of illustration, let us restrict ourselves in this subsection to the case of $N = 2$ potential generic entrants.

We investigate if focal generic g is less (or equally) likely to enter as the probability p_1 of having a competing generic increases, and the probability p_0 of having none declines (“strategic substitutes”); or alternatively, if g is more (or equally) likely to enter as p_1 increases (“strategic complements”). Substituting $p_0 = 1 - p_1$ and deriving Π_g in (3.1) with respect to p_1 ,

$$\partial\Pi_g(p_0, p_1, \delta)/\partial p_1 = (\pi_g^1 - \pi_g^0) + \delta(\Delta\pi_b^1 - \Delta\pi_b^0),$$

we can identify two effects. The first term is negative, as $\pi_g^0 > \pi_g^1$, and therefore the gains from entry of g are lower if the other is more likely to enter. This is the traditional business stealing effect from competition of other generics. The second term, though, is positive, as $|\Delta\pi_b^0| > |\Delta\pi_b^1|$. As the other generic is more likely to enter, the effect of focal generic entry on the brand firm is less detrimental, as the reduction of brand profits in the presence of another competing generic is smaller.

The overall effect depends on which of the two effects, proxied by the profits of generic entrant π_g^k and the loss in profits of the brand $|\Delta\pi_b^k|$, decreases faster with the entry of others, and thus how the ratio $\bar{\delta}_k \equiv \pi_g^k/|\Delta\pi_b^k|$ changes with k . If the generic profits decrease faster, and thus the ratios are such that $\bar{\delta}_1 < \bar{\delta}_0$, others entering is more detrimental and entry decisions exhibit strategic substitutabilities. Instead, if the brand losses decrease faster, and thus $\bar{\delta}_0 < \bar{\delta}_1$, others entering is less detrimental and entry decisions exhibit strategic complementarities. The results are summarized in the following lemma.


Lemma 2. (a) *If $\bar{\delta}_1 < \bar{\delta}_0$, the generic firm g is less (or equally) likely to enter if the other generic firm is more likely to enter (strategic substitutability).*

(b) *If $\bar{\delta}_0 < \bar{\delta}_1$, the generic firm g is more (or equally) likely to enter if the other generic is more likely to enter (strategic complementarity).*

Figure 3.7 depicts the combinations of g ’s common ownership with the brand, δ , and probability of the other entering, p_1 , for which g ’s entry is profitable (marked in the darker shade in the figure); where the left panel shows the case of strategic substitutes and the right panel the case for strategic complements. Clearly, for a given p_1 , common ownership reduces entry profitability. But the effect of the probability of the other entering, p_1 , for a given level of common ownership δ has non-trivial effects on the profitability of entering. An increase in p_1 may mean that entry switches from profitable to unprofitable in the intermediate region of δ in the case of substitutes (the left-hand panel) whereas it may switch from unprofitable to profitable in the intermediate region of δ in the case of complements (the right-hand panel). Still, in both cases, entry is profitable for any p_1 if δ is sufficiently low, i.e. entering is a dominant strategy, whereas entry is unprofitable for any

p_1 if δ is sufficiently high, i.e. not entering is a dominant strategy.

Figure 3.7: Profitable entry of g as a function of δ and p_1



Equilibrium entry decisions


Now let us consider the pure-strategy equilibrium decisions in the general case of N potential entrants as a function of their symmetric level of common ownership with the brand, δ . Considering and distinguishing between the two cases identified in the previous proposition, the proposition summarizes the overall number of entrants in equilibrium.

Lemma 3. (a) In the case of strategic substitutes ($\bar{\delta}_{N-1} < \bar{\delta}_{N-2} < \dots < \bar{\delta}_0$), the number of entrants in equilibrium is: N if $\delta \leq \bar{\delta}_{N-1}$; $N - k$ if $\bar{\delta}_{N-k} < \delta \leq \bar{\delta}_{N-k-1}$ for $k = 1, \dots, N - 1$; and 0 if $\bar{\delta}_0 < \delta$.

(b) In the case of strategic complements ($\bar{\delta}_0 < \bar{\delta}_1 < \dots < \bar{\delta}_{N-1}$), the number of entrants in equilibrium is: N if $\delta \leq \bar{\delta}_0$; N or 0 if $\bar{\delta}_0 < \delta \leq \bar{\delta}_{N-1}$; and 0 if $\bar{\delta}_{N-1} < \delta$.

Figure 3.8 depicts the number of entrants in equilibrium as a function of their symmetric level of common ownership with the brand, δ . In both cases, there exists multiple equilibria in all the intermediate regions. But in the case of strategic substitutes, the equilibrium difference is between the identity of entrants and not how many of the entrants enter. In the case of complementarities, the equilibrium number of entrants is extreme, either none or all of them shall enter. This is because, in the case of substitutes, the entry of another generic makes generic entry less profitable, whereas in the case of complements, it makes it more profitable.

Figure 3.8: Number of entrants in equilibrium as a function of δ



Still, in both cases, the equilibrium number of entrants decreases with the level of common ownership, as long as we assign a fixed probability of selecting one equilibrium over another. This is a proof of the second statement in Proposition 1 of the main text.

Proof of Lemma 2

We determine the optimal entry decision of focal generic firm g for a given probability of entry of the other generic, p_1 , i.e. the best response function.

We first note that whether profits of the focal generic increase if the other is more likely to enter depends on the level of common ownership. Indeed, in the case where $N = 2$, we can write the profits as a function of just p_1 ,

$$\Pi_g(p_1, \delta) = (1 - p_1)(\pi_g^0 + \delta\Delta\pi_b^0) + p_1(\pi_g^1 + \delta\Delta\pi_b^1)$$

and, as displayed in the text,

$$\partial\Pi_g(p_1, \delta)/\partial p_1 = (\pi_g^1 - \pi_g^0) + \delta(\Delta\pi_b^1 - \Delta\pi_b^0).$$

As this function is strictly increasing in δ ($\partial^2\Pi_g(p_1, \delta)/\partial p_1\partial\delta = \Delta\pi_b^1 - \Delta\pi_b^0 > 0$), and it has a negative intercept ($\partial\Pi_g(p_1, 0)/\partial p_1 = \pi_g^1 - \pi_g^0 < 0$), there exists δ^* such that, if $\delta \leq \delta^*$, profits are decreasing in p_1 ($\partial\Pi_g(p_1, \delta)/\partial p_1 \leq 0$) whereas, if $\delta > \delta^*$, profits are increasing in p_1 ($\partial\Pi_g(p_1, \delta)/\partial p_1 > 0$), where

$$\delta^* \equiv -(\pi_g^1 - \pi_g^0)/(\Delta\pi_b^1 - \Delta\pi_b^0).$$

Second, we determine the optimal decision in cases where the other generic uses pure-strategies:

- If $p_1 = 0$ (i.e., it does not enter for sure), g shall it find it optimal to enter if $\delta \leq \bar{\delta}_0$ as $\Pi_g(0, \delta) = \pi_g^0 + \delta\Delta\pi_b^0 \geq 0$ if and only if

$$\delta \leq \pi_g^0/|\Delta\pi_b^0| \equiv \bar{\delta}_0.$$

- Similarly, if $p_1 = 1$ (i.e., it does enter for sure), g shall it find it optimal to enter if $\delta \leq \bar{\delta}_1$ as $\Pi_g(1, \delta) = \pi_g^1 + \delta\Delta\pi_b^1 \geq 0$ if and only if

$$\delta \leq \pi_g^1/|\Delta\pi_b^1| \equiv \bar{\delta}_1.$$

Simple algebra shows that if $\bar{\delta}_1 < \bar{\delta}_0$ then $\bar{\delta}_0 < \delta^*$ whereas if $\bar{\delta}_0 < \bar{\delta}_1$ then $\delta^* < \bar{\delta}_0$. These two cases affect the strategic interaction.

Let us now consider the best response function for different levels of common ownership, δ . We first show that, if $\bar{\delta}_1 < \bar{\delta}_0$ and thus $\bar{\delta}_1 < \bar{\delta}_0 < \delta^*$, focal generic g is less (or equally) likely to enter if p_1 is greater (termed “strategic substitutes”). Still, it may be that the generic’s profits increase with the entry of the other, as long as it does not affect the decision.

- If $\delta \leq \bar{\delta}_1$ then entering is a dominant strategy. Indeed, we have that $\delta < \delta^*$ and g is less likely to enter if the probability of entering of the other is greater ($\partial\Pi_g(p_1, \delta)/\partial p_1 < 0$). As $\delta \leq \bar{\delta}_1$, g should enter for any p_1 as $\Pi_g \geq 0$ even in the most adverse case, in which the other does enter for sure, $p_1 = 1$.
- In the case in which $\bar{\delta}_1 < \delta \leq \bar{\delta}_0$, the decision to enter depends on p_1 : g should enter if the probability of the other entering is low. In formal terms, $\Pi_g > 0$ if and only if $p_1 < p_1^*$ where p_1^* is such that $\Pi_g(p_1^*, \delta) = 0$. Notice that p_1^* is well defined, as $\Pi_g(0, \delta) > 0$ (as $\delta < \bar{\delta}_0$), $\partial\Pi_g(p_1, \delta)/\partial p_1 < 0$ (as $\delta < \delta^*$) and $\Pi_g(1, \delta) < 0$ (as $\delta > \bar{\delta}_1$). In addition, note that the

threshold level of p_1^* is decreasing in the level of common ownership,

$$\partial p_1^*/\partial \delta = -[\partial \Pi_g(p_1, \delta)/\partial \delta]/[\partial \Pi_g(p_1, \delta)/\partial p_1] < 0.$$

- If $\bar{\delta}_0 < \delta \leq \delta^*$, then not entering is a dominant strategy. Indeed, g should not enter for any p_1 as $\Pi_g < 0$ even in the most favorable case, in which the other does not enter for sure, $p_1 = 0$.
- In case the levels of common ownership δ are such that $\delta > \delta^*$ then not entering is dominant. In that case g is more likely to enter if the probability of entering of the other is greater ($\partial \Pi_g(p_1, \delta)/\partial p_1 > 0$), but g should not enter for any p_1 as $\Pi_g < 0$ even in the most favorable case, in which the other enters for sure, $p_1 = 1$ as $\delta > \bar{\delta}_1$.

Second, we show that, if $\bar{\delta}_0 < \bar{\delta}_1$ and thus $\delta^* \leq \bar{\delta}_0 < \bar{\delta}_1$, focal generic g is more (or equally as) likely to enter if p_1 is greater (labeled as “strategic complements”).

- In case the levels of common ownership δ are such that $\delta < \delta^*$ then entering is dominant. In that case g is less likely to enter if the probability of entering of the other is greater ($\partial \Pi_g(p_1, \delta)/\partial p_1 < 0$) but g should p_1 as $\Pi_g > 0$ even in the most adverse case, in which the other enters for sure, $p_1 = 1$ as $\delta < \bar{\delta}_1$.
- In the case in which $\delta^* < \delta \leq \bar{\delta}_0$, entering is dominant. Indeed as $\delta > \delta^*$ g is more likely to enter if the probability of entering of the other is greater ($\partial \Pi_g(p_1, \delta)/\partial p_1 > 0$). As $\delta < \bar{\delta}_0$ g should enter for any p_1 as $\Pi_g > 0$ even in the most adverse case, in which the other does not enter for sure, $p_1 = 0$.
- In the case in which $\bar{\delta}_0 < \delta \leq \bar{\delta}_1$, the decision to enter depends on p_1 : g should enter if the probability of the other entering is high. In formal terms, $\Pi_g > 0$ if and only if $p_1 > p_1^*$ where p_1^* is such that $\Pi_g(p_1^*, \delta) = 0$. Notice that p_1^* is well defined, as $\Pi_g(0, \delta) < 0$ (as $\delta > \bar{\delta}_0$), $\partial \Pi_g(p_1, \delta)/\partial p_1 > 0$ (as $\delta > \delta^*$) and $\Pi_g(1, \delta) > 0$ (as $\delta < \bar{\delta}_1$). In addition, note that the threshold level of p_1^* is decreasing in the level of common ownership,

$$\partial p_1^*/\partial \delta = -[\partial \Pi_g(p_1, \delta)/\partial \delta]/[\partial \Pi_g(p_1, \delta)/\partial p_1] > 0.$$

- If $\delta^* > \bar{\delta}_1$ g then not entering is dominant. Indeed g should not enter for any p_1 as $\Pi_g < 0$ even in the most favorable case, in which the other does enter for sure, $p_1 = 1$.

Proof of Lemma 3

We proceed in two steps. We first determine the optimal entry decision of focal generic firm g for each entry decision of the other $N - 1$ generics. That is, we compute, as in the previous proposition, the best response function (which depends again on the level of common ownership). But here, although allowing for N generics, we concentrate on pure strategies. As we assume generics to be symmetric, the key is how many, but not which one, of the others decide to enter. In a second step, we compute the (pure-strategy) Nash equilibria.

As in the previous proposition, in case k of the other entrants enter ($k = 0, \dots, N - 1$, $p_k = 1$ and, for any $j \neq k$, $p_j = 0$), g shall it find it optimal to enter if and only if $\delta \leq \bar{\delta}_k$ as $\Pi_g = \pi_g^k + \delta \Delta \pi_b^k \geq 0$ if and only if

$$\delta \leq \pi_g^k / |\Delta \pi_b^k| \equiv \bar{\delta}_k.$$

In the case of a single potential entrant ($N = 1$ and $k = 0$), this is the optimal decision: enter if $\delta \leq \bar{\delta}_0$ and do not if $\delta > \bar{\delta}_0$. In this case, parts (a) and (b) in the statement of the proposition are the same. From now on we consider $N > 1$.

Now let us consider the two cases of the statement of the proposition. Suppose first that $\bar{\delta}_{N-1} < \bar{\delta}_{N-2} < \dots < \bar{\delta}_0$ (“strategic substitutes”). The best response function of g with respect to the number of other entrants depends, as in the previous proposition, on the level of common ownership.

- If $\delta \leq \bar{\delta}_{N-1}$ entering is a dominant strategy for g , independent of the number of other entrants, as $\delta \leq \bar{\delta}_k$ for any k .
- If $\bar{\delta}_{N-k} < \delta \leq \bar{\delta}_{N-k-1}$ for any $k = 1, \dots, N-1$, g shall enter if $N-k-1$ other generics, or less, enter, as $\delta \leq \bar{\delta}_{N-k-1} < \dots < \bar{\delta}_0$, but it shall not enter if $N-k$ other generics, or more, do enter, as $\bar{\delta}_{N-1} < \dots < \bar{\delta}_{N-k} \leq \delta$.
- Finally, if $\delta > \bar{\delta}_0$ not entering is a dominant strategy, as $\delta > \bar{\delta}_k$ for any k .

For instance in the case of two potential entrants ($N = 2$), g should enter if $\delta \leq \bar{\delta}_1$, enter if and only if the other does not enter if $\bar{\delta}_1 < \delta \leq \bar{\delta}_0$ (as $N = 2$, $k = 1$, $N-k-1 = 0$ and $N-k = 1$) and not enter if $\delta > \bar{\delta}_0$.

The equilibrium number of entrants also depends on the (symmetric) level of common ownership with the brand.

- If $\delta \leq \bar{\delta}_{N-1}$ all should enter in equilibrium, as entering is a dominant strategy.
- If $\bar{\delta}_{N-k} < \delta \leq \bar{\delta}_{N-k-1}$ for any $k = 1, \dots, N-1$, $N-k$ generics should enter in equilibrium, as entering is optimal if $N-k-1$ other generics enter and not entering is optimal if $N-k$ do so.
- Finally, if $\delta > \bar{\delta}_0$ none of them should enter as not entering is a dominant strategy.

For instance in the case of two potential entrants ($N = 2$, which implies $k = 1$), the two generics should enter if $\delta \leq \bar{\delta}_1$, one of them should enter if $\bar{\delta}_1 < \delta \leq \bar{\delta}_0$ (as $N = 2$, $k = 1$ and $N-k = 1$) and none of them should enter if $\delta > \bar{\delta}_0$.

Suppose now that $\bar{\delta}_0 < \bar{\delta}_1 < \dots < \bar{\delta}_{N-1}$ (“strategic complements”). The best response function of g with respect to the number of other entrants is now as follows:

- If $\delta \leq \bar{\delta}_0$ entering is again a dominant strategy for g , as $\delta < \bar{\delta}_k$ for any k .
- But now, if $\bar{\delta}_{N-k-1} < \delta \leq \bar{\delta}_{N-k}$ for any $k = 1, \dots, N-1$, g shall enter if $N-k$ other generics, or more, enter, as $\delta \leq \bar{\delta}_{N-k} < \dots < \bar{\delta}_{N-1}$, but it shall not enter if $N-k-1$ other generics, or less, do enter, as $\bar{\delta}_0 < \dots < \bar{\delta}_{N-k-1} < \delta$.
- Similarly, if $\delta > \bar{\delta}_{N-1}$ not entering is again a dominant strategy, as $\delta > \bar{\delta}_k$ for any k .

For instance in the case of two potential entrants ($N = 2$), g should enter if $\delta \leq \bar{\delta}_0$, enter if and only if the other does enter if $\bar{\delta}_0 < \delta \leq \bar{\delta}_1$ and not enter if $\delta > \bar{\delta}_1$.

The equilibrium number of entrants also depends on the (symmetric) level of common ownership with the brand.

- As before, if $\delta \leq \bar{\delta}_0$ all should enter in equilibrium, as entering is a dominant strategy.
- But the equilibria in the intermediate cases $\bar{\delta}_0 < \delta \leq \bar{\delta}_{N-1}$ are different: either all the N generics enter or none of them does. Indeed, if $N - 1$ generics enter, it is optimal to enter, as $\delta \leq \bar{\delta}_{N-1}$, and if 0 of them does, it is optimal not to enter either, as $\delta > \bar{\delta}_0$. Moreover, there is no equilibrium within $\bar{\delta}_0 < \delta \leq \bar{\delta}_{N-1}$ in which k generics enter, for k is such that $0 < k < N$. Indeed, if an entrant finds it profitable to enter then it should also be profitable for those that do not enter (and if one of the non-entrants find it profitable not to enter then it should also be non-profitable for one of the entrants).
- Finally, if $\delta > \bar{\delta}_{N-1}$ none of them should enter as not entering is a dominant strategy.

In the case of two potential entrants ($N = 2$, which implies $k = 1$), the two generics should enter if $\delta \leq \bar{\delta}_0$, the two or none of them should enter if $\bar{\delta}_0 < \delta \leq \bar{\delta}_1$ and none of them should enter if $\delta > \bar{\delta}_1$.

3.8.4 Robustness

Table 3.12: Robustness - Drug market fixed effects

VARIABLES	(1)	(2)	(3)	(4)	(5)
δ_S	-0.0220*** (0.00613)				
δ_C		-0.0649*** (0.0176)			
δ_L			-0.0240*** (0.00527)		
δ_{top5}				-0.00508*** (0.00139)	
δ_{top10}					-0.00269*** (0.000750)
Subsidiary (0/1)	-0.0481*** (0.0142)	-0.0467*** (0.0142)	-0.0475*** (0.0142)	-0.0240 (0.0150)	-0.0241 (0.0150)
Experience Route \div 10	0.00825*** (0.000680)	0.00824*** (0.000680)	0.00827*** (0.000680)	0.00823*** (0.000680)	0.00825*** (0.000680)
Experience ATC2 \div 10	0.0629*** (0.00810)	0.0630*** (0.00810)	0.0628*** (0.00810)	0.0631*** (0.00809)	0.0631*** (0.00810)
Experience New Drug \div 10	0.00408 (0.00296)	0.00375 (0.00290)	0.00426 (0.00293)	0.00373 (0.00288)	0.00395 (0.00294)
Breadth (ATC2) \div 10	0.00199 (0.00150)	0.00202 (0.00151)	0.00191 (0.00150)	0.00187 (0.00150)	0.00197 (0.00150)
Observations	34144	34144	34144	34144	34144
R-squared	0.121	0.121	0.121	0.121	0.121
Generic region of origin	Yes	Yes	Yes	Yes	Yes
Drug product fixed effect	Yes	Yes	Yes	Yes	Yes
Drug markets	395	395	395	395	395

Notes: OLS regression. Standard errors in parentheses are robust. The dependent variable is entry within 6 quarters. The constant term is estimated but not reported. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.13: Robustness - Probit

VARIABLES	(1)	(2)	(3)	(4)	(5)
δ_S	-0.426*** (0.126)				
δ_C		-1.212*** (0.356)			
δ_L			-0.396*** (0.125)		
δ_{top5}				-0.0842*** (0.0291)	
δ_{top10}					-0.0462*** (0.0147)
Subsidiary (0/1)	-0.918* (0.508)	-0.891* (0.505)	-0.890* (0.505)	-0.478 (0.512)	-0.459 (0.511)
Market Size	0.360*** (0.0885)	0.358*** (0.0886)	0.358*** (0.0881)	0.357*** (0.0881)	0.358*** (0.0887)
Authorized Generic (0/1)	0.117* (0.0648)	0.118* (0.0649)	0.118* (0.0649)	0.116* (0.0650)	0.116* (0.0650)
Orphan Drug (0/1)	-0.0419 (0.109)	-0.0390 (0.109)	-0.0427 (0.109)	-0.0422 (0.109)	-0.0409 (0.109)
Pediatric Drug (0/1)	0.199*** (0.0738)	0.199*** (0.0739)	0.197*** (0.0740)	0.198*** (0.0744)	0.201*** (0.0743)
Substitutes on Patent (ATC2) \div 10	-0.0636 (0.0879)	-0.0624 (0.0875)	-0.0635 (0.0875)	-0.0602 (0.0876)	-0.0626 (0.0876)
Substitutes off Patent (ATC2) \div 10	-0.0567 (0.0758)	-0.0551 (0.0761)	-0.0576 (0.0757)	-0.0596 (0.0761)	-0.0575 (0.0762)
Experience Route \div 10	0.0742*** (0.00586)	0.0744*** (0.00586)	0.0742*** (0.00589)	0.0740*** (0.00591)	0.0743*** (0.00586)
Experience ATC2 \div 10	0.386*** (0.0571)	0.389*** (0.0571)	0.386*** (0.0571)	0.388*** (0.0570)	0.390*** (0.0570)
Experience New Drug \div 10	-0.0874** (0.0376)	-0.0908** (0.0373)	-0.0902** (0.0370)	-0.0954*** (0.0368)	-0.0926** (0.0374)
Breadth (ATC2) \div 10	0.217*** (0.0275)	0.221*** (0.0275)	0.214*** (0.0274)	0.211*** (0.0275)	0.217*** (0.0276)
Observations	32994	32994	32994	32994	32994
Therapeutic field	Yes	Yes	Yes	Yes	Yes
Drug form	Yes	Yes	Yes	Yes	Yes
Submission type	Yes	Yes	Yes	Yes	Yes
Generic region of origin	Yes	Yes	Yes	Yes	Yes
Year end of exclusivity	Yes	Yes	Yes	Yes	Yes
Drug markets	395	395	395	395	395

Notes: Probit regression. Standard errors in parentheses are clustered at the drug market level. The dependent variable is entry within 6 quarters. The constant term is estimated but not reported. ** $p < 0.01$, * $p < 0.05$, * $p < 0.1$.

Table 3.14: Robustness - Logit

VARIABLES	(1)	(2)	(3)	(4)	(5)
δ_S	-0.893*** (0.267)				
δ_C		-2.691*** (0.744)			
δ_L			-0.957*** (0.266)		
δ_{top5}				-0.182*** (0.0613)	
δ_{top10}					-0.0965*** (0.0308)
Subsidiary (0/1)	-2.232** (1.116)	-2.160* (1.113)	-2.170* (1.113)	-1.270 (1.135)	-1.264 (1.133)
Market Size	0.778*** (0.191)	0.775*** (0.191)	0.778*** (0.190)	0.772*** (0.190)	0.773*** (0.191)
Authorized Generic (0/1)	0.251* (0.140)	0.251* (0.140)	0.251* (0.140)	0.247* (0.140)	0.246* (0.140)
Orphan Drug (0/1)	-0.0628 (0.235)	-0.0523 (0.235)	-0.0645 (0.236)	-0.0639 (0.236)	-0.0605 (0.235)
Pediatric Drug (0/1)	0.462*** (0.163)	0.461*** (0.163)	0.459*** (0.163)	0.459*** (0.164)	0.465*** (0.164)
Substitutes on Patent (ATC2) \div 10	-0.125 (0.192)	-0.121 (0.192)	-0.126 (0.191)	-0.120 (0.191)	-0.123 (0.191)
Substitutes off Patent (ATC2) \div 10	-0.146 (0.169)	-0.144 (0.170)	-0.148 (0.169)	-0.153 (0.170)	-0.148 (0.171)
Experience Route \div 10	0.149*** (0.0119)	0.150*** (0.0120)	0.150*** (0.0121)	0.149*** (0.0121)	0.149*** (0.0120)
Experience ATC2 \div 10	0.685*** (0.119)	0.690*** (0.119)	0.685*** (0.119)	0.690*** (0.118)	0.692*** (0.118)
Experience New Drug \div 10	-0.187** (0.0815)	-0.190** (0.0810)	-0.184** (0.0804)	-0.202** (0.0802)	-0.198** (0.0812)
Breadth (ATC2) \div 10	0.528*** (0.0588)	0.536*** (0.0586)	0.522*** (0.0586)	0.518*** (0.0589)	0.527*** (0.0588)
Observations	32994	32994	32994	32994	32994
Therapeutic field	Yes	Yes	Yes	Yes	Yes
Drug form	Yes	Yes	Yes	Yes	Yes
Submission type	Yes	Yes	Yes	Yes	Yes
Generic region of origin	Yes	Yes	Yes	Yes	Yes
Year end of exclusivity	Yes	Yes	Yes	Yes	Yes
Drug markets	395	395	395	395	395

Notes: Logit regression. Standard errors in parentheses are clustered at the drug market level. The dependent variable is entry within 6 quarters. The constant term is estimated but not reported.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.15: Robustness - Indicator for top 100 sales

VARIABLES	(1)	(2)	(3)	(4)	(5)
δ_S	-0.0276*** (0.00678)				
δ_C		-0.0789*** (0.0184)			
δ_L			-0.0286*** (0.00560)		
δ_{top5}				-0.00581*** (0.00142)	
δ_{top10}					-0.00301*** (0.000811)
Subsidiary (0/1)	-0.0549*** (0.0145)	-0.0535*** (0.0145)	-0.0540*** (0.0144)	-0.0271* (0.0148)	-0.0276* (0.0149)
Top 100 in Sales (0/1)	0.0208*** (0.00557)	0.0208*** (0.00556)	0.0207*** (0.00557)	0.0206*** (0.00558)	0.0208*** (0.00558)
Authorized Generic (0/1)	0.00260 (0.00401)	0.00268 (0.00402)	0.00267 (0.00401)	0.00261 (0.00402)	0.00260 (0.00402)
Orphan Drug (0/1)	-0.00417 (0.00703)	-0.00405 (0.00701)	-0.00416 (0.00704)	-0.00407 (0.00703)	-0.00408 (0.00703)
Pediatric Drug (0/1)	0.0129*** (0.00481)	0.0128*** (0.00481)	0.0129*** (0.00481)	0.0129*** (0.00483)	0.0129*** (0.00483)
Substitutes on Patent (ATC2) \div 10	-0.00800 (0.00616)	-0.00790 (0.00615)	-0.00789 (0.00615)	-0.00774 (0.00616)	-0.00788 (0.00617)
Substitutes off Patent (ATC2) \div 10	-0.00421 (0.00498)	-0.00418 (0.00499)	-0.00425 (0.00498)	-0.00436 (0.00500)	-0.00426 (0.00500)
Experience Route \div 10	0.00838*** (0.000856)	0.00839*** (0.000856)	0.00842*** (0.000857)	0.00834*** (0.000854)	0.00838*** (0.000857)
Experience ATC2 \div 10	0.0608*** (0.0104)	0.0609*** (0.0104)	0.0608*** (0.0104)	0.0609*** (0.0104)	0.0610*** (0.0104)
Experience New Drug \div 10	0.00457 (0.00287)	0.00413 (0.00283)	0.00466 (0.00285)	0.00384 (0.00280)	0.00410 (0.00285)
Breadth (ATC2) \div 10	0.00101 (0.00231)	0.00116 (0.00230)	0.000948 (0.00229)	0.000664 (0.00230)	0.000898 (0.00230)
Observations	34144	34144	34144	34144	34144
R-squared	0.0854	0.0854	0.0855	0.0853	0.0852
Therapeutic field	Yes	Yes	Yes	Yes	Yes
Drug form	Yes	Yes	Yes	Yes	Yes
Submission type	Yes	Yes	Yes	Yes	Yes
Generic region of origin	Yes	Yes	Yes	Yes	Yes
Year end of exclusivity	Yes	Yes	Yes	Yes	Yes
Drug markets	395	395	395	395	395

Notes: OLS regression. Standard errors in parentheses are clustered at the drug market level. The dependent variable is entry within 6 quarters. The constant term is estimated but not reported. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.16: Robustness - Using total US sales for drugs in top 100

VARIABLES	(1)	(2)	(3)	(4)	(5)
δ_S	-0.0902*** (0.0140)				
δ_C		-0.229*** (0.0396)			
δ_L			-0.0829*** (0.0121)		
δ_{top5}				-0.0195*** (0.00331)	
δ_{top10}					-0.0112*** (0.00176)
Subsidiary (0/1)	-0.121*** (0.0382)	-0.114*** (0.0382)	-0.117*** (0.0378)	-0.0371 (0.0354)	-0.0251 (0.0380)
Brand Sales USD bn.	0.0121** (0.00562)	0.0117** (0.00562)	0.0123** (0.00565)	0.0123** (0.00562)	0.0121** (0.00565)
Authorized Generic (0/1)	0.00741 (0.00850)	0.00682 (0.00859)	0.00790 (0.00857)	0.00749 (0.00862)	0.00700 (0.00850)
Orphan Drug (0/1)	0.0110 (0.0125)	0.0127 (0.0126)	0.0103 (0.0127)	0.0101 (0.0128)	0.0110 (0.0126)
Pediatric Drug (0/1)	0.0309*** (0.0112)	0.0317*** (0.0113)	0.0310*** (0.0114)	0.0314*** (0.0115)	0.0320*** (0.0112)
Substitutes on Patent (ATC2) \div 10	0.00306 (0.0101)	0.00440 (0.0102)	0.00277 (0.0102)	0.00303 (0.0102)	0.00280 (0.0101)
Substitutes off Patent (ATC2) \div 10	0.0321 (0.0219)	0.0319 (0.0220)	0.0327 (0.0222)	0.0333 (0.0224)	0.0320 (0.0219)
Experience Route \div 10	0.0143*** (0.00220)	0.0144*** (0.00221)	0.0145*** (0.00221)	0.0143*** (0.00220)	0.0144*** (0.00221)
Experience ATC2 \div 10	0.0669** (0.0266)	0.0681** (0.0266)	0.0670** (0.0267)	0.0679** (0.0267)	0.0674** (0.0266)
Experience New Drug \div 10	0.0268*** (0.00802)	0.0241*** (0.00783)	0.0259*** (0.00794)	0.0236*** (0.00772)	0.0262*** (0.00793)
Breadth (ATC2) \div 10	-0.00623 (0.00570)	-0.00582 (0.00581)	-0.00656 (0.00571)	-0.00705 (0.00573)	-0.00620 (0.00570)
Observations	8600	8600	8600	8600	8600
R-squared	0.152	0.152	0.152	0.152	0.153
Therapeutic field	Yes	Yes	Yes	Yes	Yes
Drug form	Yes	Yes	Yes	Yes	Yes
Submission type	Yes	Yes	Yes	Yes	Yes
Generic region of origin	Yes	Yes	Yes	Yes	Yes
Year end of exclusivity	Yes	Yes	Yes	Yes	Yes
Drug markets	93	93	93	93	93

Notes: OLS regression. Standard errors in parentheses are clustered at the drug market level. The dependent variable is entry within 6 quarters. The constant term is estimated but not reported. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.17: Robustness - Broader entrant set

VARIABLES	(1)	(2)	(3)	(4)	(5)
δ_S	-0.0143*** (0.00439)				
δ_C		-0.0443*** (0.0127)			
δ_L			-0.0179*** (0.00390)		
δ_{top5}				-0.00378*** (0.000957)	
δ_{top10}					-0.00173*** (0.000552)
Subsidiary (0/1)	-0.0337*** (0.00885)	-0.0332*** (0.00883)	-0.0336*** (0.00884)	-0.0157* (0.00925)	-0.0176* (0.00955)
Market Size	0.0282*** (0.00653)	0.0282*** (0.00652)	0.0282*** (0.00652)	0.0282*** (0.00652)	0.0282*** (0.00654)
Authorized Generic (0/1)	0.00419 (0.00254)	0.00421* (0.00255)	0.00421* (0.00254)	0.00419 (0.00255)	0.00419 (0.00255)
Orphan Drug (0/1)	-0.00263 (0.00399)	-0.00257 (0.00399)	-0.00262 (0.00399)	-0.00259 (0.00399)	-0.00260 (0.00399)
Pediatric Drug (0/1)	0.00790*** (0.00298)	0.00787*** (0.00298)	0.00792*** (0.00298)	0.00792*** (0.00299)	0.00793*** (0.00299)
Substitutes on Patent (ATC2) \div 10	-0.00245 (0.00346)	-0.00242 (0.00346)	-0.00242 (0.00346)	-0.00235 (0.00346)	-0.00242 (0.00346)
Substitutes off Patent (ATC2) \div 10	-0.00264 (0.00254)	-0.00262 (0.00255)	-0.00266 (0.00254)	-0.00270 (0.00255)	-0.00264 (0.00255)
Experience Route \div 10	0.00837*** (0.000791)	0.00838*** (0.000792)	0.00839*** (0.000792)	0.00836*** (0.000790)	0.00837*** (0.000792)
Experience ATC2 \div 10	0.0559*** (0.00950)	0.0560*** (0.00950)	0.0559*** (0.00950)	0.0559*** (0.00951)	0.0560*** (0.00951)
Experience New Drug \div 10	0.00447** (0.00202)	0.00429** (0.00198)	0.00474** (0.00199)	0.00437** (0.00196)	0.00439** (0.00201)
Breadth (ATC2) \div 10	1.06e-05 (0.00168)	0.000133 (0.00169)	9.73e-05 (0.00168)	-6.88e-05 (0.00169)	1.74e-05 (0.00169)
Observations	55769	55769	55769	55769	55769
R-squared	0.0817	0.0818	0.0819	0.0818	0.0817
Therapeutic field	Yes	Yes	Yes	Yes	Yes
Drug form	Yes	Yes	Yes	Yes	Yes
Submission type	Yes	Yes	Yes	Yes	Yes
Generic region of origin	Yes	Yes	Yes	Yes	Yes
Year end of exclusivity	Yes	Yes	Yes	Yes	Yes
Drug markets	395	395	395	395	395

Notes: OLS regression. Standard errors in parentheses are clustered at the drug market level. The dependent variable is entry within 6 quarters. The constant term is estimated but not reported. The sample of potential generic entrants includes all pharmaceutical companies that launched at least one generic product in our drug markets. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.18: Robustness - Entry ever

VARIABLES	(1)	(2)	(3)	(4)	(5)
δ_S	-0.0407*** (0.00848)				
δ_C		-0.122*** (0.0232)			
δ_L			-0.0424*** (0.00688)		
δ_{top5}				-0.00896*** (0.00170)	
δ_{top10}					-0.00526*** (0.00101)
Subsidiary (0/1)	-0.0783*** (0.0163)	-0.0764*** (0.0164)	-0.0770*** (0.0163)	-0.0356** (0.0166)	-0.0316* (0.0168)
Market Size	0.0465*** (0.0117)	0.0464*** (0.0117)	0.0465*** (0.0116)	0.0464*** (0.0116)	0.0465*** (0.0117)
Authorized Generic (0/1)	0.0111** (0.00500)	0.0112** (0.00501)	0.0111** (0.00501)	0.0110** (0.00501)	0.0110** (0.00501)
Orphan Drug (0/1)	0.00319 (0.00918)	0.00339 (0.00915)	0.00319 (0.00918)	0.00333 (0.00917)	0.00340 (0.00918)
Pediatric Drug (0/1)	0.0182*** (0.00598)	0.0181*** (0.00597)	0.0181*** (0.00598)	0.0182*** (0.00600)	0.0184*** (0.00600)
Substitutes on Patent (ATC2) \div 10	-0.00136 (0.00836)	-0.00122 (0.00836)	-0.00123 (0.00835)	-0.00102 (0.00837)	-0.00120 (0.00838)
Substitutes off Patent (ATC2) \div 10	-0.0143** (0.00604)	-0.0142** (0.00604)	-0.0143** (0.00602)	-0.0145** (0.00604)	-0.0143** (0.00605)
Experience Route \div 10	0.00998*** (0.00101)	0.0100*** (0.00101)	0.0100*** (0.00101)	0.00993*** (0.00101)	0.00999*** (0.00101)
Experience ATC2 \div 10	0.0840*** (0.0119)	0.0842*** (0.0119)	0.0839*** (0.0118)	0.0842*** (0.0119)	0.0843*** (0.0119)
Experience New Drug \div 10	0.00768** (0.00331)	0.00723** (0.00328)	0.00785** (0.00328)	0.00677** (0.00325)	0.00770** (0.00331)
Breadth (ATC2) \div 10	0.00542** (0.00255)	0.00572** (0.00255)	0.00533** (0.00253)	0.00495* (0.00254)	0.00551** (0.00255)
Observations	34144	34144	34144	34144	34144
R-squared	0.102	0.102	0.102	0.102	0.102
Therapeutic field	Yes	Yes	Yes	Yes	Yes
Drug form	Yes	Yes	Yes	Yes	Yes
Submission type	Yes	Yes	Yes	Yes	Yes
Generic region of origin	Yes	Yes	Yes	Yes	Yes
Year end of exclusivity	Yes	Yes	Yes	Yes	Yes
Drug markets	395	395	395	395	395

Notes: OLS regression. Standard errors in parentheses are clustered at the drug market level. The dependent variable is entry (within any time period). The constant term is estimated but not reported. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Chapter 4

Common Ownership in the US Pharmaceutical Industry: A Network Analysis¹

Chapter Abstract

This chapter investigates patterns in common ownership networks between firms that are active in the US pharmaceutical industry for the period 2004-2014. The main findings are that “brand firms” - i.e. firms that have R&D capabilities and launch new drugs - exhibit relatively dense common ownership networks with each other that further increase significantly in density over time, whereas the network of “generic firms” - i.e. firms that primarily specialize in developing and launching generic drugs - is much sparser and stays that way over the span of our sample. Finally, when considering the common ownership links between brands firms, on the one hand, and generic firms, on the other, we find that brand firms have become more connected to generic firms over time. The implications of these findings for competition in the industry are discussed.

4.1 Introduction

Investors’ holdings in multiple firms give rise to what is known as “common ownership.” Common ownership is widespread in the US pharmaceutical industry. In 2014, for instance, the largest investor in the three largest pharmaceutical companies (Johnson & Johnson, Merck & Co and Pfizer) was the same (BlackRock). This is

¹This chapter is published as: Banal-Estañol, A., Newham, M. and Seldeslachts, J. (2021) Common Ownership in the US Pharmaceutical Industry: A Network Analysis. *The Antitrust Journal*, 66(1). We thank Einer Elhauge, Sumit Majumdar and Martin Schmalz for their insightful comments. We further thank Jonas Nieto for his excellent research assistance.

the rule, not the exception. These three pharmaceutical companies share other large institutional investors, and are thus connected to each other, as well as to numerous other pharmaceutical companies, through so-called “common ownership links.”²

Common ownership links between pharmaceutical companies might have important implications for competition and innovation in this crucial industry. By bringing innovative treatments to the market, or by making treatments more widely accessible, the pharmaceutical industry makes an important contribution to global health and economic development. At the same time, the industry often generates controversies related to pricing and product development. A well-functioning pharmaceutical industry in general, and the consequences of common ownership in particular, are thus key concerns for policy making and antitrust.

In this article we study the common ownership links between firms that are active in US pharmaceutical markets in the period 2004-2014 and discuss the implications of our findings for innovation incentives, entry, pricing and collusion. There is both anecdotal and empirical evidence, reported further below, showing that large institutional investors weigh in on pharmaceutical companies’ strategic decision-making. Given that these investors are both influential and, as we will show, have ownership stakes in multiple firms within the same market, the common ownership links between pharmaceutical companies could have important implications for competition and innovation.

We make use of network analysis to describe the structure and characteristics of common ownership networks and calculate how central, or influential, actors are in the network.³ We make a distinction between “brand firms”, that have R&D capabilities and launch new drugs on to the market, and “generic firms,” that produce bioequivalent replications of brand-name drugs once these drugs come off patent. We study the evolution of common ownership networks between brand firms and generic firms separately, as well as the (bipartite) network of brand firms on the one hand and generic firms on the other. We make use of two common ownership measures, which determine links on the basis of individual or joint levels of ownership by common investors. An *individual* common ownership link between two companies occurs when there is *at least one investor* in both companies with an ownership stake

²Institutional investors manage other people’s money by buying equity in companies (such as pension funds, sovereign wealth funds, insurance companies and investment funds). They typically seek to build diversified portfolios by investing in multiple companies, often within the same industry.

³There are surprisingly few papers that make use of network analysis to study common ownership patterns. A notable exception is Vitali et al. (2011) who use network analysis to study investor networks in a large sample of transnational corporations. Network analysis has been applied to other settings in the academic literature e.g. networks in the venture capital industry see Hochberg et al. (2007); interorganizational ties see Mizruchi et al. (1993); and networks between US firms that advocate for free trade see Dreiling & Darves (2011).

of more than 5%. A *joint* common ownership link occurs when investors common to both firms *collectively* are the majority owners.

We find that, although brand companies are already fairly well connected at the start of our sample, they become almost fully connected through common ownership links at the end of the sample. This is true for both measures of common ownership, although we observe a less dramatic change when using the joint measure, in part because the network was already highly connected at the beginning of the sample. If large institutional investors do exert influence, as the anecdotal evidence below indicates, then this increasing connectivity may have a non-negligible and increasing impact on innovation incentives. If institutional investors effectively assert their power in pharmaceutical companies, this increasingly dense network might further lead to a softening of competition between brand firms' products. Furthermore, as the evolution of the network partly depends on the ownership measure used, the effects of common ownership might depend on whether common investors exert individual or joint influence.

Alongside higher levels of connectivity between brand firms, the average measure of centrality, which indicates how influential individual firms are within the common ownership network, has risen. Interestingly, at the beginning of the sample, the most central firms were not necessarily the largest (e.g. Biogen and Allergan). On the contrary, the most central firms towards the end of the sample are also the largest (e.g. Johnson & Johnson).

The network of brand companies remains, even at the end of the sample, relatively asymmetric. Indeed, some of the largest pharmaceutical companies, such as Sanofi, Novartis and Roche, remain without any strong links in 2014. This is in part because of the presence of large non-common investors in these companies. Although several brand companies, such as Johnson & Johnson and Pfizer, have a large and similar centrality value in 2014, several others have low values (or even zero). Thus, brand firm centrality has not only increased over time, as the common ownership network has become more connected, but it has also become more dispersed. The combination of a rise in centrality for the most connected companies and, at the same time, higher dispersion overall might result in these central players becoming even more powerful.

In comparison to the brand network, the generic firm network is much sparser and it becomes less connected over time. Further, as compared to brand companies, the size of the shareholdings of the top common investors in generic companies - although larger in 2004 - is smaller in 2014. Consequently, the average level of centrality for generic firms is much lower than the average for brand firms at the end of the sample. While this is unlikely to have an impact on innovation - generic

companies mainly imitate brand products - it indicates that competition between generics is less affected by common ownership.

Finally, the number of common ownership links between brand companies, on the one hand, and generic companies, on the other, has increased substantially over time. Most brand-generic pairs were not connected at the beginning of the sample, and even some of the largest brands, such as Pfizer, had zero connections with the generics. At the end of our sample there are a number of strong connections between brands and generics. Most of the large brands, such as Johnson & Johnson and Pfizer, have a large number of links by 2014. Similarly, some of the generics, such as Impax and Perrigo, have a high number of connections with brand firms, despite having limited links between each other, and with other firms within the generic ownership network. The increased brand-generic connectivity seems to have led to a decrease in generic entry, as common investors have both an incentive and the ability to delay or block generics from entering the market of a brand.⁴

This paper is structured as follows: Section 4.2 provides a background of the pharmaceutical industry and provides anecdotal evidence of investors' influence in the pharmaceutical industry. Section 4.3 presents our data and a descriptive analysis. Section 4.4 undertakes a network analysis of the common ownership links in the pharmaceutical industry. Section 4.5 lays out the antitrust implications of common ownership in the pharmaceutical industry. Section 4.6 concludes.

4.2 Background

Before analysing common ownership patterns and their implications, this section provides a brief overview of the typical pharmaceutical “life-cycle” which is important for understanding how the industry, and thus how competition and innovation therein, works.⁵ We then provide a definition of common ownership, and a few examples. Finally, we report anecdotal and empirical evidence illustrating that common investors weigh in on pharma companies' strategic decisions.

4.2.1 Pharmaceutical industry

To bring new drugs to the market pharmaceutical firms must make significant investments in research and development. In the early stages of drug development, pharmaceutical companies engage in “drug discovery” to search for and discover new compounds to treat a specific disease. Given the public nature of the drug

⁴The impact of brand-generic links through common ownership on generic entry is confirmed in Newham et al. (2018).

⁵For a more detailed overview see Lakdawalla (2018).

approval process, patents are typically issued on novel pharmacological compounds quite early in the drug development process. They cover the active compound in a specific formulation and for specific indications.

After many iterations, the final compound becomes a drug candidate. Thereafter, with one or more optimized compounds in hand, researchers turn their attention to extensive preclinical testing. In pre-clinical tests the compound is tested for toxicity and safety. After completion of pre-clinical tests pharmaceutical firms prepare for the next critical stage in the innovation process - drug development through clinical trials on humans. To be considered for FDA approval a drug must pass through three “phases” of clinical trials. This is typically a lengthy and expensive process

In general, the R&D process for each drug is centered around its intended therapeutic area - the disease the drug should target (e.g. Diabetes type II) - and its ‘Mechanism of Action’ (MoA) - the biochemical process through which the drug produces the desired effect (e.g. SGLT2 inhibitors). The combination of the MoA within a therapeutic area has been used in practice to define “relevant markets” in competition enforcement - both at the innovation and launched product stages - as drugs herein can typically be substituted by general practitioners and patients.

During the process of drug research and development there is competition in the “innovation space.” Pharmaceutical companies engage in a race with other firms who are working on compounds to treat the same disease with a similar MoA. As rivals are often working in parallel on similar targets, often applying the same fundamental knowledge sourced from open science, the solutions they come up with may be similar. Pharmaceutical companies typically want to be the first to market with a drug that uses a new type of technology in order to profit from a first-to-market advantage.

Companies that produce novel drugs must apply for FDA approval through the new drug application (NDA) procedure. Drugs that are declared safe and effective, and are successfully approved by the FDA, are then launched on the market. Novel “brand-name” drugs are afforded a number of regulatory protections, including the patent on the key compound or active ingredient, which provide the company with a monopoly for their specific drug for a period of time. Nevertheless, once in the market, the drug will compete with other treatments that are substitutable from a therapeutic perspective, although not identical (“brand-brand” and/or “intermolecular” competition).

Once the regulatory protections afforded to the drug have expired, the market is open for generic entry. Generic firms produce bioequivalent copies of brand drugs and are typically much lower in price. The process by which generic manufacturers can seek approval from the FDA is set out in the Hatch-Waxman Act. The act

allows the generic applicant to apply for FDA approval by filing an abbreviated new drug application (ANDA) whereby the generic applicant can rely on the efficacy and safety data generated by the original innovator. The Hatch–Waxman Act also provides incentives for generic manufacturers to challenge patents in court, under “Paragraph IV.” Once launched on the market, generic drugs compete directly with the brand drug as they are essentially the same product (“intramolecular” competition) (Regan, 2008). In our analysis we distinguish between “brand firms” that have R&D capabilities and launch new drugs, and “generic firms” that primarily specialize in generic drugs.

In the US, drug prices are negotiated on between individual health insurance plans and the pharmaceutical company. While consumers may face some out-of-pocket expenditures for drugs, the cost of medical treatments is primarily paid by health insurance companies. High prescription drug prices are a concern for policy makers.⁶ A number studies do not find that “brand-brand” competition effectively lowers list prices (Sarpatwari et al., 2019) Generic competition, on the other hand is crucial for lowering prices. For products with a single generic producer, the generic average market price is 39% lower than the brand average market price before generic competition. With six or more competitors, generic prices show price reductions of more than 95% compared to brand prices⁷ Accordingly, promoting generic entry is an important policy goal for the FDA.⁸

4.2.2 Institutional investors and common ownership

Common ownership exists when an investor has a stake in two or more firms. Table 4.1 shows the top five investors in the three largest pharmaceutical companies - which are all brand firms - that operate in US markets in the period 2004-2014 (see the data section below for more details on our sample). From this table it is clear that there are a number of institutional investors, such as Vanguard and State Street, that are common owners with shareholdings in all three firms in both 2004 and 2014. BlackRock holds the number one position, with a stake of 5-7%, in all

⁶E.g. see Kuchler, H. Why prescription drugs cost so much more in America, September 19, 2019. Financial Times. Available at: <https://www.ft.com/content/e92dbf94-d9a2-11e9-8f9b-77216ebelf17>

⁷See FDA website, New Evidence Linking Greater Generic Competition and Lower Generic Drug Prices. Available at: <https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/generic-competition-and-drug-prices>

⁸See FDA website, Statement from FDA Commissioner Scott Gottlieb, M.D., on new policy to improve access and foster price competition for drugs that face inadequate generic competition [Press release]. 19 February 2019. Available at: <https://www.fda.gov/news-events/press-announcements/statement-fda-commissioner-scott-gottlieb-md-new-policy-improve-access-and-foster-price-competition>

three companies in 2014 (in 2004 Barclays Global Investors, which was taken over by Blackrock in 2009, was number one or two in all three companies). A comparison between 2014 with 2004 also shows the growth of Vanguard, both in terms of the size of its shareholdings and position.

Table 4.1: Top five investors in top brand firms

Johnson & Johnson			
2004		2014	
State Street Global	5%	BlackRock	6%
Barclays Global Investors	4%	Vanguard Group	6%
Fidelity Investments	3%	State Street Global	5%
Robert Wood Johnson Foundation	2%	Royal Bank of Canada	2%
Vanguard Group	2%	Fidelity Investments	2%
Merck & Co			
2004		2014	
Barclays Global Investors	4%	BlackRock	6%
State Street Global	3%	Capital World Investors	6%
Fidelity Investments	3%	Wellington Management	5%
Vanguard Group	2%	Vanguard Group	5%
Capital Group	2%	State Street Global	4%
Pfizer			
2004		2014	
Fidelity Investments	4%	BlackRock	7%
Barclays Global Investors	4%	Vanguard Group	5%
State Street Global	3%	State Street Global	4%
Vanguard Group	2%	Capital World Investors	2%
Wellington Management	2%	Wellington Management	2%

Table 4.2 shows the top five investors in the three largest generic firms that operate in US markets in 2004 and 2014. Here too we see that BlackRock is an important common owner with shareholdings in Endo International and Perrigo in 2014. However, in comparison to the relatively stable ownership structure of brand companies in Table 4.1, we see more changes in the identity and size of the shareholdings of the top shareholders in generic firms. We further note that, especially in 2004, the largest investor in each company has a sizeable stake. For instance, Kelso & Company has a stake of 66% in Endo in 2004, and J.P. Morgan Chase has a stake of 27% in Sun Pharmaceutical in 2004. The largest shareholders in brand firms have much smaller stakes (around 5-6%). Furthermore, the identity of these top investors is different to the top investors in the largest brand firms, especially for Sun Pharmaceutical.

Table 4.2: Top five investors in top generic firms

Endo International			
2004		2014	
Kelso & Company	66%	Capital Group	10%
Black Diamond Capital	8%	Janus Capital Group	9%
Royce & Associates	6%	BlackRock	7%
Barclays Global Investors	4%	Vanguard Group	6%
Fidelity Investments	3%	Blue Ridge Capital	4%
Perrigo			
2004		2014	
Wellington Management	13%	BlackRock	7%
Royce & Associates	10%	Vanguard Group	6%
Jandernoa (Michael J)	9%	Fidelity Investments	5%
Barclays Global Investors	7%	State Street Global	4%
Perkins Investment	6%	Wellington Managment	3%
Sun Pharmaceutical			
2004		2014	
J.P. Morgan Chase	27%	Shanghvi (Dilip Shantilal)	11%
ABF Espana Gestion	14%	Viditi Investment	10%
Arisaig Partners (Asia)	14%	Tejaskiran Pharmachem Industries	9%
Aberdeen Asset Management	14%	Family Investment	9%
HDFC Asset Management	4%	Quality Investment	9%

4.2.3 Institutional investors' influence in pharma

Despite having shareholdings of “only” 5-7%, there is growing evidence that institutional investors such as BlackRock and Vanguard engage in active discussions with company management and boards with a view to influence companies' long-term strategies (McCahery et al., 2016). Specifically, in pharmaceutical markets, institutional investors with common holdings can be seen to take an active interest in the strategic decisions of companies. We now provide some anecdotal evidence of this.

In 2016, a group of representatives of major US institutional investors including Fidelity Investments, T. Rowe Price. and Wellington Management called a meeting with top biotech executives and pharma lobbyists to demand firm leaders do a better job defending their pricing.⁹ The meeting took place at a hotel conference room in Boston.

In 2019, BlackRock stated in their annual stewardship report that they engaged with a number of pharmaceutical companies including Abbott, Abbvie, Bristol-Myers Squibb, Pfizer, Novartis, Merck, GlaxoSmithKline, Johnson & Johnson, Sanofi,

⁹See Chen, C. Mutual Fund Industry to Drugmakers: Stand Up and Defend Yourself, Bloomberg News, 2016. Available at <https://www.bloomberg.com/news/articles/2016-05-09/top-funds-said-to-tell-pharma-leaders-to-defend-drug-pricing>.

Biogen, Allergan, Teva Pharmaceutical and Takeda.¹⁰

Similarly, State Street reported in their 2019 annual stewardship report that they engaged with 64 pharmaceutical companies.¹¹ The head of corporate governance at State Street Global Advisors stated that “Our size, experience, and long term outlook provide us with corporate access and allow us to establish and maintain an open and constructive dialogue with company management and boards.”¹²

More recently, in relation to the COVID-19 crisis, institutional investors have openly pushed for firms to collaborate with rivals and share information. In April 2020, a number of asset managers, including BlackRock and Fidelity, announced that “they want drug companies to put aside any qualms about collaborating with rivals.”¹³ BlackRock held talks with pharmaceutical companies to discuss ways to develop and deploy treatments by “working with industry competitors.” Separately, a group of 50 investors with over \$2.5 trillion in assets requested that companies share their findings related to the vaccine and agree not to enforce the relevant patents. Since then a number of alliances have formed to collaborate on treatments and vaccines for COVID-19.

Institutional investors have also been involved in merger decisions in the pharmaceutical industry. BlackRock is reported to have actively pushed for a merger between the pharmaceutical firms AstraZeneca and Pfizer. BlackRock, the largest institutional shareholder in AstraZeneca and also a top five shareholder in Pfizer at the time, “urged the British pharma giant’s board to eventually re-engage in talks with Pfizer Inc. over a possible deal.”¹⁴

¹⁰See BlackRock Investment Stewardship Annual Report 2019. Available at <https://www.blackrock.com/corporate/literature/publication/blk-annual-stewardship-report-2019.pdf>.

¹¹See State Street Stewardship Report 2019. Available at <https://www.ssga.com/library-content/products/esg/annual-asset-stewardship-report-2018-19.pdf>

¹²See Kumar, R. Passive investment, active ownership, State Street, 2014. Available at <https://www.ft.com/content/7c5f8d60-ba91-11e3-b391-00144feabdc0>.

¹³See Attracta Mooney, A. and Mancini, D. Drugmakers urged to collaborate on coronavirus vaccine, Financial Times, April, 2020, available at <https://www.ft.com/content/b452ceb9-765a-4c25-9876-fb73d736f92a>; Levine, M. Investors Want a Cure, Not a Winner, Bloomberg, April, 2020, available at <https://www.bloomberg.com/opinion/articles/2020-04-24/investors-want-a-cure-not-a-winner>

¹⁴See Plumridge, H., AstraZeneca Shareholder Backs Board Rejection of Pfizer Bid, Wall Street Journal, 2014, available at <https://www.wsj.com/articles/astrazeneca-shareholder-blackrock-sides-with-board-on-rejecting-pfizer-bid-1400791061>; Serafin, P. & Childs, M. BlackRock Is Said to Encourage Pfizer-AstraZeneca Talks, Bloomberg, 2014, available at <https://www.bloomberg.com/news/articles/2014-05-22/blackrock-is-said-to-encourage-pfizer-astrazeneca-talks>

4.3 Data and descriptive statistics

Our data comprises of publicly owned pharmaceutical firms (of any country of origin) that were active in the US pharmaceutical market between 2004 and 2014.¹⁵ Information on which firms are active in the US pharma market is obtained from the FDA Orange Book.¹⁶ We obtain the ownership structure of the companies in our sample from the Thomson Global Ownership Database. This database includes holdings of each shareholder in publicly listed firms for every year-quarter. For US-listed firms Thomson Reuters collects ownership information from 13F, 13D and 13G filings, and forms 3, 4, and 5. For companies listed outside the US, information is sourced from stock exchange filings, trade announcements, company websites, company annual reports and financial newspapers. For each firm, for each quarter, in the period 2004-2014 we extracted data on the shareholders that own at least 1% of the shares of the firm, and computed yearly ownership averages of each shareholder in each firm.

This dataset has considerable advantages over to Thomson’s Spectrum database used by most other papers on US common ownership. The Thomson’s Spectrum database is limited to 13F filings, which contains only large investors in US companies, whereas some pharma companies are not listed on a US stock market. Moreover, the Thomson’s Spectrum database shows holdings assigned to the owner that filed the 13F. This is what is commonly referred to as an “as-filed view.” Our database utilizes a “money-manager view.” With this view, the database combines together one or more filings to link the holdings to the actual firm that manages the investments. In other instances, it might break apart a single filing in order to accomplish the same. The holdings would then be assigned to one or more of the managers listed on the file.¹⁷

We use data from the FDA Orange Book to classify firms as “brand” or “generic” firms based on the type of drug that they have launched in the past. For each company and each year, we calculate the firm’s share of successful NDA applications (launched brands) relative to successful ANDA applications (launched generics). If a company operates subsidiaries, we aggregate drug counts at the parent com-

¹⁵This is the same database as used in the paper of Newham et al. (2018). Available at: <http://ssrn.com/abstract=3194394>. The data ends in 2014 due to the workload of dynamically assigning ultimate owners to subsidiaries; see also footnote 23 and references therein.

¹⁶The FDA Orange Book provides data on all launched pharmaceutical products in the United States. We drop conglomerates such as GE and Procter & Gamble from the sample as these firms focus on multiple markets and have launched relatively few pharmaceutical products given their large size. In total the sample consists of 157 distinct pharmaceutical firms.

¹⁷For a detailed explanation of our data and dynamic assignment of ultimate owners, see data repository: <https://www.openicpsr.org/openicpsr/project/120781/version/V1/view> attached to the paper Albert Banal-Estanol et al. (2020).

pany level. For each year we calculate the share of generic drugs out of all drugs launched by each company. Thereafter, we calculate the average generic share of each company during the years in which the company was active, within the time span 2004-2014. We categorize companies based on this measure. Firms with an average generic share of 90% or more are classified as “generic firms.” Remaining firms are classified as “brand firms.”¹⁸. Our dataset also contains information on the total market value of the firm.

Table 4.3 presents the ten largest common shareholders for our sample of brand firms at the start of our sample (2004) and at the end of our sample (2014). Firstly, it is clear that the largest common investor in 2004 is Barclays. Barclays has a stake of at least 1% in 48 brand companies in 2004. In our sample there are 85 brand firms in total in 2004, thus Barclays holds a stake of at least 1% in more than half of all brand firms in 2004. In 2009 BlackRock and Barclays merged which had an impact on BlackRock’s size. BlackRock moves from being number 6 in 2004 to being number 1 in 2014 with a stake in 68 brand companies in 2014.

¹⁸Our categorization aims to label “generic firms” as those firms that have limited R&D capabilities and focus almost entirely on producing generic drugs. A number of firms engage in the production of both brand and generic drugs, and may do so within the same company or may separate the activities in different subsidiaries. For example, while the company Novartis is primarily focused on developing brand drugs, its subsidiary Sandoz produces generic medications. Hybrid firms, such as Novartis, that have strong R&D capabilities and have an average generic share of less than 90%, are classified as brand firms in our analysis. Our data shows that these hybrid companies show very similar common ownership patterns to the pure brand companies, which is why we classify them together

Table 4.3: Top 10 common investors in brand firms

Investor	No. of shareholdings >1%	No. of shareholdings >5%	Average size of shareholding	No. of companies where investor is the largest
2004				
Barclays Global Investors	48	4	3%	1
Fidelity Investments	41	16	5%	11
State Street Global	41	2	2%	2
Vanguard Group	39	0	2%	0
Wellington Management	31	11	5%	6
BlackRock	25	0	2%	1
Capital Group	24	11	6%	5
Northern Trust Global	23	0	1%	0
HarbourVest Partners	23	3	4%	2
Deutsche Bank	21	0	2%	0
2014				
BlackRock	68	46	6%	14
Vanguard Group	65	32	4%	2
Fidelity Investments	53	26	5%	9
State Street Global	48	4	4%	1
Wellington Management	32	13	5%	3
Northern Trust Global	29	0	1%	0
Invesco	23	4	4%	1
T. Rowe Price	23	9	5%	3
Mellon Financial Corp.	21	0	1%	0
Royal Bank of Canada	19	1	2%	0

In our sample there are 86 brand firms in total in 2014, thus BlackRock holds a stake of at least 1% in 79% of all brand firms in 2014. In the firms where BlackRock has an ownership stake of at least 1%, the size of their stake is 2% on average in 2004, and increases to 6% in 2014. This is enough to place BlackRock as *the largest* shareholder in 14 companies in 2014 (up from 1 company in 2004). It is also evident that there has been very little change in the identity of the top five largest common owners for brand firms (apart from Barclays changing into Blackrock due to its merger). The top owners are BlackRock (Barclays), Fidelity Investments, State Street Global, Vanguard Group and Wellington Management.

Table 4.4 presents the ten largest common shareholders for our sample of generic firms at the start of our sample (2004) and at the end of our sample (2014). Comparing Table 4.4 to Table 4.3, we can see some clear differences in terms of the identity and size of the holdings of the common investors. Among the top common investors in 2004 is Franklin Templeton (with an average shareholding of 9%), UTI Asset Management (with an average shareholding of 24%) and HSBC Holdings (with an average shareholding of 12%). Thus, in 2004 we find that common investors in generic firms have large shareholdings in a selective set of firms. In 2014, these common investors with large stakes disappear or take a cut in the average size of

their shareholding. For example, in 2014, the average shareholding size of Franklin Templeton declines to 5%. Common investors have less coverage of generic firms in comparison to brand firms. In our sample there are a total of 29 generic firms in 2004, and 35 generic firms in 2014. Vanguard and BlackRock — the two largest common investors in generic firms in 2014 — have stakes in 11 generic pharmaceutical firms (31% of all generic companies).

Table 4.4: Top 10 common investors in generic firms

Investor	No. of shareholdings >1%	No. of shareholdings >5%	Average size of shareholding	No. of companies where investor is the largest
2004				
Franklin Templeton	14	4	9%	2
UTI Asset Management	12	9	24%	6
Fidelity Investments	8	1	3%	1
Vanguard Group	8	0	2%	0
HSBC Holdings	8	5	12%	1
Barclays Global Investors	8	2	5%	0
State Street Global	7	0	2%	0
Invesco	6	0	2%	0
Reliance Capital	6	3	8%	0
J.P. Morgan Chase	6	3	10%	1
2014				
Vanguard Group	11	4	4%	0
BlackRock	11	7	5%	2
Fidelity Investments	9	1	3%	0
State Street Global	7	0	3%	0
Dimensional Fund Advisors	7	0	2%	0
Life Insurance Corp. of India	6	2	4%	0
Franklin Templeton	6	1	5%	0
Norges Bank Investment	6	0	2%	0
HDFC Asset Management	5	1	3%	0
Capital Group	5	2	5%	2

4.4 Network analysis

In this section, we provide an analysis of the evolution of the common ownership links in the pharmaceutical industry. We make use of network analysis, which uses graph theory to describe the structure and characteristics of networks of actors by focusing on the links that exist between them. Graphs are made up of “nodes” which are connected by “edges” or “links”. In our setup, the nodes represent the firms whereas the edges represent the common ownership links that exist between pairs of firms.

We proceed in three steps. We first provide a graphical analysis of the common ownership links that exist within and between the top brand and top generic

firms. Subsequently, we investigate the determinants of such links by analysing the “investor networks” created by the top three investors in the industry. Finally, we analyse which brand and generic firms are the most influential (i.e. the most “central”) in the common ownership networks of the pharmaceutical industry.

4.4.1 Common ownership links between top firms

We first depict the evolution of the common ownership links amongst (i) the top 20 brand firms, (ii) the top 20 generic firms, and (iii) between the top 20 brand firms on the one hand and top 20 generic firms on the other hand. In all our graphs, the size of the nodes represent the value of the company, relative to the other companies in the same network, whereas the weight of the edges represent how strong the common ownership connections are. We make use of two common ownership measures, which determine links on the basis of (i) individual levels or (ii) joint levels of ownership.¹⁹

4.4.1.1 Common ownership networks among brand firms

Figure 4.1 provides a comparison of the network structure of the 20 most valuable (“top 20”) brand firms, which are also the 20 most valuable firms overall, at the beginning and end of the sample period, i.e. in 2004 and 2014 (Panels A and B, respectively). To ease the comparison, we depict the firms that were in the top 20 in both years in the same position (in green circles). We also include the top 20 companies in 2004 that drop from the top 20 by 2014 (in blue diamonds) and, vice-versa, those that appear in the top 20 in 2014 but were not in the top 20 in 2004 (in purple diamonds). As a measure of the common ownership link between two firms, we compute the number of individual investors whose ownership stake is larger than 5% in both firms, i.e. the number of common investors with more than 5% in both firms.

A link between two firms exists if they have at least one such common investor. The weight of the link between two firms depends on the number of such common investors that the two firms share.²⁰

Figure 4.1 shows that the top brand firms have become more connected over time, according to this measure of common ownership. As shown by Panel A, several pairs of firms already had common investors, i.e. with more than 5% in both firms, in 2004.

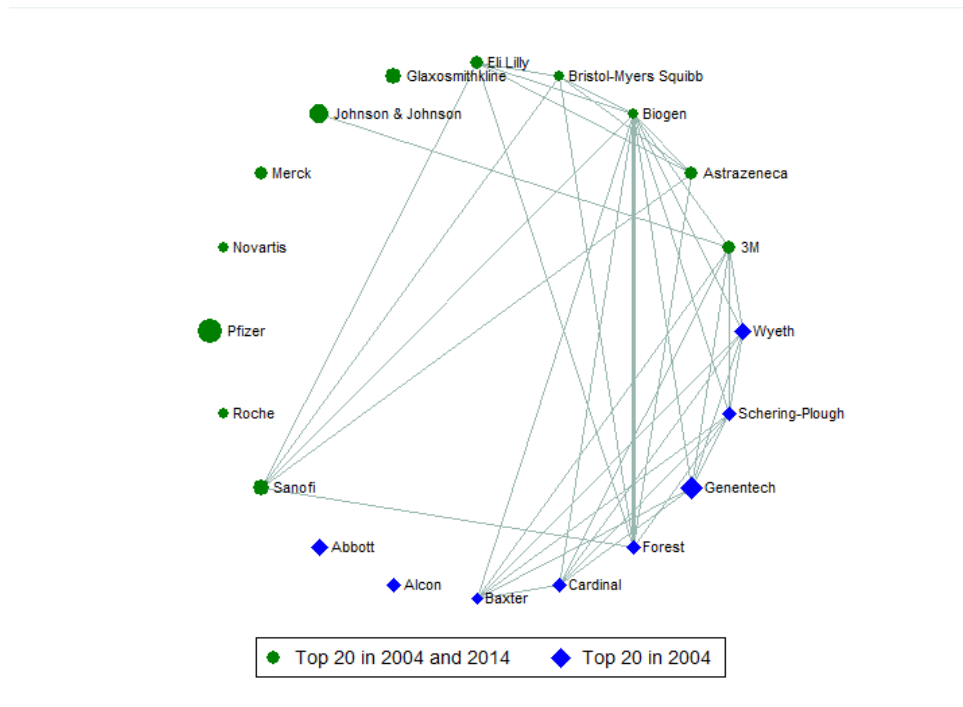
¹⁹All network plots are made using “nwcommands” See Thomas U. Grund, *nwcommands: Software Tools for the Statistical Modeling of Network Data in Stata* (2014). Available at: <http://nwcommands.org>.

²⁰In formal terms, and denoting by $s_{i,j}$ the ownership share of an investor i in firm j , the “weight” of the link between any pair of firms j and j' is given by $\sum_i I(\min(s_{i,j}, s_{i,j'}) > 0.05)$ where $I(x)$ is the indicator function that takes a value of 1 if the condition x is satisfied and a value of 0 if it is not. If the weight is 0 the link between the pair of firms does not exist.


But the links that existed in 2004 had relatively low weight, i.e. the firms have few investors in common. Some of the largest firms, such as Pfizer or GlaxoSmithKline, had no connections at all. There are, however, some firms that are highly connected. Perhaps surprisingly, the most connected firms, such as Baxter or Cardinal, or the ones with stronger links, such as Biogen and Forest, are relatively small within the set of top-20 companies. Similarly, if anything, most of the (smaller) firms that drop from the top 20 by 2014 are more connected than those that remain.²¹ In sum, the network in 2004 is not only sparser as compared to 2014, but also more asymmetric.

As shown by Panel B, the network becomes almost fully connected by 2014. Some firms, such as Pfizer, go from having no connection in 2004 to being almost fully connected with all the other firms in 2014. The connections between firms also become stronger. For example, in 2014 Johnson & Johnson and 3M have three common investors with more than 5% in both firms. Interestingly though, some firms, such as Sanofi, Novartis and Roche, remain without any links in 2014. Although to some extent present, the institutional investors in these firms have ownership stakes that do not reach the 5% threshold, in part because of the presence of large non-common investors such as L'Oreal in Sanofi, the Sandoz Family and the Novartis foundation in Novartis, and Novartis itself as a shareholder in Roche. In sum, top brand firms become, according to the individual measure of ownership, more connected over time with a few notable exceptions.

²¹The majority of companies that exit the top 20 in 2014 were acquired. Schering Plough was acquired by Merck in 2009. Genentech was acquired by Roche in 2009. Forest was acquired by Actavis (now Allergan) in 2014. Novartis acquired a majority stake in Alcon in 2010. Wyeth was acquired by Pfizer in 2009. Abbott, Cardinal and Baxter still exist as independent companies.



(a) 2004



(b) 2014


Figure 4.1: Common ownership network of the top 20 brand firms - Individual ownership

Notes: The size of the nodes indicates the value of the firm. The weight (thickness) of the edges represents the strength of the connections. A link between two firms exists if they have at least one common investor with more than 5% in both firms. The weight of the link between two firms depends on the number of such common investors that the two firms share.

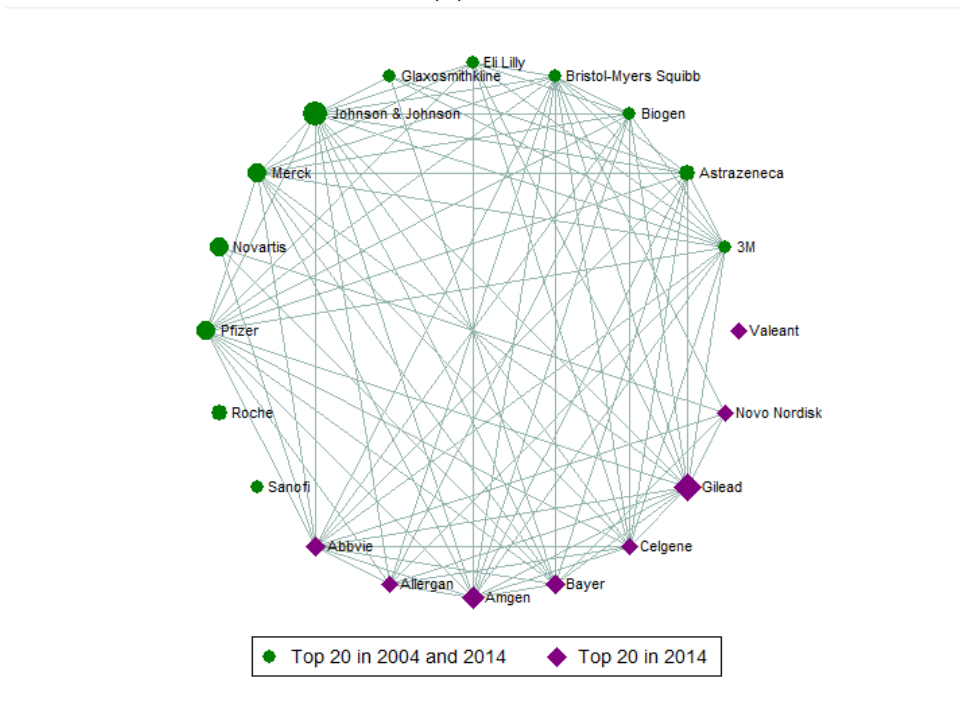
Figure 4.2 shows the network for a measure of joint ownership of the common investors. This measure compares the ownership stakes of all the common investors in relation to the ownership stakes of all the investors in our database. We consider two companies linked if the common investors ($>1\%$ in the two firms) own, on average, in the two firms, more shares than the non-common investors ($>1\%$ in just one of the two firms); that is, if the sum of the ownership stakes of all the common investors is greater than 50%. Note that there is no measure of the strength of the links in this network; the link just exists or not (in this sense, it is an example of an “unweighted network”). We make use of the same set of top 20 most valuable brand firms in 2004 and 2014, as in Figure 4.1.²²

Figure 4.2 shows that the common investors own more than half of all the (large) shareholders in many pairs of firms, both in 2004 (Panel A) and 2014 (Panel B). The network becomes even more connected over time. Novartis, for instance, had no connections in 2004. But, in 2014, the common investors of Novartis and Bayer, for instance, have more than 50% of the shares in both firms; the two firms become thus connected according to our joint measure. This is true despite the fact that Novartis and Bayer do not share any single individual investor holding more than 5% in both firms (as shown in Figure 4.1 Panel B). In general, though, we observe a less dramatic change when using the joint measure of common ownership than the individual one. Thus, the effects of the evolution of common ownership may depend on whether common investors have individual influence or if they do (or can) exert joint influence. For both measures, some firms, such as Sanofi and Roche, remain without any links in 2014.

²²In formal terms, and denoting by $s_{i,j}$ the ownership share of an investor i in firm j , a link between any pair of firms j and j' exists if $(\sum_{i \in C} (s_{i,j} + s_{i,j'})/2) > 0.5$ where C is the set of “common investors” in that pair of firms j and j' , i.e. those investors i with $\min(s_{i,j}, s_{i,j'}) > 0.01$.



(a) 2004



(b) 2014

Figure 4.2: Common ownership network of the top 20 brand firms - Joint ownership

Notes: The size of the nodes indicates the value of the firm. A link between two firms exists if the common investors (>1% in the two firms) own, on average, in the two firms, more shares than the non-common investors (>1% in just one of the two firms).


4.4.1.2 Common ownership networks among generic firms

Figure 4.3 replicates the network analysis of Figure 4.1 for the 20 most valuable generic firms. First, we again use the number of common investors whose ownership stake is larger than 5% in both firms as our measure of a common ownership link.


Figure 4.3 shows that the generic firms, contrary to the brand firms in Figure 4.1, became less connected in 2014 relative to 2004. Sun Pharma, for instance, lost all the connections it had in 2004, despite the fact that it became relatively larger. Overall, the level of connectivity of the generic firms is substantially lower than the brands in both years. Whereas the top brands are almost fully connected in 2014, the network of the generic firms is sparse. Very few firms have connections with other firms, and even fewer have connections with more than one investor. Only Perrigo, Impax and Endo have relatively strong links with each other.

Figure 4.4 shows the generic network of common ownership using the joint shareholding measure. We again take into account the ownership stakes of all the common investors in relation to the ownership stakes of all the investors in our database. We consider two companies linked if the sum of the ownership stakes of all the common investors in the two firms is, on average, greater than 50%.

Figure 4.4 shows that the generic firms became less connected in 2014 relative to 2004, when applying our joint measure of common ownership. This is also what we found when using the individual measure. Sun Pharma lost all of the connections that it had in 2004 despite the fact that it became relatively larger. Overall, the level of connectivity is even lower when using the joint measure as opposed to the individual measure of common ownership. Very few firms have connections in 2004, and even fewer have connections in 2014. The exception is Perrigo, which slightly increased its number of connections.




(a) 2004



(b) 2014

Figure 4.3: Common ownership network of the top 20 generic firms - Individual ownership

Notes: The size of the nodes indicates the value of the firm. The weight (thickness) of the edges represents the strength of the connections. A link between two firms exists if they have at least one common investor with more than 5% in both firms. The weight of the link between two firms depends on the number of such common investors that the two firms share.



(a) 2004



(b) 2014

Figure 4.4: Common ownership network of the top 20 generic firms - Joint ownership


Notes: The size of the nodes indicates the value of the firm. A link between two firms exists if the common investors (>1% in the two firms) own, on average, in the two firms, more shares than the non-common investors (>1% in just one of the two firms).

4.4.1.3 Common ownership networks between brand and generic firms


We now provide an analysis of the bipartite network of brands and generics. Bipartite networks are a particular class of networks, whose nodes are divided into two sets, and only connections between two nodes in different sets are allowed. As in the previous analysis, we use two measures of common ownership: (i) individual ownership, where the links reflect the number of investors whose ownership stake is larger than 5% in both firms, and (ii) joint ownership, where a link exists if the sum of the ownership stakes of all the common investors is greater than 50%. Note again that the size of the nodes represents the value of the firm relative to the firms in the same network.

Figure 4.5 shows that the brands and the generics became significantly more connected over time when looking at individual levels of ownership. As shown by Panel A, most brand-generic pairs were not connected in 2004, and in case they were, they only had one investor in common. Even the largest brands, such as Pfizer, had zero connections with the generics. Instead, as shown by Panel B, the number and the strength of the connections between brands and generics increased in 2014. Most of the large brands, such as Johnson & Johnson and Pfizer, have a large number of links. Some generics, such as Impax and Perrigo, have a high number of connections too.

Figure 4.6 shows that, when considering common ownership networks based on the joint measure, the same pattern emerges. Whereas in 2004 there were very few links between brand and generic companies, in 2014, these links were much more numerous (although fewer when compared to the common ownership network based on individual ownership).




(a) 2004




(b) 2014

Figure 4.5: Bipartite network of brands and generics - Individual ownership

Notes: The size of the nodes indicates the value of the firm. The weight (thickness) of the edges represents the strength of the connections. A link between two firms exists if they have at least one common investor with more than 5% in both firms. The weight of the link between two firms depends on the number of such common investors that the two firms share.



(a) 2004



(b) 2014

Figure 4.6: Bipartite network of brands and generics - Joint ownership

Notes: The size of the nodes indicates the value of the firm. A link between two firms exists if the common investors (>1% in the two firms) own, on average, in the two firms, more shares than the non-common investors (>1% in just one of the two firms).

4.4.2 Investor networks

This section investigates the determinants of the common ownership links identified in the previous section. We analyse in particular the evolution of the “investor networks” created by the shareholdings of the top three individual institutional investors of 2014 (Blackrock, Vanguard and Fidelity; see tables 4.3 and 4.4) in both brand and generic firms.

4.4.2.1 Brand firms’ investor network

Figure 4.7 represents the investor networks of Blackrock, Vanguard and Fidelity in the top 20 brand firms, in the beginning (2004) and end of our sample (2014), respectively. Each figure shows a “radar plot” of the ownership stakes. The axis tick marks represent the levels of 2.5%, 5%, 7.5% and 10%.²³


A comparison of Panel A with Panel B shows the significant growth of Blackrock and Vanguard over time. Blackrock’s growth is partly due to the merger with Barclays Global Investors in 2009. In 2004, Blackrock only had ownership stakes below 5%. In 2014 Blackrock owns significant stakes in many of the top pharmaceutical companies, usually in the range of 5-7.5%, but in some cases even close to 10%. In 2004, Vanguard’s stakes are all below 2.5%. In 2014, Vanguard’s ownership stakes are consistently around 5%.

Fidelity owns a much lower number of blocks than Vanguard and BlackRock, although they tend to be of a larger size in 2004. The holdings of Fidelity appear more stable over time and have not experienced the same growth as Vanguard and BlackRock, which have surpassed Fidelity in both number and average size of holdings.

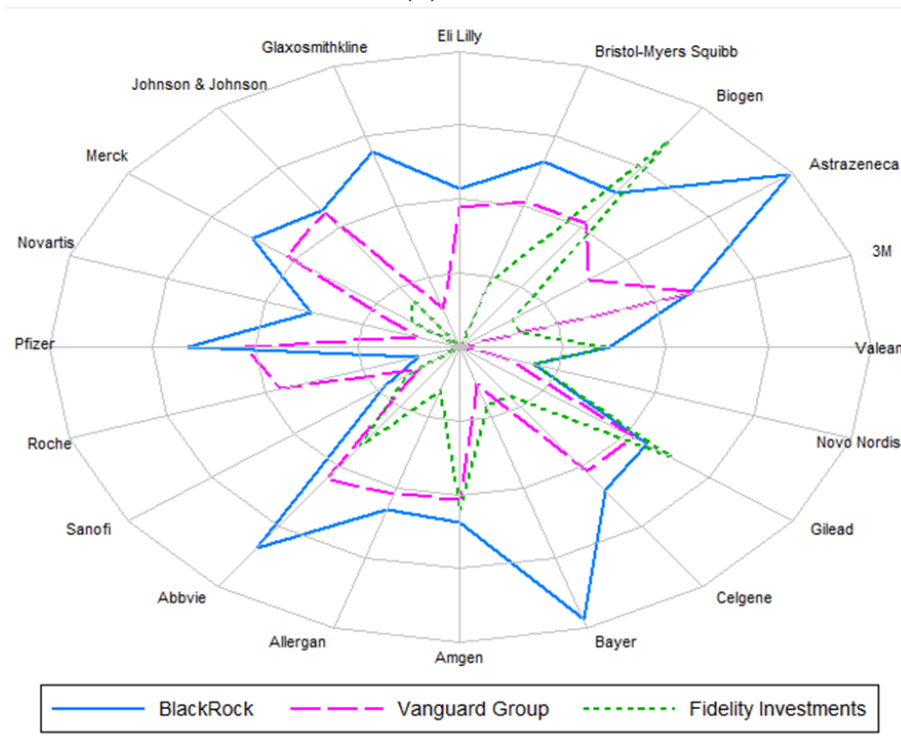
4.4.2.2 Generic firms’ investor network

Figure 4.8 represents the investor networks of the top three investors in the top 20 generic firms in 2004 and 2014, respectively. Comparing Figure 4.8 with Figure 4.7 shows much smaller investor networks in the generics than in the brands. While increasing over time, in 2014 Blackrock and Vanguard own significant stakes in just five of the top 20 generic firms. Fidelity owns even fewer and smaller blocks in 2014 than it did in 2004.

²³All radar plots are made using RADAR’: Adrian Mander, 2007. RADAR: Stata module to draw radar (spider) plots, Statistical Software Components S456829, Boston College Department of Economics, revised 02 Sep 2018.



(a) 2004




(b) 2014


Figure 4.7: Investor networks in brand firms

Notes: The axis tick marks for each circle represent shareholding levels of 2.5%, 5%, 7.5% and 10pct.

4.4. NETWORK ANALYSIS



(a) 2004



(b) 2014

Figure 4.8: Investor networks in generic firms

Notes: The axis tick marks for each circle represent shareholding levels of 2.5%, 5%, 7.5% and 10%.

4.4.3 Centrality in the brand and generic networks

We now analyse which are the most influential brand and generic firms in their respective common ownership networks. In network analysis, influence is measured by how “central” an individual node’s position is in the network, based on the existence and strength of its links with other nodes (Freeman, 1977; Freeman, 1978).

We proceed as follows. We first provide a definition of two standard measures of centrality: degree and closeness centrality. As both of them depend on network size, throughout this section, we consider the network of the 85 most valuable brands and the network of the 25 most valuable generics in each year. Although the identity of the firms in each network changes over time, the number (and thus the size of the network) remains constant throughout the sample period.

We then provide a description of the centrality measures of the top 20 brand and top 20 generic firms within each of their networks, both at the beginning and end of the sample. Finally, we show the evolution of the mean and dispersion of the measure of degree centrality for the whole set of firms in each of the two networks.

4.4.3.1 Definitions

We construct two centrality measures based on the concepts of degree and closeness centrality. These concepts capture slightly different aspects of the firms’ roles in the common ownership network. We provide a definition of each:

- **Degree centrality** measures the number of relationships an actor in the network has. The more ties, the more opportunities to interact and so the more influential, or central, the actor is. Firms that have ties to many other firms may be in an advantaged position. Since they have many ties, they are less dependent on any other firm for information, for instance. Formally, degree centrality counts the number of unique ties each firm has; that is, the number of unique firms with which a firm has a link with. Naturally, as networks become more connected, the average degree centrality across firms increases.
- **Closeness centrality** is a measure based on the distance between nodes.²⁴ Nodes with high closeness centrality are close to all other nodes, that is, they can reach all other nodes in only a few steps. In contrast, nodes with low closeness centrality are far away from all other nodes. For unconnected nodes, we set the distance to all the other nodes as the maximum distance in the network plus 10.

²⁴For a formal definition see Section 13.2 in Thomas U. Grund (2014). *nwcommands: Software Tools for the Statistical Modeling of Network Data in Stata*.

4.4.3.2 Firm-level and mean centrality measures

Table 4.5 shows the average degree and closeness centrality for each top 20 brand company (calculated within the sample of 85 firms in the brand network), as compared to the levels in 2004. The average value at the bottom of the table is the average amongst the top 20 brand companies. We order firms by their size (market value) in 2014. We again make use of the two measures of common ownership: individual and joint ownership.

As we can see in Table 4.5, the average degree centrality for both ownership measures has more than doubled between 2004 and 2014. There are differences in which firms are the most central depending on the measure used. If we focus on the individual ownership measure, in 2004, Biogen and Allergan have the highest levels of degree centrality with values of 25 and 24 respectively. These two firms also have high closeness values in 2004 when using the individual ownership measure. In 2014, the most central firm is Biogen with a value of 51 for degree centrality using the individual measure.

When using the joint ownership measure, in 2004, Pfizer and 3M have the highest levels of degree centrality with values of 23 and 21 respectively. In 2014, the most central firm is Johnson & Johnson with a value of 40 for degree centrality using the joint measure. This indicates that how common ownership links are measured plays an important role in determining which actors are the most central.

Table 4.5 also shows that many of the top 20 brand firms have a similar number of connections in 2014 based on the measure degree centrality for the individual measure i.e. levels of degree centrality lie between 46 and 51. Still, some others, including large firms such as Novartis and Roche, are not connected at all and thus have a level of degree centrality of 0, both in 2004 and in 2014. The levels of degree centrality in 2004 were substantially lower than in 2014: Johnson & Johnson had one connection whereas Pfizer had none. The differences across firms in terms of closeness centrality for both measures are lower than for degree centrality. In addition, the differences between 2004 and 2014 are smaller in the case of closeness than in the case of degree centrality.

Table 4.5: Levels of centrality for the 20 brand companies

Firm	Individual Ownership Measure				Joint Ownership Measure			
	Degree centrality		Closeness		Degree centrality		Closeness	
	2004	2014	2004	2014	2004	2014	2004	2014
Johnson & Johnson	1	46	0.12	0.22	16	40	0.11	0.15
Pfizer	0	46	0.07	0.22	23	31	0.11	0.15
Merck	0	49	0.07	0.22	19	30	0.11	0.15
Gilead	0	49	0.07	0.22	20	33	0.11	0.15
Novartis	0	0	0.07	0.07	0	3	0.07	0.14
Amgen		49		0.22		27		0.15
Roche	0	0	0.07	0.07	0	0	0.07	0.07
Astrazeneca	10	46	0.12	0.22	7	21	0.11	0.15
Biogen	25	51	0.13	0.22	14	25	0.11	0.15
Glaxosmithkline	0	45	0.07	0.22	1	8	0.10	0.14
Bayer	1	47	0.07	0.22	0	14	0.07	0.14
Abbvie		46		0.22		31		0.15
Bristol-Myers Squibb	10	48	0.12	0.22	15	33	0.11	0.15
3M	16	46	0.13	0.22	21	40	0.11	0.15
Sanofi	10	0	0.12	0.07	0	0	0.07	0.07
Eli Lilly	10	47	0.12	0.22	10	15	0.11	0.14
Celgene	15	46	0.13	0.22	3	33	0.10	0.15
Valeant	0	4	0.07	0.19	0	5	0.07	0.14
Novo Nordisk	11	8	0.13	0.19	0	0	0.07	0.07
Allergan	24	48	0.13	0.22	13	14	0.11	0.14
Average	7.39	36.05	0.10	0.19	9.00	20.15	0.10	0.14

Table 4.6 shows the degree and closeness centrality of the top 20 generic firms within the 25-generic firm network in 2014, as compared to the levels of 2004. We again order firms by 2014 market value, and include the averages at the bottom of the table.

The levels of degree centrality for the generics are substantially lower than for the brand firms. For both measures, many generics have a degree of zero in 2014, including the largest generic firm in our sample, Sun Pharma. The generic firm with most connections in 2014, Endo, has 7 when using the individual measure, i.e. 29% of the maximum number of connections possible in the generic network (24). By comparison, 15 out of the top 20 brand firms have more than 45 connections, i.e. 54% of the maximum number of connections possible in the brand network (84). Moreover, the average degree of centrality of generics is lower in 2014 than it was in 2004.

Table 4.6: Levels of centrality for the 20 generic companies

Firm	Individual Ownership Measure				Joint Ownership Measure			
	Degree centrality		Closeness		Degree centrality		Closeness	
	2004	2014	2004	2014	2004	2014	2004	2014
Sun Pharma	3	0	0.10	0.08	4	0	0.11	0.08
Perrigo	3	6	0.08	0.11	0	2	0.08	0.09
Endo	1	7	0.08	0.11	0	0	0.08	0.08
Lupin	6	0	0.10	0.08	5	0	0.11	0.08
Dr. Reddy's	0	0	0.07	0.08	0	0	0.08	0.08
Teva	0	1	0.07	0.11	0	0	0.08	0.08
Cipla	6	1	0.10	0.09	6	0	0.11	0.08
Hikma		0		0.08		0		0.08
Taro Pharma	1	0	0.08	0.08	0	0	0.08	0.08
Ranbaxy	6	1	0.10	0.09	6	0	0.11	0.08
Aurobindo	5	0	0.10	0.08	4	0	0.11	0.08
Akorn	0	6	0.07	0.11	0	0	0.08	0.08
Glenmark		0		0.08		0		0.08
Torrent	6	0	0.10	0.08	6	0	0.11	0.08
Haemonetics	2	6	0.08	0.11	0	1	0.08	0.09
Impax	0	6	0.07	0.11	0	1	0.08	0.09
Wockhardt	6	0	0.10	0.08	5	0	0.11	0.08
Zhejiang Huahai	0	0	0.07	0.08	0	0	0.08	0.08
Alembic		0		0.08		0		0.08
Ajanta Pharma		0		0.08		0		0.08
Average	2.81	1.70	0.09	0.09	2.25	0.20	0.09	0.08

In sum, the measures of centrality are substantially higher in 2014 as compared to 2004 for the brand firms. For the generic firms, the opposite is true. Degree centrality is not only much lower than for the brand firms in both years, but it is also lower in 2014 than it was in 2004. In the following subsection, we investigate more systematically the evolution, over time, of average degree centrality for both measures of ownership.


4.4.3.3 Evolution of the mean and dispersion of the centrality measures over time

We now investigate the evolution of the mean and dispersion of centrality over time in the brand and generic networks. We again make use of the two measures of common ownership: individual and joint ownership. For simplicity, we focus on one of the measures of centrality, degree centrality (the pattern is similar for closeness centrality).

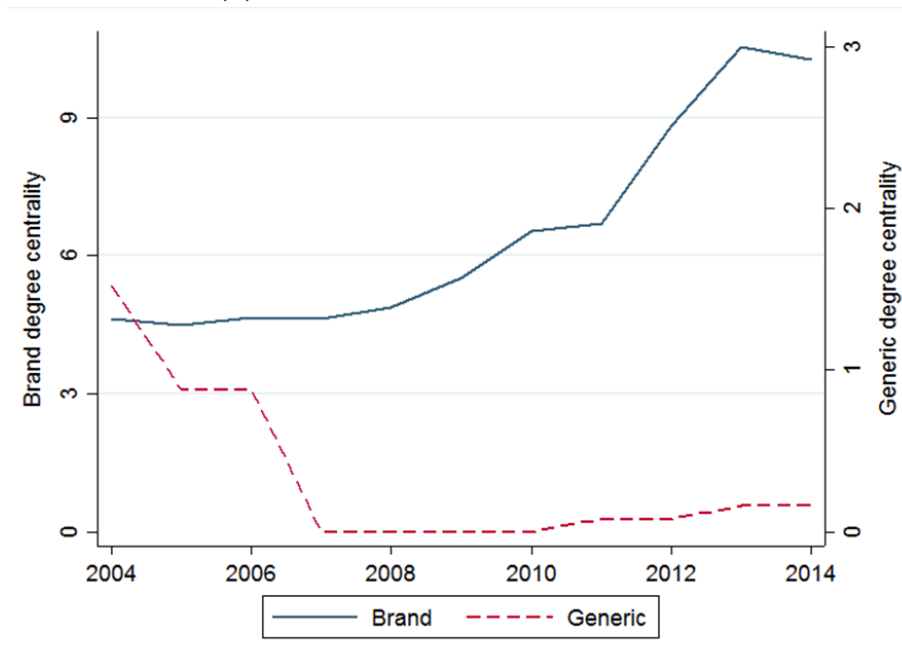
Figure 4.9 shows the average degree centrality for the 85 brand and the 25 generic companies over the 2004-2014 period (brand on the right axis, and generics on the left), for the two measures of common ownership: (individual in Panel A and joint

in Panel B).

Note that there are important differences between the two sets of firms. Whereas the average degree of the brand firms has increased substantially, the average degree centrality of the generic firms has decreased over time. This is true for the individual measure of ownership but it is especially the case for the joint measure of ownership.



(a) Individual ownership measure




(b) Joint ownership measure

Figure 4.9: Average degree centrality of brand and generic firms over time (2004-2014)


Figures 4.10 and 4.11 display the histograms of the measure of degree centrality in both 2004 and 2014 for the brand firms and generic firms respectively. Panel A shows the individual ownership measure and Panel B shows the joint ownership measure.

Figure 4.10 shows that the dispersion of degree centrality across the brand firms has increased in 2014 as compared to 2004, in both measures of ownership. A relatively large number of firms still have zero connections and thus a degree centrality of zero in 2014. But the highest levels of degree centrality become higher by 2014 relative to 2004.

As can be seen from Figure 4.11, the conclusions that can be drawn about the dispersion of the generic firm network are sensitive to the measure of common ownership used. If we use the individual common ownership measure, it appears that the dispersion of degree centrality across the generic firms has increased slightly in 2014 as compared to 2004. A larger number of firms have zero connections and thus a degree centrality of zero. The most connected generics have a slightly higher number of connections in 2014 as compared to 2004. However, if we use the joint measure, we find that centrality is limited to the range of 0-2 in 2014, whereas in 2004 some generics had centrality measures in the range of 4-6. These differences are explained by the fact that for generics, in 2004, there are common owners present with large stakes (see Table 4.4) which create linkages in 2004 when we use the joint ownership measure. With the absence of these investors in 2014, the common ownership network for generic firms on the basis of the joint ownership measure is much sparser (see Figure 4.4 Panel B).




(a) Individual ownership measure




(b) Joint ownership measure

Figure 4.10: Histogram of degree centrality for brands within the 85-brand firm network in 2004 and 2014



(a) Individual ownership measure



(b) Joint ownership measure

Figure 4.11: Histogram of degree centrality for generics within the 25-generic firm network in 2004 and 2014

4.5 Antitrust implications

Our empirical analysis shows, generally, that the common ownership network among brand companies has become denser and more complete over time, whereas that of

the generics is much sparser and becomes sparser over time. Finally, the bipartite network between brand companies, on the one hand, and generic companies, on the other, has become denser. This section discusses the antitrust implications of these patterns.²⁵

We first discuss the implications of the dynamics of the brand firm network for innovation, as brand companies engage in innovation investments with the aim to patent new drugs –and enjoy rents from the resulting temporary monopoly. We then analyze the consequences of the evolution of the bipartite network between brand and generic companies on generic entry in markets where the brand no longer enjoys regulatory protection. Thirdly, we discuss the implication of common ownership for prices in the pharmaceutical industry. High drug prices are a major concern for policy makers in the US: prescription drugs, responsible for 10% of all healthcare costs, represent one of the fastest-growing areas of healthcare spending (Martin et al., 2017). Finally, we briefly discuss implications for collusion.

4.5.1 Innovation

R&D is crucial for bringing new drugs to the market. Thus, whether common ownership positively or negatively affects innovation in the pharmaceutical industry is a key concern for policy makers. Common ownership between brand companies may, on the one hand, enhance information sharing, generate synergies, and increase the incentives to invest in R&D. On the other hand, common ownership may also incentivize firms to innovate in a way that avoids head-on competition between each other in the innovation space. We briefly discuss each of these possibilities in turn.

The increasingly dense common ownership network that we observe among brand companies may be good for innovation for the following reasons. First, the common ownership links may *facilitate information sharing* between connected firms. This can bring in substantial benefits. Indeed, in the early stages of development, firms select which R&D projects to bring into their R&D portfolio and test numerous lead molecules. At this stage, connections with other firms may provide an opportunity for brand firms to share each other’s knowledge bases. Indeed, Kostovetsky and Manconi (2020) find a higher intensity of patent citations among firms that share institutional owners, suggesting that common institutional investors can facilitate the diffusion of information among their portfolio firms. In a similar vein, Ghosh and Morita (2017) show that cross-ownership, which has elements in common with common ownership (see also below in the section on pricing), can induce knowledge transfer between connected firms, thereby increasing consumer surplus and/or total

²⁵For a full discussion of legal theories to tackle common ownership see Elhaug (2020).

surplus under certain conditions.

Common ownership links may also lead to more informal or formal innovation collaborations, with the associated benefits. Indeed, sharing scientific personnel and/or research labs that result in a combination of complementary assets may lead to synergies. Similarly, collaboration may lead to the reduction of wasteful innovation duplication. He and Huang, for instance, find evidence suggesting that institutional cross-ownership facilitates explicit forms of collaboration, such as within-industry joint ventures and strategic alliances, and that this improves innovation productivity (He and Huang, 2017).²⁶ Geng et al. (2017) find furthermore that shareholder ownership overlap across firms with patent complementarities correlates significantly with higher investment in innovation and more success with patents.

Further, common ownership links between innovating pharma companies can increase innovation by mitigating technology spillover problems. Indeed, companies often hold back on costly innovation efforts since competitors may be able to imitate and free ride on these efforts. If companies are commonly owned, then innovation spills over to companies within the same network and may thus benefit the same owners. Supporting this line of reasoning, Lopez and Vives (2019) show theoretically that horizontal common-ownership links can mitigate firms' well-known disincentives to innovate that can arise because of the technological spillovers. Anton et al. (2018) confirm theoretically and empirically that common ownership may incentivize firms to engage in more R&D. In particular, common ownership increases R&D when technological spillovers are large relative to product market spillovers. If the reverse occurs, i.e., when product market spillovers are larger, then common ownership reduces R&D.

Common ownership may also *reduce competition in innovation*. For example, common ownership might negatively affect the number and/or the selection of R&D projects pursued. As drugs pass through clinical trials, firms may re-optimize their portfolio and decide which drugs to submit for FDA approval. Many development projects are terminated, not due to safety or efficacy concerns, but due to commercial considerations. Large pharmaceutical firms often invest in 10-15 distinct research programs that run simultaneously. In an effort to reduce competition, firms with common investors may *jointly* pursue a similar line of research or *terminate competing projects*. This is potentially to the detriment of consumers if it means that fewer drug variants are available.

Recent research indicates that one of the motives for pharmaceutical firms to engage in M&As is to neutralize potential competition. The idea is that an incumbent

²⁶Although there is evidence that research joint ventures, in turn, may facilitate collusion in product markets see Duso et al. (2014) and Helland and Sovinsky (2019).

-i.e., a company that has already launched a drug- has an incentive to acquire and terminate projects in the development process if these projects have “overlap” with its launched product (where overlap is defined as the same mechanism of action within a therapeutic class). These acquisitions, where the incumbent acquires a nascent or potential competitor in order to neutralize the competition have been termed “killer acquisitions.” Cunningham et al. (2019) find that projects acquired by firms that have an overlapping drug are 23.4% less likely to have continued development activity.

The presence of common ownership between two firms with overlapping drugs may mitigate the need for a merger to achieve a similar effect. A recent paper that looks at common ownership links in pharmaceutical start-ups by venture capital (VC) companies finds precisely this effect (Li et al., 2020). In particular, the paper examines how a start-up responds after seeing a competitor make progress on a related drug project. If the two start-ups share a common VC, the lagging start-up is less likely to advance its own project, which reduces competition between the start-ups. The authors find that these anticompetitive effects are mostly present for technologically similar projects, early-stage projects, and with VCs involved that have larger equity stakes and less-diversified portfolios.

In sum, high common ownership among brand companies can have both positive and negative effects on innovation in the pharmaceutical sector. Current theoretical and empirical research highlights both sides. Research in this dimension is a promising avenue for future research, especially in terms of identifying whether and under which circumstances common ownership of firms with projects that have overlapping mechanisms of action and similar therapeutic classes leads to better or worse innovation outcomes.

4.5.2 Entry

Patented markets are the main source of revenue for brand companies. When the patent expires - or when it is challenged in court²⁷ - and generic companies enter, revenues for the brand decline dramatically (by as much as 90%). Therefore, brand companies have a strong incentive to deter generic entry, or at least to delay generic entry as long as possible. Entry induces losses to the brands and gains to the generics that are highly asymmetric: a brand company loses much more after entry than a generic profits after entry. Therefore, the joint payoff for brand and generic in holding off entry is clearly positive.

²⁷A brand’s patent validity can be brought to court through a Paragraph IV challenge, which is the section of the Hatch-Waxman act under which generic entrants dispute pharmaceutical patents e.g see Helland and Seabury (2016).

Scott Morton (2002) reviews how direct ownership links between brand and generic firms influences the likelihood of generic entry. She finds that generics owned by the original innovator (i.e., the brand company) are less likely to enter the market. This hints that an investor with shares in both the brand and generic may benefit from steering the generic away from entering. Therefore, entry decisions of generics may crucially depend on the joint ownership of generic and brand firms. Shareholdings in the brand provide common investors with incentives to steer decisions towards joint profits and shareholdings in the generic provide investors with the ability to influence such decisions.²⁸

Newham et al. (2018) find that this is indeed the case. They analyze generic firms' entry decisions into pharmaceutical markets opened up by the end of regulatory protection. They find that a higher level of common ownership between a brand firm and a potential generic entrant is robustly linked with a lower probability of generic entry, and that this effect is economically significant in the sense that overall common ownership at the market level decreases the total number of generics in that market.²⁹ This means that the increasingly dense bipartite network between brand companies, identified in the previous section, is likely to lead to less generic entry.³⁰

4.5.3 Pricing

4.5.3.1 Unilateral effects

Commonly owned brand firms that commercialize drugs that are therapeutically similar might have less incentives to unilaterally compete³¹ due to various mechanisms.³² Indeed, as O'Brien and Salop (2000) note, the anticompetitive effects of common ownership are similar to that of cross ownership in that common ownership can be understood to be ownership in one firm, coupled with cross ownership in the others

First, firms that are largely owned by shareholders who also have sizeable stakes in competitors might just simply act in these shareholders' interest, which leads them

²⁸Also, see Posner et al. (2017).

²⁹Related, Xie and Gerakos (2017) find that common ownership between brand and generic is positively associated with the two parties entering into a settlement agreement where the generic manufacturer stays out of the market. Hovenkamp and Lemus (2017) further confirm that settlements after Paragraph IV challenges cause generics to stay out of the market.

³⁰There is also evidence from other industries that ownership structures affect entry. Majumdar (2017) documents the relationship between horizontal ownership and entry in the local telecommunication exchange segment in the US, and finds that dominant ownership controllers experienced lower entry in their territories.

³¹Previous research in the airline and banking industries has pointed towards a positive relationship between common ownership and prices see Azar et al.(2019) and Azar et al. (2018).

³²For an overview of the mechanisms by which large horizontal shareholdings are likely to influence corporate management see Elhauge (2021).

- rather than maximizing their own profits - to maximize the return of their shareholders' portfolios, in whose interest it might be to soften price competition (Azar, 2017). Further, while there is evidence that institutional investors engage in active discussions with companies' management (see McCahery et al., 2016), investors do not need to actively intervene to have an impact on the firms' decisions. They may apply "selective omission" by encouraging actions that increase both firm value and portfolio profits and remaining silent when this is not the case (Hemphill and Kahan, 2020). Further, they may design payment schemes for the top management to shape their incentives in a way that leads to softer product market competition. Anton et al. (2020) find that higher firm-level common ownership is linked to less performance-sensitive incentives for CEOs and other top managers, which in turn may lead to softer competition.

Increases in common ownership links between brand and generic companies, as we show in the bipartite network in the previous section, may also indirectly raise drug prices. Indeed, common ownership should reduce generic entry and, as shown by previous research, the reduction of generic companies in the market increases prices. In sum, both the increasingly linked brand network and brand-generic bipartite graph suggest that price competition might have softened. This is an interesting area for future research: to study how the link between common ownership affects prices through the channel of entry; see e.g., Grabowski and Vernon (1992) for a study that links generic entry to drug prices and Suzuki (2013) for a study that looks at the impact of differences in market conditions (regulation) on prices through the channel of entry.

4.5.3.2 Coordinated effects

Our empirical results indicate that common networks among generics are sparse and, if anything, have become sparser over time. While managers of commonly owned firms may unilaterally engage in anti-competitive behavior, common ownership might also induce coordinated action. Economic theory predicts that communication can facilitate both coordination and monitoring defection from a common strategy. While many forms of private communication are illegal, public information disclosure could serve as an alternative coordinating and monitoring mechanism to achieve tacit collusion, as suggested by e.g. OECD (2012). Indeed, Pawliczek et al. (2019) find that higher horizontal shareholding levels increase firm disclosures of information that can help firms to coordinate.

Rock and Rubinfeld (2020) provide a summary of how common ownership has an impact on coordinated effects; we provide some elements of that discussion here.

A key issue is how ownership structure can affect the likelihood that a coordinated outcome will be achieved, i.e., the relevant question is how common shareholders can have an influence in coordinating outcomes. The article discusses a variety of ways in which a common owner will be more conducive to collusion, by being, for example, a better “cartel ringmaster” or “cartel initiator.” On the other hand, there are also a variety of ways in which a common owner can be a poorer cartel organizer than a non-common owner.

Among generics, where we find sparse networks of common ownership, a large cartel operating between the years of 2006 and 2016 is currently being investigated.³³ It may thus be that, in the pharmaceutical industry, common ownership and explicit collusion are substitutes. However, we should be very careful when making this connection: the generic pharmaceutical industry has a number of other characteristics that make cartels more likely, for example homogenous products, and frequent interaction at industry trade fairs.

4.6 Conclusion

This paper documents the common ownership networks between companies that operate in US pharma markets during the period 2004-2014. We show that common ownership networks between brand companies are rather dense and complete, especially at the end of our sample. Furthermore, the common ownership links between brand and generic companies have become notably stronger.

While there is little direct evidence yet how these common ownership networks might impact competition and innovation in pharmaceutical markets (with the notable exception of the impact on generic entry), the presence of large institutional investors in the industry is so wide-spread that it would be hard to believe that they have no material impact. The further investigation of their influence in pharma markets is an exciting topic for future research.

³³Rowland, C. Investigation of generic ‘cartel’ expands to 300 drugs, *The Washington Post*, December 10, 2018.

Chapter 5

The Interaction between Industry Payments to Physicians, Insurance and Drug Costs: Evidence from Medicare Part D

Chapter Abstract

High and growing prescription drug costs in the United States are a major concern for policy makers. This chapter focuses on the extent to which promotional gifts and other transfers made to physicians by pharmaceutical companies causes physicians to prescribe more expensive medicines. In the analysis data from a federal database on the universe of industry payments between 2014 and 2017 is linked to prescribing behavior in Medicare Part D. We develop a novel empirical strategy that uses data on the prescription behavior of physicians in Vermont, where a strict ban on industry payments to physicians is in place, combined with machine learning techniques to construct the counterfactual outcome for physicians who receive payments in the nearby states of New Hampshire and Maine. We find that a gift ban, such as the one implemented in Vermont in 2009, has the potential to result in a 3% decline in the total cost to treat diabetes. We investigate heterogeneity in the treatment effect and find that physicians who have a high share of patients with a low-income subsidy, and thus lower out-of-pocket expenditures, prescribe relatively more brand drugs and expensive drugs in response to industry payments. Our findings illustrate how industry payments interact with insurance to drive up health care costs.

5.1 Introduction

Health care cost containment is an important policy goal in the US. In 2018, US health care spending grew 4.6 percent, reaching USD 3.6 trillion.¹ One of the fastest growing areas of health care spending is prescription drugs (Martin et al., 2016). Prescription drug costs per patient for diabetes, the market of interest in this study, have doubled from 2013 to 2017. While consumers and healthcare payors pay for drugs, physicians control access. Consequently, pharmaceutical companies invest heavily in marketing drugs to physicians with the aim to influence physicians' prescribing decisions and increase drug sales.² In the United States, the promotion of drugs to physicians, includes not only the provision of information related to the product, but also entails free meals, gifts, fees for consulting and speaker events, travel and lodging, and payments for education and training. If such transfers result in the prescription of costlier medications, banning payments to physicians may be a helpful means to reduce health expenditure. This paper quantifies the extent to which payments to physicians leads to the prescription of costlier drugs and tests whether physicians who have patients with lower out-of-pocket expenditures, prescribe relatively more expensive medications in response to industry payments.

We make use of a publicly available federal database on industry payments to physicians. With the enactment of the Physician Payments Sunshine Act, healthcare product manufacturers are required to fully disclose payments more than USD 10 to physicians. In September 2014, the first batch of payment data was made available to the public via the online "Open Payments" portal. The data indicates that payments to physicians are substantial: In 2017, industry payments totaling USD 8.31 billion were made to more than 600,000 physicians and 1000 teaching hospitals (openpaymentsdata.cms.gov).

Physicians act as (imperfect) agents for their patients. When a physician is deciding between two substitutable drugs of differing prices, for example an expensive brand drug and a cheaper generic alternative, a personal reward in the form of a gift or an expensive dinner if he/she prescribes the brand may shift the physician's preference towards the costlier medication. If physicians, in addition to caring about their patients' clinical health, care about their patients' financial health, a trade-off arises between the cost of the medication to the patient and the physician's expected reward. If, on the other hand, the patient is fully insured against pre-

¹CMS National Health Expenditure Data. Available online: <https://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/NationalHealthExpendData/NationalHealthAccountsHistorical>

²Even accounting for direct-to-consumer advertising 90% of promotional expenditures are directed towards physicians (Donohue et al., 2007).

scription drug costs, this trade-off does not exist. Thus one might expect that as patients' out-of-pocket expenditures decline, physicians are more likely to prescribe costlier medication as a consequence of industry payments. We test this hypothesis.

In our analysis we combine four years of data (2014-2017) from Open Payments matched with data on prescribing behavior from Medicare Part D at the physician level. We focus on prescriptions and payments related to treatments for diabetes. The data allows us to observe the rate at which physicians prescribe brand vs. generic anti-diabetics and the average cost of a dose (30-day supply) at the physician level. Variation in out-of-pocket expenditures is driven by the share of a physician's patients that have a low-income subsidy. In contrast to regular Medicare beneficiaries, beneficiaries with a low-income subsidy (LIS) face substantially lower out-of-pocket expenses for prescription drugs (Yala et al., 2014).

There are a number of challenges that arise when aiming to identify the causal effect of industry payments on prescribing decisions and the manner in which payments interact with patients' out-of-pocket expenditures. Firstly, payments to physicians are not random. Pharmaceutical firms typically monitor physicians' prescriptions and specifically target high-volume or high-cost prescribers for payments (Fugh-Berman and Ahari, 2007). This strategic behavior by firms tends to make the correlation between a prescriber's payments and prescriptions positive even if payments have no effect on prescribers' choices (Carey et al., 2020). Finally, patients with low out-of-pocket expenditures as a result of the low-income-subsidy may have specific characteristics, such as more severe medical conditions, which could cause higher treatment costs independently of insurance coverage or payments to physicians. Finally, industry payments may have spillover effects to physicians who do not receive payments through peer networks (Agha and Zeltzer, 2019).

To combat these challenges, this paper proposes a novel identification strategy. We make use of data on the prescription behavior of physicians in Vermont, where a strict ban on industry payments to physicians is in place, combined with supervised machine learning techniques (specifically, regularized regression) to construct the counterfactual outcome for physicians who receive payments in the nearby states of New Hampshire and Maine. This approach allows us to control for the manner in which prescribing costs change with key explanatory variables in the complete absence of industry payments. To determine the average treatment effect we compare predicted prescribing behavior with actual prescribing behavior at the physician level, and test how this effect varies with observable characteristics, including the share a physician's patients' that have a low-income subsidy.

Identification of the causal treatment effect relies on the assumption of 'unconfoundedness' i.e. that we can control for all relevant confounders that affect both the

treatment and outcome variable (Rubin, 1990). The application of machine learning is crucial in this regard as it allows us to flexibly control for a very large number of covariates. Identification also relies on the standard assumptions of ‘common support’ and ‘no spillovers.’ For these reasons we train the machine learning algorithm using data from physicians in Vermont where industry payments are banned.

We find that physicians who receive payments prescribe a significantly higher share of brand drugs and have higher average drug costs per dose. Receipt of an industry payment, on average, causes the brand prescription rate of a physician to increase by 5 percentage points in the year in which the payment is received, corresponding to a 20% increase in the unconditional average brand prescription rate. The drug cost for a 30-day dose is USD 21 higher for physicians who receive industry payments, corresponding to a 22% increase in the unconditional average drug cost. The effect of a payment increases significantly with the share of patients who have a low-income subsidy (and thus out-of-pocket costs). The average increase in the brand prescription rate for physicians who have a high share of patients with a low-income subsidy (>75%) is 9 percentage points and the average increase in the drug cost per dose is 29 USD. This suggests that greater insurance coverage makes it easier for physicians to appease pharmaceutical companies.

During the time period of the sample (2014-2017) a total of just under USD 900,000 was spent by the pharmaceutical industry on payments related to anti-diabetics to 245 physicians in the states of New Hampshire and Maine. Back-of-the-envelope calculations indicate that holding prices and prescribing volumes constant, the elimination of payments would have resulted in savings of just under USD 10 million during this period as a result of substitution towards cheaper medication. This represents a 3% decline in the total prescription costs for diabetes in New Hampshire and Maine. The increase in drug costs is more than 10 times what the industry spent on payments during this time period which clearly indicates that payments to physicians are very profitable for the pharmaceutical industry as a whole.

To the extent that payments may affect the quantity of a specific drug that physicians prescribe, this is a conservative estimate of the savings to healthcare payors from the elimination of industry payments to physicians. This financial cost needs to be traded-off against the possible benefits of industry payments to physicians, for example, as industry claims, payments may be associated with information that helps physicians to make better prescription decisions.

This paper is structured as follows: Section 5.2 outlines the contribution of this paper to the literature; Section 5.3 explains the data sources, dataset construction and undertakes a descriptive analysis of the data; Section 5.4 lays out our empir-

ical strategy; Section 5.5 contains the estimation results; Section 5.6 provides a discussion of the results; and Section 5.7 concludes.

5.2 Contribution to the literature

This paper contributes to the literature on insurance coverage and drug expenditure, and the literature on the impact of industry payments on physicians' prescription choices. While a number of papers have sought to quantify the effect of payments on physicians' prescribing behavior, to the best of our knowledge this paper is the first to study heterogeneity in the effect of industry payments to physicians with respect to a measure of patients' out-of-pocket costs. To identify the causal effect of industry payments we make use of machine learning techniques. Thus, this paper also contributes to the emerging body of work that applies machine learning to aid causal inference.

Several papers document a positive relationship between expansion in insurance coverage and increasing drug utilization and expenditures (see for example Ghosh et al., 2019; Duggan and Scott Morton, 2006; Danzon & Pauly, 2002). This is in line with standard theories of insurance and moral hazard (Pauly, 1968; Zweifel and Breyer, 1997). Higher insurance may reduce an individual's motive to prevent illness, leading to greater medical expenditures. More importantly, once an individual has fallen ill, individuals with insurance face lower effective prices, which might lead them - or their physicians - to elect for greater levels of treatment and/or more expensive treatment than they would have received if they were uninsured.

Relatedly, previous studies have investigated the extent to which physicians take patients' out-of-pocket expenditure and price sensitivity into consideration when prescribing drugs. Using data on physician prescription decisions in Sweden during the years 1992 and 1993, Lundin (2000) finds that physicians are more likely to prescribe brand-name versions of a drug (where a generic is available) to patients with lower out-of-pocket expenditures. Carrera et al. (2018) find that physicians take the price sensitivity of patients into consideration in their prescription decisions; this is evidenced by a greater increase in the prescriptions of cholesterol-lowering statins to low-income patients following the release of generic statins.

Physicians' prescription choices can also be influenced by rent-seeking motives. In Taiwan, where physicians prescribe and dispense drugs, Liu et al. (2009) find that profit incentives affect physicians' prescribing decisions. Iizuka (2012) shows that Japanese doctors respond to drug prices when choosing between brand-name drugs and generic drugs (only) if they expect to pocket some share of drug expenditure. Making use of a controlled field experiment in China, Lu (2014) finds that

when doctors are provided with an incentive to promote drug sales, prescriptions for insured patients cost more than those for uninsured patients.

Using US data, several papers study the effect of gifts and payments from pharmaceutical firms to physicians on prescription patterns. For example in a cross-sectional study using industry payment data from the Open Payments Program from 2013, DeJong et al. (2016) find that receipt of an industry-sponsored meals is correlated with an increased rate of prescribing the brand name medication being promoted. Physicians who receive payments from industry are also found to have higher average prescribing costs per patient and a larger share of brand drug prescriptions (Perlis & Perlis, 2016; Qian et al., 2017). The problem with studies of this nature is that they do not show evidence of a *causal* relationship between payments and prescription choices. Fugh-Berman and Ahari (2007) discuss how drug firms commonly monitor physicians' prescribing and specifically target high-volume prescribers for payments. This strategic behavior by pharmaceutical firms tends to make the correlation between a prescriber's payments and prescriptions positive even if payments have no effect on prescribers' choices.

One approach to correct for selection in the literature has been to use instrumental variables. This approach relies crucially on the validity of the instrument. There is no clear "good" instrument for payments to physicians and a number of different instruments have been used in the literature. Fernandez & Zejcirovic (2018) and Engelberg et al. (2014) use the distance between the county where the physician is located and the pharmaceutical company headquarters as an instrument for industry payments. This instrument is limited by the fact that it is only relevant for counties and physicians that that are within a reasonable distance from the headquarters, and company headquarters are few.³ Grennan et al. (2018) use regional variation in the strictness of conflict of interest policies at academic medical centers as an instrument for industry payments. The instrument is only relevant for physicians that are within a reasonable distance from an academic medical center, and given that the strictness of academic medical center policies may directly affect physicians prescription behavior, the instrument is only exogenous conditional on a rich set of controls. Payments received by a physician's colleagues has also be used as an instrument for industry payments (Aramal-Gracia, 2020; Agha and Zeltzer, 2019). A concern with this instrument is that changes in colleagues' prescription behavior, as a result of industry payments, may have a direct impact on the focal physician's prescription behavior making it plausible that the assumption of strict exogeneity is violated (Agha and Zeltzer, 2019).

³The mean distance to the closest headquarters of a physician is about 800km (Fernandez & Zejcirovic, 2018).

Another approach has been to include physician fixed effects to take out persistent unobserved differences across physicians (as in Carey et al., 2020 and Datta and Dave, 2016). The limitation with this approach is that in cases where the physician-firm relationship pre-dates the payments data, the estimates using a fixed effects strategy only measure the incremental effect of an additional payment. Since physician-firm relationships typically involve repeated interaction, the fixed effect approach is not able to capture the full effect of the relationship. A further identification strategy that has been used in the literature is matching.⁴

Our identification strategy differs from the existing literature on industry payments. We make use of data from the state of Vermont to construct the counterfactual outcome for physicians who receive industry payments. Machine learning is applied in order to specify a rich model that uses the data on the prescribing behavior of (unpaid) physicians in Vermont to predict the outcomes for paid physicians in neighboring states. The use of machine-learning to aid causal inference is a growing field in economics (see e.g. McCaffrey et al., 2004; Hansen & Kozbur 2014; Wager & Athey, 2017; Chernozhukov et al., 2017; Athey & Imbens, 2019; Abrell et al., 2019). Abrell et al. (2019) use machine learning to analyze the emissions and cost impacts of the UK Carbon tax. Our approach bears similarities with the method developed by Abrell et al. (2019) in that we use ML to predict the outcome in the absence of treatment. However, whereas Abrell et al. (2019) use machine-learning to construct the counterfactual in a situation where there is no control group, in our case, we data from the state of Vermont, where industry payments are banned, serves as the control. Abrell et al. (2019) predict the outcome for a counterfactual value of the treatment using an ML-trained model in which one can control the treatment variable. Our strategy, on the other hand, uses ML to predict prescribing choices in the complete absence of the treatment variable. The treatment effect is estimated as the difference between the actual outcome with treatment, and the predicted outcome without treatment.

5.3 Data

This section introduces the two main datasets used in the empirical analysis, Medicare Part D and Open Payments, and provides a brief overview of the anti-diabetics market. Thereafter we explain the construction of the final dataset and present some descriptive statistics.

⁴For example, see the robustness check in Agha and Zeltzer (2019). In the methodology section of this paper we discuss the advantages of our approach over matching (see Section 5.4.4).

5.3.1 Medicare Part D

Medicare is a US public health insurance program for the elderly (over the age of 65) and disabled that covers over 50 million beneficiaries (Cubanski et al., 2018). Medicare Part D, introduced in 2006, is a voluntary program that provides insurance for prescription drugs. A large share of Medicare beneficiaries are enrolled in Part D. In 2018, 43.4 million beneficiaries (74% of all Medicare beneficiaries) were enrolled in Part D (Cubanski et al., 2018).

Utilization of drugs in the Part D program is a function of physicians' prescribing decisions. Data on the prescriptions dispensed under the Medicare Part D Program is available online from The Centers for Medicare and Medicaid Services (CMS).⁵ The Medicare Part D Detailed Prescriber Public Use File (PUF) provides data on prescriptions at the physician-drug-year level. A further database, the Medicare Part D Prescriber Summary PUF contains additional information at the physician-year level. In both datasets, physicians can be identified by their National Provider Identifier (NPI) and so the two datasets can be easily combined.⁶ In the Detailed Prescriber file each drug is identified by its brand name and generic name.

The Detailed Prescriber file provides information on the volume and total cost of physician's prescriptions of a specific drug in a given year. Measures of prescription volume include the total number of unique beneficiaries with at least one claim for the drug, the total number of claims and the total number of standardized 30-day fills. The total drug cost consists of "the ingredient cost of the medication, dispensing fees, sales tax, and any applicable administration fees and is based on the amount paid by the Part D plan, Medicare beneficiary, government subsidies, and any other third-party payers" (Centers for Medicare and Medicaid Services, 2019). To protect the privacy of Medicare beneficiaries, any aggregated records which are derived from 10 or fewer claims are excluded.

The out-of-pocket costs Medicare beneficiaries are the drug copays which are stipulated by their plan. For patients covered by insurance, typical drug copay costs for diabetes range from \$10 to \$50, depending on the drug. If the patient has a multi-drug regimen, copays are estimated to total around \$200 a month or more.⁷ In particular, out-of-pocket costs for insulin can be very high. Among insulin users without Part D low-income subsidies (LIS), average annual per capita out-of-pocket spending on insulin was \$580 in 2017.⁸

⁵<https://data.cms.gov/part-d-prescriber/archived-data>

⁶A NPI is a unique 10-digit identification number issued to health care providers in the US by the Centers for Medicare and Medicaid Services. An NPI is permanently associated with a specific individual regardless of any changes in practice location or additional speciality training.

⁷This estimate is obtained from <https://health.costhelper.com/diabetes-medication.html>.

⁸This figure is taken from <https://www.kff.org/medicare/issue-brief/insulin-costs-and->

Enrollees with sufficiently low income and assets receive a low-income-subsidy (LIS). In 2018, more than 12 million Part D enrollees (29%) received low-income subsidies (Cubanski et al., 2018). Patients with a LIS benefit from zero or much lower out-of-pocket expenditures for premiums, deductibles and medication costs (Yala et al., 2014).⁹ The share of beneficiaries with a LIS at the physician-year level is provided in the Prescriber Summary file. This variable provides important variation in patients' out-of-pocket costs across physicians.

Additional variables in the Prescriber Summary file include the prescriber's full name, address (street address, 9-digit zip code and state), speciality and gender. It also provides aggregate prescription measures for all drugs, the number of claims for beneficiaries covered by Medicare Advantage (MAPD) plans¹⁰, the average age of beneficiaries, the number of female and male beneficiaries and the beneficiary average Hierarchical Condition Category (HCC) risk score. The HCC risk score is determined by CMS using demographic information and diagnoses on Medicare fee-for-service claims to measure each enrollee's medical risk status, with higher scores going to enrollees with more (or more severe) health conditions or demographic risk factors. Thus, risk scores provide a proxy for patients' health status.

5.3.2 Diabetes

We focus on prescriptions for anti-diabetic drugs. This market is of particular interest because diabetes is becoming increasingly prevalent in the US, as well as globally, and costs to treat Medicare beneficiaries with diabetes have risen steadily overtime (see Figure 5.1). To the extent that payments to physicians may lead to higher drug costs, the potential savings in this market could be large. Moreover, since diabetes disproportionately affects older people, the sample of prescriptions to Medicare beneficiaries is likely to cover a substantial portion of prescriptions in the US diabetes market.

Diabetes is a chronic disease in which the body's ability to produce or respond to the hormone insulin is impaired, resulting in abnormal metabolism of carbohydrates and elevated levels of glucose in the blood. Diabetes is a prevalent public health concern in the US as well as worldwide. As of 2015, 30.3 million Americans - 9.4% of the US population - had diabetes (Centers for Disease Control and Prevention, 2017). Diabetes is on the rise; among people with Medicare, one third (33%) had

coverage-in-medicare-part-d/


⁹Yala et al. (2014) find that the average out-of-pocket costs for LIS beneficiaries are 74% lower than of the out-of-pocket costs for non-LIS beneficiaries with gap coverage (\$148 vs. \$570).

¹⁰Medicare-eligible individuals can acquire coverage for prescription drugs either through standalone Part D plans or bundled with medical and hospital coverage in the form of "Medicare Advantage" plans.

diabetes in 2016, up from 18% in 2000 (Cubanski et al., 2019). The risk of diabetes increases with age and of the estimated 30.3m Americans with diabetes in 2015, just under half are 65 or above.

Anti-diabetic medication has seen substantial price rises in recent years, which is a major focus for policy-makers. In particular, high insulin prices have attracted attention from the policy makers and the media.¹¹ Based on the full Medicare Part D dataset for the US, Figure 5.1 shows how both total Medicare expenditures on anti-diabetic treatments and average treatment costs per beneficiary have increased over time. Total Medicare expenditures on both brand and generic treatments amounted to USD 17.6 billion in 2017. Brand name drugs are responsible for the lion’s share of costs making up for 95% of total drug costs, despite only accounting for 30% of 30-day supply prescriptions for anti-diabetics in 2017.

Figure 5.1: Costs to treat diabetes



Current therapy for diabetes is almost entirely drug-based in combination with lifestyle intervention (weight-loss, appropriate diet, exercise). Type 1 diabetes requires insulin therapy. In the case of type 2 diabetes, which accounts for 90-95% of all diabetes cases, the first line treatment is Metformin (which is available as a generic) and thereafter no single drug is highlighted according to the Standards of Medical Care in Diabetes (Centers for Disease Control and Prevention, 2017; American Diabetes Association 2015).¹² Thus, manufacturers are motivated to influence physician prescription decisions through providing monetary incentives.

¹¹For example see Prasad (2019) for BBC News “The human cost of insulin in America.”

¹²Appendix Section 5.8.2 provides a detailed overview of the available treatments for diabetes and lists all the treatments in the final dataset.

5.3.3 Payments to physicians

Data on industry payments to physicians related to anti-diabetic treatments is obtained from the CMS Open Payments website. Open Payments is a national disclosure program created by the Affordable Care Act (ACA) and managed by the CMS that collects and publishes information about financial relationships between the healthcare industry and providers. The creation of Open Payments follows the passing of Section 6002 of the ACA of 2010, also known as Physician Payments Sunshine Act. The Act aims to promote transparency and accountability by providing information on gifts, payments and other transfers of value on the publicly accessible Open Payments website.

Each year, during a submission period, manufacturers fully-disclose payments of \$10 or more in value made to physicians and teaching hospitals. These transfers are categorized into three payment types: 1) general payments (food and beverage, speaker fees, consulting, travel and lodging, and gifts and entertainment); 2) payments for research purposes and 3) physician ownership and investment interests. After submission, there is a 45-day period during which payment recipients can review and dispute errors before public release. Thereafter the data are corrected and then uploaded online. The published data include the identities of the payment recipients and the paying firms, date of payment, associated product, payment amount, and nature of payment. The first batch of data became publicly available in September 2014. The 2013 program year includes only data collected from the second half of year, whereas all subsequent program year publications contain data from the full year.

The Sunshine Act came on top of an existing regulatory environment to prevent undue influence of the medical industry at the state level. While Open Payments provides the first national disclosure program, prior to its implementation six states had enacted similar disclosure programmes (Gorlach and Pham-Kanter, 2013).¹³ A handful of states also have laws that ban certain types of industry gifts and payments or place restrictions on the value of such transfers.¹⁴ One of the most comprehensive statutory gift bans was implemented by Vermont in 2009. Vermont's gift ban (18 V.S.A §4631a) prohibits most gifts, including free meals, to physicians who regularly

¹³These states include: Maine (2006), West Virginia (2004), Minnesota (1994), Massachusetts (2009), Vermont (2002), and District of Columbia (2004).

¹⁴States statutory gift bans payments prior to the Sunshine Act include Vermont (2009), Colorado (2007), Minnesota (1994) and Massachusetts (2009) (Gorlach and Pham-Kanter, 2013). Maine enacted a ban of certain types of pharmaceutical payments at the end of our sample, in mid-2017. However clear rules only came into effect in June 2020. In our data we do not see a decline in the value or number of payments to physicians in Maine in 2017 vs. 2016 and 2015. (See: <https://www.policymed.com/2020/06/maine-finalizes-its-physician-gift-ban-rules-to-exclude-pharmacists-speaker-fees-expenses-accredited-education-and-market-research.html>)

practice in Vermont. The law is actively enforced and violators have been made to pay penalties in the past.¹⁵ While certain types of payments are still “allowable” under the ban, the law severely limits payments to physicians.

The empirical analysis will make use of the payment ban in Vermont to construct a counterfactual of how physicians who receive payments would have prescribed in the absence of payments. For identification, we require that the physicians operate in states that are broadly similar to Vermont, aside from not banning payments, thus we focus on identifying the effect of payments on the prescription choices of physicians in the states of New Hampshire and Maine.

5.3.4 Dataset construction

The final dataset comprises of physicians in the states of Vermont, New Hampshire or Maine, who prescribe anti-diabetic medication to Part D beneficiaries, matched with payments data for the years 2014 to 2017. The construction of the dataset follows four main steps. First, the set of all anti-diabetic drugs is established. Next, we isolate prescription data from Medicare Part D on physicians that prescribe anti-diabetic drugs in the states of Vermont, New Hampshire or Maine. In a third step, Part D prescriptions are matched with the payments data at the physician-drug-year level. Finally, the dataset is aggregated to the physician-year level. This subsection describes each of these steps in detail and provides information on the final sample.

The set of all approved treatments for diabetes (both brand and generic) is identified using the FDA Orange Book matched with Anatomical Therapeutic Chemical (ATC) codes.¹⁶ We select drugs with the ATC code “A10 - Drugs used in diabetes”. The complete list of anti-diabetic treatments is matched with the drug names in Part D and Open Payments using string-matching algorithms, and a cross-walk file is created with standardized drug names.

Information on the prescriptions of anti-diabetic drugs, including the name and address of the prescribing physician, is extracted from the Medicare Part D database. This sample provides the universe of anti-diabetic medications prescribed to patients enrolled in Medicare Part D during this time span.¹⁷ The sample is restricted to physicians located in Vermont, New Hampshire or Maine.

Part D prescriptions are matched with the Open Payments data contained in the

¹⁵In 2013, Novartis was reported to have paid \$36,000.00 in civil penalties to settle a total of six gift-ban violations; each violation consisted of providing a meal to a health care provider (Hams & Wilkinson, 2013).

¹⁶The FDA Orange Book provides data on all launched pharmaceutical products in the US since 1982.

¹⁷Note that physicians may prescribe anti-diabetic medication to other patients that are not enrolled in Part D. This information is not publicly available and is not included in the dataset.

general payments file at the physician-drug-year level for the years 2014 to 2017.¹⁸ The dataset is aggregated to the physician-drug-year level such that payment values reflect the sum of all payments associated with a specific brand drug in a given year. The data is matched with Part D on the basis of drug name and physician name. Since there is no common physician ID that connects Part D and Open payments¹⁹, the datasets are matched on the basis of full name and 9-digit zip code. This is complemented by a manual check in cases where physicians in Part D did not directly match to the Open Payments database.

On the basis of this intermediate dataset, where information on payments is visible at the drug-level, a few stylized facts concerning payments to physicians can be established. Firstly, all payments are made in connection with a brand drug. There are no payments related to the promotion of a generic drug. A list of pharmaceutical companies and associated brand drugs with positive transfers to physicians in the sample is presented in Table 5.8 in the Appendix. Secondly, payments to physicians involve repeat interaction. During the years 2014 to 2017, for the subset of physicians receiving payments, the average number of payments that relate to the same brand drug is 9 and the average number of payments from the same company is 12. Thirdly, the average value of payments tends to increase with the total cost of the drug. As illustrated by Table 5.1, drugs in the lowest cost tercile have significantly lower average payment values than those in the top cost tercile (USD 295 vs. USD 1,119), although there is substantial variability in the size of payments across all cost levels.

Table 5.1: Payment value (USD) by drug cost

	mean	sd	observations
Lowest cost tercile of diabetes drugs	295.57	2,732.04	509
Middle cost tercile of diabetes drugs	499.45	2,718.09	459
Top cost tercile of diabetes drugs	1,199.40	5,030.51	421
Total	636.89	3,601.67	1,389

Lastly, the dataset at the physician-drug-year is aggregated to the physician-year level. The dataset is supplemented with information on yearly median household income and population density at the 5-digit zip code level from the United States Census Bureau’s American Community Survey, and information on diagnosed diabetes rates and obesity at the county and year level from the Centers for Disease

¹⁸Each payment in Open Payments is linked to a specific drug, in cases where multiple drugs are listed the payment value is split equally amongst all listed drugs.

¹⁹The Part D data uses as its unique ID each provider’s NPI number. The Open Payments system uses a randomly generated unique ID. By law, the government could not release NPI numbers with the Open Payments data, but it could, and did, release their name and address.

Control and Prevention.²⁰ A list of all relevant variables in the dataset, a description and their source is provided in the Appendix. The final sample is restricted to physicians. Nurses and physician assistants are dropped due to the fact that they never receive payments from pharmaceutical companies. Observations with missing information for any variable are dropped. Since the information provided in Part D is limited for cases where a drug is prescribed to 10 or fewer unique beneficiaries, physicians who prescribe anti-diabetic medication to a total of 10 or fewer beneficiaries per year end up being dropped. These physicians account for only 8% of all claims in the data, thus the final dataset covers the majority of prescriptions by physicians for diabetes treatments. The final sample consists of 5,704 observations and 1,862 physicians.

5.3.5 Descriptive analysis

In this section we present summary statistics and descriptive analysis using the final sample of physicians. Approximately 17% of physicians in New Hampshire and Maine receive a payment at some point in time during the time span 2014 to 2017. Not all physicians receive a payment every year. In total, 543 physician-year observations are associated with a positive payment value. Table 5.2 presents summary statistics for the sample split by physicians in New Hampshire and Maine who receive a payment in a given year (column 1), physicians in New Hampshire and Maine who do not receive any payments in a given year (column 2) and physicians in Vermont who do not receive any payments.²¹ On average paid physicians prescribe brand drugs for 33% of 30-day supply prescriptions. This is in contrast to unpaid physicians who prescribe brand drugs in 23% of cases. The average for physicians in Vermont, where many payments are banned, lies between these value at 27%. A comparison of means also indicates that physicians who receive payments are more likely to be male than those who do not.

Table 5.3 provides summary statistics on the payment variables. On average, a paid physician receives a total of 14 separate payments per year in connection with anti-diabetic treatments. The majority of these transfers are in-kind (such as free meals, travel and accommodation). The average total value of payments received for anti-diabetic treatments in a year is USD 1,610. There is significant variance in the size of yearly payments; with the maximum earnings for an observation in the sample totalling USD 98,443 per year. Cash payments tend to be much higher in value than in-kind transfers.

²⁰ Available online: <https://gis.cdc.gov/grasp/diabetes/DiabetesAtlas.html>

²¹ Despite the ban, a small fraction (4%) of physicians in Vermont received a positive payment at some point in the time span 2014 to 2017. These observations are excluded from the analysis.


Table 5.2: Summary statistics for physician characteristics (2014-2017)

	Received Payment		No Payment		Vermont	
	mean	sd	mean	sd	mean	sd
Brand prescription rate, anti-diabetics	0.33	0.16	0.23	0.17	0.27	0.14
Total 30-day supply, anti-diabetics	871.32	976.07	501.78	364.36	514.34	354.07
Drug cost '000 USD, anti-diabetics	155.23	351.64	53.43	85.55	62.95	94.36
Specialist in diabetes (0/1)	0.37	0.48	0.33	0.47	0.35	0.48
Family practice (0/1)	0.62	0.49	0.62	0.49	0.61	0.49
New practitioner (0/1)	0.10	0.30	0.15	0.35	0.11	0.31
Male practitioner (0/1)	0.84	0.37	0.59	0.49	0.62	0.49
Beneficiary count for anti-diabetics	78.93	137.81	40.37	39.62	38.34	39.87
Anti-diabetics claim share	0.09	0.11	0.06	0.07	0.06	0.06
Insulin claim share of anti-diabetics	0.24	0.15	0.21	0.17	0.27	0.15
Beneficiaries over 65 for anti-diabetics	0.15	0.24	0.12	0.26	0.12	0.27
Share of male beneficiaries	0.45	0.08	0.41	0.11	0.43	0.11
Total beneficiary count	372.57	176.78	328.63	141.29	321.66	124.92
MAPD claim share	0.21	0.14	0.20	0.14	0.09	0.07
LIS claim share	0.52	0.21	0.46	0.22	0.46	0.17
Average age of beneficiaries	70.19	4.18	70.78	4.25	71.29	3.34
Average risk score of beneficiaries	1.18	0.21	1.19	0.26	1.13	0.21
Median household income '000 USD	53.50	17.20	57.59	19.13	54.81	12.69
Population density '000	0.76	1.43	0.93	1.52	0.64	1.36
Percent diagnosed with diabetes	8.48	1.14	8.11	1.05	6.88	0.89
No. diagnosed with diabetes '000	12.69	8.22	13.32	8.24	4.19	2.12
Percent obese	29.38	3.62	28.53	3.32	25.09	4.07
No. obese '000	37.20	24.42	39.62	24.81	13.31	7.17
No. other physicians in same county	74.15	45.49	81.18	46.62	30.48	18.89
No. other physicians in same ZIP	9.71	8.34	12.69	10.72	7.23	6.07
Observations	543		4,162		999	

Table 5.3: Summary statistics for payments related to anti-diabetic drugs

	mean	sd	min	max	observations
No. payments per year	13.91	29.94	1	340	543
No. cash payments per year	1.07	6.06	0	59	543
No. in-kind payments per year	12.84	25.77	0	281	543
Value (USD) payments per year	1,610.28	8,471.34	1	98,443	543
Value (USD) cash payments per year	1,291.62	7,057.26	0	77,956	543
Value (USD) in-kind payments per year	318.67	1,531.51	0	20,487	543

The two key outcome variables of interest are the physician’s brand prescription rate which measures the share of 30-day prescriptions attributable to brand drugs, and the physician’s average cost for a 30-day prescription which is calculated by dividing total cost by total 30-day supply. Figure 5.2 compares how average brand prescription rates (a) and drug cost per 30-day supply (b) vary across physicians in Vermont and paid physicians for different levels of LIS claim share. Firstly, it is evident that for both a low (≤ 0.5) and a high (>0.5) share of LIS claims, paid physicians prescribe more brand drugs and more expensive drugs. Secondly, the difference is greater for physicians with a high LIS share than it is for physicians with a low LIS share. This provides an indication that there may be an interaction between patient’s out-pocket costs and the effect of industry payments.



(a) Brand prescription rate

(b) Drug cost per 30-day supply

Figure 5.2: Comparison of physicians’ prescription outcomes

5.4 Empirical strategy

In this section we explain the empirical methodology and identification strategy. First, we draw on the Rubin causal framework to characterize the causal ‘target’ or ‘treatment’ parameter of interest (Rubin, 1974). The target parameter is distinct from ‘nuisance’ parameters which are not of interest per se, but need to be modelled to identify the treatment effect. We then explain how we apply machine learning (ML) tools to model the nuisance parameters and thereby uncover the causal effect of payments to physicians.

5.4.1 Causal model

We model the prescription outcome Y_{it} of a physician i in period t as being generated according to

$$Y_{it} = f(X_{it}, H_{it}) + \theta D_{it} + U_{it}$$

where $D_{it} \in \{0, 1\}$ is a binary treatment variable equal to 1 if physician i receives an industry payment in period t and zero otherwise, X_{it} is a vector of observable characteristics, referred to as covariates, and H_{it} is a vector of unobserved covariates. U_{it} is a random error term with mean zero, $E[U_{it}] = 0$. U_{it} is independent of the covariates and treatment:

$$U_{it} \perp (X_{it}, H_{it}, D_{it}) \forall i, t$$

The covariates affect the outcome variable via the function $f(X_{it}, H_{it})$. The function f is invariant to the treatment. Given $E[U_{it}|X_{it}, H_{it}, D_{it}] = 0$, θ has the interpretation of the treatment effect parameter. For purposes of illustration θ is assumed to be constant for all physicians and across time. In our application we will model θ as a function of the observable characteristics i.e. $\theta = g(X_{it})$ where X_{it} will include the LIS claim share amongst other variables.

Let Y_{1it} denote the outcome for physician i with treatment ($D_{it} = 1$) and Y_{0it} denote the outcome without treatment ($D_{it} = 0$). The effect of an industry payment on the outcome Y_{it} can then be calculated as:

$$\theta = Y_{1it} - Y_{0it}$$

The fundamental problem of causal inference (Holland, 1986), often also referred to as the missing data problem (Rubin, 1974) is that we can never observe both outcomes for an individual at the same time. Here, we observe the prescription outcomes for physicians that have received payments from pharmaceutical companies but we do not observe these same physicians' prescription outcomes in the absence of payments. We propose to overcome the missing data problem by predicting the outcome without treatment for an individual using supervised machine learning. Given the possibility to predict $f(X_{it}, H_{it})$ we can re-write the treatment effect as:

$$\theta = Y_{1it} - f(X_{it}, H_{it}) - U_{it}$$

In order to calculate counterfactual outcomes, the following two assumptions concerning the interaction between the covariates and treatment have to be satisfied:

ASSUMPTION 1: Observed covariates are independent of changes in the treatment variable: $X_{it} \perp D_{it}$.

ASSUMPTION 2: Unobserved covariates are conditionally independent of the treatment variable given the observed controls: $H_{it} \perp D_{it} | X_{it}$.

Assumption 1 rules out an effect of the treatment variable on observed covariates. If D influences X, there would be an indirect effect on the outcome, which would bias our estimate of the treatment effect. Assumption 2 rules out effects of the treatment variable on unobserved variables after controlling for the observed variables. Again, if D would influence H, there would be an indirect effect on the outcome. It is important to note that assumption 2 does not rule out an effect of unobserved controls. It only implies that once we include all observed covariates into the model, the impact of unobserved variables is independent of the treatment (Abrell et al., 2019). It is worth highlighting that for assumption 2 to hold we need to be able to include all relevant confounders in X which affect both the treatment and outcome variable. The application of ML is instrumental in this regard as it allows us to flexibly control for a large number of covariates.

5.4.2 Using machine learning for prediction

To construct the counterfactual outcome we need an estimator \hat{f} for the function f that produces reliable *out-of-sample* predictions. To do so, we will make use of modern prediction methods and the fact that we observe outcomes in a state where payments are banned. Further details on the specific ML algorithms that will be used are provided in the following subsection.

ML algorithms seek to maximize predictive performance by minimizing the *test* mean squared prediction error (MSPE) which is defined as the average squared prediction error among observations not previously seen (Varian, 2014). Superior predictive performance is typically achieved by a combination of flexibility and simplicity, often described as the ‘bias-variance trade off’. Standard econometric techniques such as ordinary least squares (OLS) aim at minimizing the bias while allowing for high variance which can lead to ‘over-fitting’: The model represents the sample data very well, however it is unlikely to explain out-of-sample observations equally well. ML methods solve a bias-variance trade-off by introducing hyper- or tuning parameters in the estimation function. Given the optimal set of tuning parameters λ^* that minimizes the test MSPE, the true outcome for an untreated physician can be written as the sum of the predicted value and the prediction error $\xi(X_{it}, H_{it})$:

$$Y_{0it} = \hat{f}^{\lambda^*}(X_{it}) + \underbrace{f(X_{it}, H_{it}) - \hat{f}^{\lambda^*}(X_{it})}_{=:\xi(X_{it}, H_{it})} + U_{it}$$

To estimate \hat{f}^{λ^*} data is required on how the outcome variable relates to the observable covariates in the absence of treatment. The data we have comprises of the state of Vermont, that has a ban on all payments to physicians and other states that permit payments. For simplicity, let us just consider two states: state A that allows payments and state B that bans all payments. If we make the assumption that the function $f(X_{it}, H_{it})$ is invariant across the two states, we can write the corresponding population models for the state A and state B follows:

$$Y_{it}^A = \gamma_t^A + f(X_{it}^A, H_{it}^A) + \theta D_{it} + U_{it}^A$$

$$Y_{it}^B = f(X_{it}^B, H_{it}^B) + U_{it}^B$$

where γ_t^A is a state-year constant for the state without a ban to control for state-specific time-invariant characteristics and state-year-specific shocks.

ASSUMPTION 3: The function $f(X_{it}, H_{it})$ is invariant across the state that allows payments to physicians and the state that bans payments to physicians.

With this set-up, it is clear that we can use data from the state with a payment ban to estimate \hat{f}^{λ^*} in absence of D . An alternative option would be to use (or include) data on untreated physicians in state A to estimate the prediction function. However, untreated physicians that share a peer network with treated physicians may experience spillover effects (Agha and Zeltzer, 2019). If untreated physicians in state A are indirectly affected by payments we cannot cleanly predict how Y_{it} relates to X_{it} in the complete absence of industry payments. Another concern when using only data from untreated physicians from state A to estimate \hat{f}^{λ^*} is that unpaid physicians may have very different X values due to selection, resulting in poor out-of-sample predictions for treated physicians due to “insufficient covariate overlap” (Samii, Paler and Daly, 2016). We thus estimate \hat{f}^{λ^*} using only data from state B . Correspondingly, there should be sufficient covariate overlap between state A and state B .

ASSUMPTION 4: The range of observed covariates X is similar across the state that allows payments to physicians and the state that bans payments to physicians.

Taken together, assumptions 3 and 4 imply that states A and B should be similar, apart from the fact that payments are banned in state B . For this reason, in the application, we focus on estimating the treatment effect for physicians in the nearby

states of New Hampshire and Maine. We do not consider New York in the analysis because it is a far larger state than Vermont and therefore less comparable. We also exclude Massachusetts since this state has its own particular regulation of payments to physicians.

The final step towards uncovering the treatment effect θ is to compare the observed outcomes for the physicians in state A with their predicted outcomes in the absence of industry payments, where the difference can be specified as follows:

$$Y_{it}^A - \hat{f}^{\lambda^*}(X_{it}^A) = \gamma_t^A + \theta D_{it} + \epsilon_{it}$$

where $\epsilon_{it} := \xi(X_{it}^A, H_{it}^A) + U_{it}^A$. The error term ϵ_{it} comprises of the unobserved prediction error for physician i in state A , and the random noise term U_{it}^A . Given that f is the same for both treated and untreated physicians and assumption 2 holds, the prediction $\hat{f}^{\lambda^*}(X_{it}^A)$ and the prediction error $\xi(X_{it}^A, H_{it}^A)$ do not depend on whether the physician is treated or not.²² To obtain an unbiased estimate of θ we only require the further assumption that the prediction error has mean zero:

ASSUMPTION 5: The prediction error has mean zero, $E[\xi(X_{it}^A, H_{it}^A)] = E[\xi(X_{it}^B, H_{it}^B)] = 0$.

Hence, using OLS, we can estimate:

$$Y_i - \hat{f}^{\lambda^*}(X_i) = \hat{\gamma}_t^A + \hat{\theta} D_i \tag{1}$$

where $\hat{\theta}$ provides an estimate of the treatment effect. One final concern is that if there are spillovers from paid to unpaid doctors in the state where payments are permitted, then this would be incorporated into the state-year constant term and thus would affect the estimate of $\hat{\theta}$. In order to avoid this, when we estimate (1), we drop unpaid physicians that practice close (in same zip code at the 5-digit level) to paid physicians.

5.4.3 Estimators

Machine learning is suited to prediction problems, which have been traditionally been viewed as distinct from causal questions (Mullainathan and Spiess, 2017). However, recently a number of approaches have been developed that apply machine

²²In the empirical analysis, we will show that after controlling for the treatment, the mean of the error term is not significantly different from zero for physicians in the ban state, treated physicians and untreated physicians.

learning to aid in the identification of causal treatment effects (see e.g. McCaffrey et al., 2004; Hansen & Kozbur 2014; Wager & Athey, 2017; Chernozhukov et al., 2017; Athey & Imbens, 2019; Abrell et al., 2019). In problems of causal inference, a distinction can be made between the causal estimand of interest or target parameter, and nuisance parameters, which are not of interest per se, but are modelled in order to estimate the target parameter. Estimation of complex nuisance parameters can be thought of as a prediction problem for which ML methods are particularly well-suited (Chernozhukov et al., 2017).

In our application, θ is the target parameter and $f(X_{it}, H_{it})$ is the nuisance model. To predict physicians' prescription outcomes in the absence of payments we will use *supervised* machine learning. Supervised machine learning tools map an input to an output based on example input-output pairs in a 'training' dataset. In our application, data on physicians' characteristics and prescription outcomes in Vermont where payments are banned are used to train the ML algorithm.

The class of prediction models that we will apply are regularized linear regression models, also known as shrinkage or penalized regression methods. While regularized linear regression is only one of many methods in the toolbox of machine learning, it has some properties that make it attractive for empirical research. For one, it is a straightforward extension of linear regression. Similarly to ordinary least squares (OLS), regularized linear regression minimizes the sum of squared deviations between observed and model predicted values, but imposes a regularization penalty aimed at limiting model complexity. By reducing model complexity, regularized regression methods tend to outperform OLS in terms of out-of-sample prediction performance. The degree and type of penalization is determined by tuning parameters. The tuning parameters are typically chosen using K-fold cross-validation in order to optimize out-of-sample prediction performance.²³

By introducing tuning parameters we are allowing the data to determine which model provides the best prediction. This is different to standard econometric practice where we assume a data generating process and focus on estimating the model. Machine learning methods combine model selection with estimation. Economic intuition nevertheless plays an important role when determining which covariates make sense to include in the model.

²³The aim of cross-validation is to directly assess the performance of a model on unseen data. To this end, the data is repeatedly divided into a training and a test data set. The models are fit to the training data and the test data is used to assess the predictive performance. With K-fold cross-validation the data is divided into K equally sized folds. Each fold is used as a test set once while the remaining folds are used to fit the model. A summary measure of the model's performance across all test sets is used to select the best model. Thus, cross-validation can be used to compare models with different tuning parameters and select the tuning parameters that yield the best performance, e.g., the smallest out-of-sample mean squared prediction error.

We now provide a short overview of the three main regularized regression methods that will be applied in this paper.

Ridge regression (Tikhonov, 1963; Hoerl and Kennard, 1970) shrinks the regression coefficients, so that variables, with minor contribution to the outcome, have their coefficients close to zero. The shrinkage of the coefficients is achieved by penalizing the regression model with a penalty term called L2-norm, which is the sum of the squared coefficients. The tuning parameter $\lambda > 0$ controls the relative impact of the penalty. When $\lambda = 0$, the penalty term has no effect, and Ridge regression will produce the same coefficients as OLS. Given p regressors and n observations, Ridge regression minimizes the following loss function:

$$\hat{\beta}^{ridge}(\lambda) = \arg \min_{\beta \in \mathbb{R}^p} \left\{ \frac{1}{n} \sum_{i=1}^n (y_i - x_i' \beta)^2 + \lambda \sum_{j=1}^p \beta_j^2 \right\}$$

LASSO stands for Least Absolute Shrinkage and Selection Operator (Tibshirani, 1996). It shrinks the regression coefficients toward zero by penalizing the regression model with a penalty term called L1-norm, which is the sum of the absolute coefficients. The penalty has the effect of forcing some of the coefficient estimates to be exactly equal to zero, thus LASSO performs variable selection. It minimizes the following loss function:

$$\hat{\beta}^{lasso}(\lambda) = \arg \min_{\beta \in \mathbb{R}^p} \left\{ \frac{1}{n} \sum_{i=1}^n (y_i - x_i' \beta)^2 + \lambda \sum_{j=1}^p |\beta_j| \right\}$$

Elastic Net (Zou and Hastie, 2005) produces a regression model that is penalized by a linear combination of the LASSO and Ridge penalties, where λ_1 and λ_2 determine the relative weights. Elastic Net minimizes the following loss function:

$$\hat{\beta}^{elastic}(\lambda_1, \lambda_2) = \arg \min_{\beta \in \mathbb{R}^p} \left\{ \frac{1}{n} \sum_{i=1}^n (y_i - x_i' \beta)^2 + \lambda_1 \sum_{j=1}^p |\beta_j| + \lambda_2 \sum_{j=1}^p \beta_j^2 \right\}$$

5.4.4 Comparison with matching

In this section we discuss how our methodology compares with matching and highlight the advantages of our approach. Our methodology uses data on the prescribing behavior of physicians in Vermont to predict individual counterfactuals for each observation in Maine and New Hampshire. An alternative methodology would be to use a matching approach. The goal of matching is to pair treated units with non-treated units that have similar observable characteristics. The outcome values of these matched observations are then used to compute the counterfactual outcome

without treatment for the observation at hand.

With high dimensional data and many covariates finding a good match becomes increasingly difficult. In this context, matching is typically performed using propensity scores which are defined as the probability of receiving the treatment given the observed covariates (Rosenbaum and Rubin, 1983). Given the data at hand, it would be possible to match paid physicians in New Hampshire and Maine with physicians in Vermont based on their propensity score. The propensity score function can be estimated using the sample of physicians in New Hampshire and Maine where payments are possible.

Relative to such a matching approach, our methodology displays a number of advantages that make it better suited to the data and research question at hand. Propensity score matching results in matched pairs that are not necessarily similar across all of their covariates; rather, the goal of propensity score matching is that the sub-sampled treatment and control groups are similar to each other *on average* across all covariates. Further, misspecification of the propensity score model can lead to bad matches (King and Nielsen, 2019). Bad matches are particularly likely to be a concern with our data given that we have multiple treatments - some physicians receive in-kind payments, some receive cash and the size and number of transfers varies. Modelling the propensity score becomes complicated when there are multiple treatment levels. Our approach focuses on modelling the outcome variable rather than selection into the treatment. This provides us with a more precise counterfactual at the individual level, which in turn enables us to more accurately study how the treatment effect varies with different treatment levels and physician characteristics.

Propensity score matching also requires the researcher to make a number of modelling decisions, such as which variables to include when estimating the propensity score and which matching algorithm to apply. As there is no matching algorithm which dominates in all data situations, this decision needs to be made by the researcher and the final results will be sensitive to this choice (Caliendo and Kopeinig, 2005). With our approach, the choice of model is systematic and transparent as we include all possible confounders and choose the model that best fits the data. This prevents cherry-picking the model that gives the most appealing results and model selection can be easily replicated.

5.5 Estimation and results

This section describes the estimation process and compares the out-of-sample performance of all three penalized regression models, Ridge, LASSO and Elastic Net.

Thereafter we present the final estimates of the treatment effect.

5.5.1 Estimation and out-of-sample performance of the ML models

We model f for the two outcome variables of interest: 1) the brand prescription rate and 2) the average cost per 30-day prescription. X_{it} includes the full set of explanatory variables in the dataset (variables 12 - 33 in Table A2), indicators for each year and indicators for which quartile the physician falls into in terms of total 30-day prescriptions of anti-diabetics at the state-level, as well interactions between all variables. In total X_{it} contains 406 variables. If we included the full set of variables in a standard OLS regression model, the parameters would be poorly estimated. Additionally, many of the regressors will be multi-collinear resulting in poor out of sample predictions using OLS. The ML-estimators outlined above avoid this problem by shrinking the regression coefficients and performing variable selection.

X_{it} includes variables that affect both the treatment and outcome variable. As explained in Section 5.4, it is important that the included covariates are not affected by the receipt of an industry payment. If this was the case there would be an indirect effect on the outcome variable. The set of covariates contains fixed characteristics about the physician including their gender, specialization and whether or not he/she is a new practitioner. We include the share of anti-diabetic claims out of all claims as a further measure of the physician's specialization in diabetes drugs. X_{it} includes information about the region where the physician practices including the median household income, the population density, the percentage of adults diagnosed with diabetes, the number of adults diagnosed with diabetes, the percentage of obese adults, the number of obese adults, the number of other physicians practising in the same county and same 5-digit ZIP code. It includes detailed information about the patients of the physician including the share of male beneficiaries, the share of beneficiaries over the age of 65, average age of beneficiaries, the share with a MAPD, the share with a LIS and the average risk score (HCC) of beneficiaries. We include aggregate measures of the physician's prescribing behavior that are arguably unaffected by industry payments. We include the physician's total number of unique beneficiaries and the number of unique beneficiaries with a prescription for an anti-diabetic treatment. By including a count of the number of unique beneficiaries we are making the assumption that industry payments do not affect the decision to prescribe or not, but rather influence the decision of which drug to prescribe. We include the share of prescriptions that are for insulin. Given that insulin is the

only treatment option for diabetes type 1, it is reasonable to assume that payments would not affect the decision to prescribe insulin vs. another type of drug (but could affect which brand of insulin to prescribe). Finally we include the quartile that the physician falls into in terms of total 30-day prescriptions of anti-diabetics at the state-level. Prescription volumes affect selection in the treatment. Under the assumption that payments affect the drug selected and not aggregate volumes, this measure should not be affected by receiving a payment. A full description of these variables can be found in the Appendix Table 5.9.

To estimate the models, the Vermont dataset is randomly split into a ‘training’ and ‘test’ set. 80 percent of observations ($N = 800$) are used for the training set and 20 percent ($N=199$) are retained in the test or ‘hold-out’ set. The models are fit to the training set. For each method (LASSO, Ridge and Elastic Net) we use 5-fold cross-validation to select the optimal tuning parameters out of a grid of possible values.²⁴ For each method, the final model (with the optimal tuning parameters) is used to predict the outcome variable for the unseen test data. On the basis of these predictions the test MSPE is calculated to evaluate the out-of-sample predictive performance of each model.

Table 5.4 and Table 5.5 present the out-of-sample performance for the brand prescription rate and the drug cost per 30-day supply respectively. Lower MSPE and higher R-squared values are indicative of good predictive performance. The R-squared across all models and for both outcome variables is relatively high indicating that the included covariates are able to explain a high proportion of the variability in the outcome variable. The models are able to predict the brand prescription rate more accurately than the drug cost per 30-day supply. The best model, that minimizes the test MSPE, for the brand prescription rate is LASSO. Ridge is best able to predict the drug cost per 30-day supply. For each outcome variable the final model is used to calculate $\hat{f}^{\lambda^*}(X_i)$ for all observations, including those in Maine and New Hampshire.

	MSPE	S.E. for MSPE	R-squared
LASSO	0.00364	0.00041	0.84259
Ridge	0.00369	0.00042	0.84019
Elastic Net	0.00370	0.00042	0.84007

Table 5.4: Out-of-sample performance - Brand prescription rate

²⁴This is implemented in R using the *caret* package. For more information see <https://cran.r-project.org/web/packages/caret/caret.pdf>

	MSPE	S.E. for MSPE	R-squared
LASSO	1127.11116	149.57987	0.78206
Ridge	1077.74990	115.58213	0.79267
Elastic Net	1086.94209	120.77792	0.79237

Table 5.5: Out-of-sample performance - Drug cost per 30-day supply

5.5.2 Effect of industry payments

Equipped with the predicted outcome values for all physicians we estimate equation (1) using OLS for both outcome variables. The results are presented in Table 5.6 and Table 5.7. To account for the fact that predicted values have been used in the regression we present the bootstrapped standard errors.

First, we regress the difference between the true and predicted brand prescription rate for each observation on state-year dummies and a treatment indicator variable which takes the value 1 if physician i receives any positive payment in year t . Column (1) in Table 5.6 presents the result. Receipt of an industry payment, on average, causes the brand prescription rate of a physician to increase by 5 percentage points in the year in which the payment is received. This corresponds to a 20% increase in the unconditional average brand prescription rate in the states where payments are permitted. In Table 5.7 column (1) we do the same for the drug cost per 30-day supply. We find that the drug cost per dose is USD 21 higher for physicians who receive industry payments, corresponding to a 22% increase in the unconditional average drug cost.

Next, we investigate heterogeneity in the average treatment effect. We find that for physicians with a low LIS claim share (less than 25%) the average size of the effect of a payment is a 2 percentage point increase in the brand prescription rate, whereas for physicians with a high LIS claim share the effect is an 8.8 percentage point increase (see Table 5.6 column 2). The average drug cost per 30-day supply is also higher for physicians with a high LIS claim share (see Table 5.7 column 2).

A potential concern is that there may be variables that are correlated with having a high LIS share that may affect responsiveness to the treatment. For example, if pharmaceutical companies actively target doctors with a high LIS claim share and pay them more frequently, the variability in the treatment effect with LIS claim share may be driven by the fact that physicians with a high LIS claim share receive more payments. Another potential confounding factor is median income in the area where the physician works. If physicians with a high share of patients with a low income subsidy also operate in impoverished areas, the physicians themselves may

earn less and this may affect how they respond to payments from pharmaceutical companies.

In column (3) in Table 5.6 and Table 5.7 we control for potential confounding factors. After including potential confounders, the coefficient on the share of beneficiaries with a LIS remains positive and significant. A 10 percentage point increase in the share of LIS patients increases the brand prescription rate by 1 percentage point, all else held constant. A 10 percentage point increase in the share of LIS patients increases the drug cost per 30-day supply by 1.7 USD. Additionally, we find that male physicians are more likely to prescribe more expensive drugs and brand drugs in response to receiving payments from pharmaceutical companies. Specialists (in the treatment of diabetes), by contrast, react less strongly. An increase in the number of in-kind payments is significantly and positively associated with an increase in the average cost per 30-day supply. In sum, our analysis corroborates the hypothesis that physicians with a higher share of patients that have lower out-of-pocket costs are more likely to prescribe more expensive drugs and brand drugs in response to receiving a payment.

In order to test the validity of our approach and provide support for assumption 5 made in section 5.2, we compare the residuals of the regression in column (3) for physicians in Vermont, paid physicians and unpaid physicians. After controlling for the treatment, we find that the average error ϵ_{it} for all three groups is not significantly different from zero. Figure 5.3 and 5.4 in the Appendix show the distribution for the error term for paid and unpaid physicians in the states where payments are permitted. In the Appendix Section 5.8.3 we provide an illustrative model which produces insights on the interaction between drug costs, insurance coverage and payments to physicians that are in line with our findings.

Finally, we use our estimates to provide a rough estimate of the potential cost savings from a policy to eliminate payments to physicians. During the time period of the sample (2014-2017) a total of just under USD 900,000 was spent by the pharmaceutical industry on payments related to anti-diabetics to 245 physicians in the states of New Hampshire and Maine. These physicians account for 18% of prescription volumes. Multiplying the average increase in the cost per 30-day prescription by the total number of 30-day prescription by paid physicians indicates that, holding prices and prescribing volumes constant, the elimination of payments would result in savings of just under USD 10 million during the sample period. The total cost of all prescriptions for diabetes medication in New Hampshire and Maine during the sample period amounts to USD 306 million. Thus by eliminating payments, total prescription costs could be reduced by approximately 3%. Higher savings come from physicians with a higher share of patients receiving a low-income

subsidy. Specifically, breaking down where this cost increase comes from, we find that 72% of the cost increase is driven by physicians who have an LIS claim share of above 50%, 19% of the cost increase is driven by physicians who have an LIS claim share between 25% and 50% and the remaining 9% is attributable to physicians with a low share of LIS patients between 0 and 25%. On the basis of our analysis, it also clear that payments to physicians are very profitable for the pharmaceutical industry as a whole - the increase in drug costs is more than 10 times what the industry spent on payments. In the following section we discuss our findings.

Table 5.6: Results - Brand prescription rate

	(1)	(2)	(3)
D - I(Payment>0)	0.0514*** (0.00456)		
D X (LIS claim share ≤ 0.25)		0.0205* (0.0106)	
D X ($0.25 < \text{LIS claim share} \leq 0.5$)		0.0334*** (0.00928)	
D X ($0.5 < \text{LIS claim share} \leq 0.75$)		0.0602*** (0.00755)	
D X ($0.75 < \text{LIS claim share} \leq 1$)		0.0878*** (0.0216)	
D X LIS claim share			0.100*** (0.0154)
D X Median household income			-0.120*** (0.0239)
D X Male			0.0810*** (0.0124)
D X Specialist			-0.0317** (0.0124)
D X No. in-kind payments			0.000399 (0.000302)
D X No. cash payments			-0.00137 (0.00115)
Observations	3775	3775	3775
R-squared	0.173	0.180	0.199
State-year FE	Yes	Yes	Yes

Notes: Bootstrap standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5.7: Results - Drug cost per 30-day supply

	(1)	(2)	(3)
D - I(Payment>0)	21.12*** (2.083)		
D X (LIS claim share \leq 0.25)		20.91*** (4.794)	
D X (0.25 < LIS claim share \leq 0.5)		13.38*** (3.608)	
D X (0.5 < LIS claim share \leq 0.75)		23.83*** (3.371)	
D X (0.75 < LIS claim share \leq 1)		28.60*** (6.536)	
D X LIS claim share			17.09* (8.996)
D X Median household income			-12.40 (10.85)
D X Male			26.40*** (7.104)
D X Specialist			-14.18*** (4.277)
D X No. in-kind payments			0.216** (0.101)
D X No. cash payments			-1.466*** (0.457)
Observations	3775	3775	3775
R-squared	0.0889	0.0915	0.105
State-year FE	Yes	Yes	Yes

Notes: Bootstrap standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5.6 Discussion

In this section we discuss our findings in relation to other studies, we state some limitations of this research, and outline implications for policy makers. Our results are not directly comparable to the literature. Previous studies have focused on different drugs, time spans and outcomes measures. Although not directly comparable, our findings are within the range of previous results. Using the full sample of all prescribed drugs in Medicare Part D for the year 2013, Qian et al. (2017) find that on average receipt of industry payments of less than USD 100 is associated with a 2% reduction in annual generic drug prescribing rate and receipt of payments of more than USD 500 is associated with a 5% reduction in annual generic drug prescribing rate. Perlis and Perlis (2016) find that being in the top quintile of payment receipt is associated with an incremental prescribing cost per patient ranging from USD 27 (general surgery) to USD 2931 (neurology). Both Qian et al. (2017) and Perlis and Perlis (2016) highlight that their results are associations, and cannot be interpreted causally.

Carey et al. (2020) measure the effect of *specific* payments from companies on the prescription of *that* company's drugs. Using data on all drugs in the Open Payments database, they find that a single payment raises expenditures on the *paying* drug by USD 121 during the first year. Focusing on anticoagulants drugs, Agha and Zeltzer (2019) estimate the effect of a payment for a *specific* drug on the volume prescribed of *that* drug and find that small payments for food 8 percent increase over the average prescription volume. Using data on payments from medical device companies to physicians combined with hospital discharge datasets, Amaral-Garcia (2020) assesses the impact of payments on treatment provided to heart attack patients arriving at the Emergency Room in Florida hospitals. She finds that patients treated by doctors who interact with the industry are more likely to have higher medical device costs (up to 16% increase) and total hospital costs (up to 3% increase). Thus, our findings are in line with the general conclusion of other studies; namely that industry payments influence the prescription decisions of physicians and can lead physicians to choose more expensive treatments.

We contribute to the existing literature by providing an estimate of the causal effect of payments on the cost of treatment for diabetes and show that this effect varies with insurance coverage. Our estimates suggest that, holding prices and prescribing volumes constant, issuing a payment ban similar to Vermont's can result in a 3% decline in total prescription costs for diabetes. These savings arise due to physicians substituting towards cheaper anti-diabetic drugs. In the absence of industry payments, physicians are motivated to choose the best and cheapest treatment for

their patient because the patient will face some out-of-pocket expenditures. Thus, our estimate of the potential savings from a payment ban also assumes that current levels of insurance coverage are held constant. If out-of-pocket costs are reduced, drug costs will increase.

If physicians alter the quantity they prescribe in reaction to industry payments e.g. if a physician prescribes two of a specific type of pill instead of one, or if he/she prescribes a treatment for two months instead of one, then our estimate provides a lower bound on the savings from a payment ban. Often specific drugs are associated with a set recommended dosage, thus it is reasonable to assume that savings due to a payment ban will be driven largely by reductions in the cost *per 30-day dose* as opposed to a reduction in the aggregate number of doses prescribed.

Further savings would be possible if drug prices fell in response to a payment ban, however addressing this effect is outside of the scope of this paper. To investigate price changes would require a structural model approach that takes into account how the pricing decisions of firms is influenced by payments to physicians. Pricing decisions in the pharmaceutical industry are determined by a number of features and are bargained between insurance plans and pharmaceutical suppliers, thus this would be not be a straightforward exercise. Moreover, prices are typically set nationally and so incorporating price effects would be more important when assessing a nationwide ban on payments as opposed to a state-level ban. A further limitation of this study is that we focus on diabetes treatments, thus we cannot comment on the exact magnitude of cost declines for other therapeutic fields.

The financial cost to the public health care system needs to be traded-off against the possible benefits of industry payments to physicians, for example, as industry claims, payments may be associated with information that helps physicians to make better prescription decisions. In fact, there is little evidence that payments help physicians to make better prescribing decisions. Focusing on three major therapeutic classes, Carey et al. (2020) find that physicians who receive payments tend to prescribe lower quality drugs after the payment, but that the effect is very small. Amaral-Garcia (2020) finds no significant impact of industry interactions on the healthcare outcomes of patients in relation to treatment provided to heart attack patients. To credibly assess whether paid physicians in our dataset make better prescribing decisions would require access to data at the patient level which is not publicly available.

5.7 Concluding remarks

In this paper we quantify the extent to which industry payments to physicians lead to higher drug costs using a novel identification strategy. Furthermore, we test the hypothesis that physicians with patients who have lower out-of-pocket expenditures are likely to prescribe relatively more costly drugs in reaction to industry payments. If physicians care about their patients' financial health, in addition to their physical health, then we expect that physicians will trade off the drug cost faced by the patient with the personal reward associated with the prescription of more expensive brand medications. We focus on the prescription of anti-diabetic medication, and find that receipt of payments related to anti-diabetic medication increases the average brand prescription rate by 5 percentage points, and the average drug cost per dose by 21 USD. Physicians with a higher share of patients with a low income subsidy, and who therefore face lower out-of-pocket expenditures, prescribe relatively more expensive medication in response to receiving a payments. This effect remains after controlling for other potential confounders that could drive heterogeneity in the treatment effect.

Back-of-the-envelope calculations based on our estimates suggest that, holding prices and prescribing volumes constant, issuing a payment ban similar to Vermont's can result in a 3% decline in total prescription costs for diabetes. These savings arise due to physicians substituting towards cheaper anti-diabetic drugs. In the absence of industry payments, physicians are motivated to choose the best and cheapest treatment for their patient because of the fact that patients face some out-of-pocket expenditures. Given that existing research does not point to strong informational benefits of payments to physicians, a ban on industry payments is likely to be an effective way to contribute towards health care cost containment in the US.

5.8 Appendix

5.8.1 Tables and figures

Firm	Brand
ASTRAZENECA	BYDUREON
ASTRAZENECA	FARXIGA
ASTRAZENECA	ONGLYZA
BOEHRINGER INGELHEIM	GLYXAMBI
BOEHRINGER INGELHEIM	JARDIANCE
BOEHRINGER INGELHEIM	JENTADUETO
BOEHRINGER INGELHEIM	TRADJENTA
ELI LILLY	BASAGLAR
ELI LILLY	HUMALOG
ELI LILLY	HUMULIN
ELI LILLY	TRULICITY
GLAXOSMITHKLINE	TANZEUM
JANSSEN (SUBSIDIARY OF PFIZER)	INVOKAMET
JANSSEN (SUBSIDIARY OF PFIZER)	INVOKANA
MERCK	JANUMET
MERCK	JANUVIA
NOVO NORDISK	LEVEMIR
NOVO NORDISK	NOVOLOG
NOVO NORDISK	TRESIBA
NOVO NORDISK	VICTOZA
SANOFI	LANTUS
SANOFI	TOUJEO SOLOSTAR
SANTARUS	GLUMETZA
TAKEDA	KAZANO

Table 5.8: Firms and brands with payments to physicians in the sample

No.	Variable	Description	Source
1	No. payments per year	Total number of payments for physician i in year t	Open Payments
2	No. in-kind payments per year	Total number of in-kind payments for physician i in year t	Open Payments
3	No. cash payments per year	Total number of cash payments for physician i in year t	Open Payments
4	Value (USD) payments per year	Value (USD) of payments for physician i in year t	Open Payments
5	Value (USD) in-kind payments per year	Value (USD) of in-kind payments for physician i in year t	Open Payments
6	Value (USD) cash payments per year	Value (USD) of cash payments for physician i in year t	Open Payments
7	State	State where the physician is located: NH, ME, VT	Part D Prescriber Summary Table, Open Payments
8	Year	Years: 2013, 2014, 2014, 2016, 2017	Part D Detailed Data File, Open Payments
9	Brand prescription rate, anti-diabetics	Share of 30-day supply brand anti-diabetic drugs out of total 30-day supply of anti-diabetic drugs for physician i in year t	Part D Detailed Data File
10	Total 30-day supply, anti-diabetics	Total 30-day supply prescriptions for anti-diabetic drugs for physician i in year t	Part D Detailed Data File
11	Drug cost '000, anti-diabetics	Total drug cost in 1000s for anti-diabetic drugs for physician i in year t . The drug cost is based on the amount paid by the Part D plan, the beneficiary, government subsidies, and any other third-party payers.	Part D Detailed Data File
12	Specialist in diabetes (0/1)	Indicator taking the value 1 if speciality is recorded as Endocrinology, Diabetes, Internal Medicine and/or Specialist	Part D Prescriber Summary Table, Open Payments
13	Family practice (0/1)	Indicator taking the value 1 if speciality is recorded as Family Practice, Family Medicine and/or General Practice	Part D Prescriber Summary Table, Open Payments
14	New practitioner (0/1)	Physicians with a provider enumeration year (date of NPI assignment) later than or equal to 2008	NPPES NPI Registry
15	Male practitioner (0/1)	Indicator taking the value 1 if physician gender is male	Part D Detailed Data File

Table 5.9: Variable definitions

No.	Variable	Description	Source
16	Beneficiary count for anti-diabetics	Number of unique Medicare Part D beneficiaries prescribed an anti-diabetic drug by physician i in year t	Part D Detailed Data File
17	Anti-diabetics claim share	Share of claims for anti-diabetic drugs out of all claims for physician i in year t	Part D Detailed Data File
18	Insulin claim share of anti-diabetics	Share of claims for insulin out of all claims for anti-diabetic drugs for physician i in year t	Part D Detailed Data File
19	Beneficiaries over 65 for anti-diabetics	Share of beneficiaries age 65 and older with at least one claim for an anti-diabetic drug for physician i in year t	Part D Detailed Data File
20	Share of male beneficiaries	Share of male beneficiaries for physician i in year t	Part D Prescriber Summary Table
21	Total beneficiary count	Number of unique Medicare Part D beneficiaries prescribed any drug by physician i in year t	Part D Prescriber Summary Table
22	MAPD claim share	Share of total claims attributable to beneficiaries covered by MAPD plans for physician i in year t	Part D Prescriber Summary Table
23	LIS claim share	Share of total claims attributable to beneficiaries with a Part D low-income subsidy for physician i in year t	Part D Prescriber Summary Table
24	Average age of beneficiaries	Average age of beneficiaries for physician i in year t (Beneficiary age is calculated at the end of the calendar year or at the time of death)	Part D Prescriber Summary Table
25	Average risk score of beneficiaries	Beneficiary average Hierarchical Condition Category (HCC) risk score for physician i in year t . The risk score is based on disease and demographic risk factors.	Part D Prescriber Summary Table
26	Median household income '000	Median household income in the past 12 months in dollars for the area (5-digit ZIP code) in which physician i is located in year t	American Community Survey
27	Population density '000	Total population divided by area in square miles for the area (5-digit ZIP code) in which physician i is located in year t	American Community Survey

No.	Variable	Description	Source
28	Percent diagnosed with diabetes	Diagnosed Diabetes, Adults (20+) with Diabetes, Age-Adjusted Percentage for the county in which physician i is located in year t-1	Centers for Disease Control and Prevention
29	No. diagnosed with diabetes '000	Diagnosed Diabetes, Adults Aged 20+ Years, Crude Number for the county in which physician i is located in year t-1	Centers for Disease Control and Prevention
30	Percent obese	Obesity, Adults Aged 20+ Years, Age-Adjusted Percentage for the county in which physician i is located in year t-1	Centers for Disease Control and Prevention
31	No. obese '000	Obesity, Adults Aged 20+ Years, Crude Number for the county in which physician i is located in year t-1	Centers for Disease Control and Prevention
32	No. other physicians in same county	Number of other physicians in sample in same county for physician i in year t	Part D Detailed Data File
33	No. other physicians in same ZIP	Number of other physicians in sample in same 5-digit ZIP code area for physician i in year t	Part D Detailed Data File

Figure 5.3: Distribution of error term - Brand prescription rate

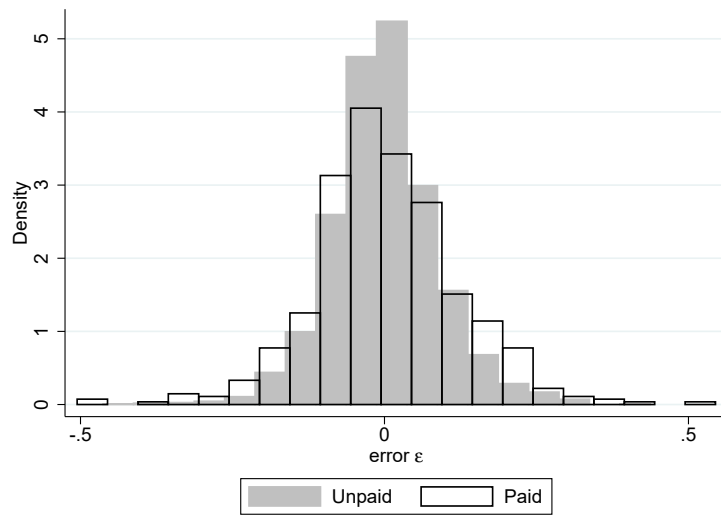



Figure 5.4: Distribution of error term - Drug cost per 30-day Supply



5.8.2 Treatments for diabetes

This Appendix provides an overview of the different treatments for diabetes and provides more information on how physicians prescribe treatments for diabetes.

Type 1 patients are treated exclusively with insulin. Typically the treatment involves several insulins simultaneously. A patient can be treated with “rapid”, “long” or “intermediate” insulin. An interview with a diabetologist revealed that physicians may change the type of insulin prescribed to a patient within these groups due to reasons such as side effects, patient intolerance, insurance reimbursement, and due to the introduction of new insulin that is generally perceived as a better product by the physician.

Type 2 patients are treated in a more complex way. The first line treatment is Metformin which is used for as long as the body tolerates it, thereafter different drugs are added to the treatment regimen. After Metformin no single drug is highlighted according to the Standards of Medical Care in Diabetes (Centers for Disease Control and Prevention, 2017; American Diabetes Association 2015). Thus, physicians can choose between prescribing drugs from different drug classes for example DPP-4 inhibitors or SGLT-2 inhibitors. Therapy often combines several drugs from different classes. Type 2 patients with a severe condition are treated with insulin therapy.

Diabetes drugs can be grouped in several classes:

- **Sensitizers**

Insulin sensitizers improve sensitivity to insulin so that glucose can be absorbed.

Biguanides/Metformin: Glucophage (1995q1)²⁵, Riomet (2003q3), Fortamet (2004q2), Glumetza (2005q2), Actoplus Met (2005q3)

TZDs (Thiazolidinediones): Avandia (1999q2), Actos (1999q3), Avandamet (2002q4), Avandaryl (2005q4), Duetact (2006q3)

- **Insulins**

Insulin treatment is used to keep blood sugar levels within the target range. Insulin must be used for patients with Type 1 diabetes. Insulin is usually administered via injection.

Rapid and intermediate acting insulins: Humulin (1982q4), Novolin (1991q2), Humalog (1996q2), Novolog (2000q2), Apidra (2004q2), Afrezza (2014q2)

Long acting insulins: Lantus (2000q2), Levemir (2005q2), Toujeo (2015q1), Tresiba (2015q3), Basaglar (2015q4), Xultophy (2016q4), Soliqua (2016q4)

- **GLP-1 receptor agonists**

GLP-1 receptor agonists bind to the membrane GLP-1 receptor, preventing uptake of GLP-1 from the blood. This raises the level of blood GLP-1, stimulating insulin secretion and suppressing glucagon secretion.

Drugs: Byetta (2005q2), Victoza (2010q1), Bydureon (2012q1), Tanzeum (2014q2), Trulicity (2014q3), Adlyxin (2016q3)

- **DPP-4 inhibitors**

DPP-4 inhibitors block the action of the DPP-4, an enzyme that inactivates the incretin GLP-1 that helps the body produce more insulin.

²⁵Date of FDA approval in parentheses

Drugs: Januvia (2006q4), Janumet (2007q1), Onglyza (2009q3), Kombiglyze XR (2010q4), Tradjenta (2011q2), Jentadueto (2012q1), Nesina (2013q1), Oseni (2013q1), Kazano (2013q1),

- **SGLT-2 inhibitors**

SGLT-2 inhibitors are a new group of oral medications used for treating type 2 diabetes, approved in 2013. They inhibit the sodiumglucose transport proteins (SGLT-2) that help re-absorb glucose into the blood, and pass out the excess glucose as urine.

Drugs: Invokana (2013q1), Farxiga (2014q1), Invokamet (2014q3), Jardiance (2014q3), Xigduo (2014q4), Glyxambi (2015q1), Synjardy (2015q3)

- **Alpha-glucosidase inhibitors**

These agents do not have a direct effect on insulin secretion or sensitivity, but rather slow the digestion of starch so that glucose enters the bloodstream more slowly.

Drugs: Precose (1995q3), Glyset (1996q4)

- **Secretagogues**

Secretagogues are drugs that increase insulin output from the pancreas.

Sulfonylureas: Glucotrol (1984q2), Diabeta (1984q2), Glynase (1992q1), Glyburide Micronized (1992q2), Amaryl (1995q4), Glucovance (2000q3), Metaglip (2002q4)

Non-sulfonylurea secretagogues/Meglitinides: Prandin (1997q4), Starlix (2000q4), Prandimet (2008q2)

- **Injectable amylin analogues**

Amylin agonist analogues slow gastric emptying and suppress glucagon. They have all the incretins actions except stimulation of insulin secretion. Currently, pramlintide (trade name Symlin) is the only clinically available amylin analogue.


Drugs: Symlin (2005q1)

Table 5.10 provides a list of all diabetes treatments that are in the final dataset and which drug class they belong to. Figure 5.5 provides insight into which drug classes are the most frequently prescribed. Figure 5.6 shows the break down of total drug costs by drug class. It is evident that while the most commonly prescribed drug class is sensitizers, the prescription of insulin is what drives total drug costs.

Firm and drug name	Drug class
PFIZER GLYSET	Alpha-glucosidase
GENERIC COMPANY PRECOSE	Alpha-glucosidase
ASTRAZENECA SYMLIN	Amylin agonists
MERCK JANUMET	DPP-4 inhibitors
BOEHRINGER INGELHEIM TRADJENTA	DPP-4 inhibitors
MERCK JANUVIA	DPP-4 inhibitors
TAKEDA KAZANO	DPP-4 inhibitors
GENERIC COMPANY NESINA	DPP-4 inhibitors
BOEHRINGER INGELHEIM JENTADUETO	DPP-4 inhibitors
ASTRAZENECA ONGLYZA	DPP-4 inhibitors
ASTRAZENECA KOMBIGLYZE XR	DPP-4 inhibitors
TAKEDA OSENI	DPP-4 inhibitors
TAKEDA NESINA	DPP-4 inhibitors
GLAXOSMITHKLINE TANZEUM	GLP-1 agonists
NOVO NORDISK VICTOZA	GLP-1 agonists
ASTRAZENECA BYDUREON	GLP-1 agonists
ASTRAZENECA BYETTA	GLP-1 agonists
ELI LILLY TRULICITY	GLP-1 agonists
NOVO NORDISK TRESIBA	Insulins - Long
NOVO NORDISK LEVEMIR	Insulins - Long
ELI LILLY BASAGLAR	Insulins - Long
SANOFI TOUJEO SOLOSTAR	Insulins - Long
SANOFI LANTUS	Insulins - Long
ELI LILLY HUMULIN	Insulins - Rapid & intermediate
NOVO NORDISK NOVOLOG	Insulins - Rapid & intermediate
MANNKIND AFREZZA	Insulins - Rapid & intermediate
ELI LILLY HUMALOG	Insulins - Rapid & intermediate
NOVO NORDISK NOVOLIN	Insulins - Rapid & intermediate
SANOFI APIDRA	Insulins - Rapid & intermediate
BOEHRINGER INGELHEIM GLYXAMBI	SGLT-2 inhibitors
JANSSEN (SUB. PFIZER) INVOKAMET	SGLT-2 inhibitors
BOEHRINGER INGELHEIM SYNJARDY	SGLT-2 inhibitors
BOEHRINGER INGELHEIM JARDIANCE	SGLT-2 inhibitors
JANSSEN (SUB. PFIZER) INVOKANA	SGLT-2 inhibitors
ASTRAZENECA FARXIGA	SGLT-2 inhibitors
ASTRAZENECA XIGDUO	SGLT-2 inhibitors
NOVARTIS STARLIX	Secretagogues - Non-sulfonylureas
GEMINI LABORATORIES PRANDIN	Secretagogues - Non-sulfonylureas
GENERIC COMPANY STARLIX	Secretagogues - Non-sulfonylureas
GENERIC COMPANY PRANDIN	Secretagogues - Non-sulfonylureas
GENERIC COMPANY GLYBURIDE (MICRONIZED)	Secretagogues - Sulfonylureas
SANOFI AMARYL	Secretagogues - Sulfonylureas
GENERIC COMPANY GLYBURIDE BRAND	Secretagogues - Sulfonylureas
GENERIC COMPANY GLUCOTROL	Secretagogues - Sulfonylureas
BRISTOL MYERS SQUIBB GLUCOVANCE	Secretagogues - Sulfonylureas
PFIZER GLUCOTROL	Secretagogues - Sulfonylureas
GENERIC COMPANY METAGLIP	Secretagogues - Sulfonylureas
GENERIC COMPANY AMARYL	Secretagogues - Sulfonylureas
GENERIC COMPANY GLUCOVANCE	Secretagogues - Sulfonylureas
SANTARUS GLUMETZA	Sensitizers - Metformin
TAKEDA ACTOPLUS MET	Sensitizers - Metformin
ANDRX LABS FORTAMET	Sensitizers - Metformin
SUN PHARMACEUTICAL RIOMET	Sensitizers - Metformin
GENERIC COMPANY ACTOPLUS MET	Sensitizers - Metformin
GENERIC COMPANY METFORMIN BRAND	Sensitizers - Metformin
BRISTOL MYERS SQUIBB GLUCOPHAGE	Sensitizers - Metformin
TAKEDA ACTOS	Sensitizers - TZDs
GENERIC COMPANY ACTOS	Sensitizers - TZDs


Table 5.10: Diabetes treatments in the sample

Figure 5.5: Prescriptions by drug class



Notes: This figure is created using data from Vermont, New Hampshire and Maine.

Figure 5.6: Total drug cost by drug class



Notes: This figure is created using data from Vermont, New Hampshire and Maine.

5.8.3 Illustrative model

This Appendix provides an illustrative model which yields insights that are in line with the empirical findings of this paper. The aim of this simple model is to highlight the interaction between industry payments for brand name prescriptions and insurance coverage. The underlying assumption is that physicians act on behalf of their patients, weighing up both their patient's financial and clinical health. However, physicians are not perfect agents, and also take into account the expected personal benefits, in the form of gifts and other transfers from pharmaceutical companies, that are associated with prescribing certain drugs.

We consider a physician's decision when faced with two clinically substitutable drugs. In particular we focus on a physician's choice between prescribing a brand or generic version of a specific drug, since here it is clear that the drugs are equally effective. However, the set-up can be equally applied to the case of any two substitutable drugs with differing prices, where the higher priced drug is perceived to be at least as effective as the lower priced alternative. First, we analyse how the probability to prescribe the more expensive brand version of a drug varies with insurance coverage in a context where industry payments are banned. Thereafter we extend the model to the case where physicians receive payments for brand prescriptions.

Given that both drugs are substitutable from a treatment perspective, the physician takes into account the patient's preferences for the brand vs. generic version of the product, U^B and U^G , and the respective prices the patient faces for each, P^B and P^G . We assume that brand drugs are always at least as expensive as their generic counterparts such that $P^B \geq P^G$ and are perceived by the patient to be at least as effective as the generic $U^B \geq U^G$. The latter assumption is in line with survey evidence from the US that patients (and sometimes physicians too) can have negative perceptions about the efficacy and quality of generic medications (Shrank et al., 2011; Shrank et al., 2007).²⁶ One potential explanation is that brand firms with an established reputation are perceived as producing safer and more reliable drugs.²⁷ Prices are assumed to be determined exogenously of the model. Patients can be partially or fully insured against drug costs. Let $\tau \in (0, 1)$ denote insurance coverage such that the out-of-pocket expenditures incurred by the patient are $(1 - \tau)P$. It follows that the physician will prescribe the brand drug if:

$$U^B - (1 - \tau)P^B \geq U^G + (1 - \tau)P^G \quad (5.1)$$

Without a loss of generality, we set the price of the generic P^G to zero and define $K =: U^B - U^G$, where it is clear that $K \geq 0$. Thus, the physician will prescribe the brand drug when

$$P^B \leq P^* \equiv \frac{K}{1 - \tau}. \quad (5.2)$$

When the patient has no insurance coverage the cut-off will be equal to K ; the patient's perceived additional benefit from the brand vs. the generic. As the level of insurance coverage increases, the cut-off value P^* increases at an increasing rate as,

²⁶In a survey of patients by Shrank et al. (2007) over 60% of respondents stated that they preferred brand medication to generic medication. Over 23% of physicians surveyed by Shank et al. (2011) expressed negative perceptions about efficacy of generic drugs and almost 50% reported negative perceptions about quality of generic medications.

²⁷In one instance, the FDA determined that a generic anti-depressant performed less well than its branded counterpart, likely due to differences in their "extended release" coatings (Thomas, 2012). A widely publicized recall of generic acetaminophen in 2006 resulted from the discovery that some pills could contain metal fragments (Associated Press, 2006).

$$\frac{\partial P^*}{\partial \tau} = \frac{K}{(1-\tau)^2} \geq 0 \quad (5.3)$$

$$\frac{\partial^2 P^*}{\partial \tau^2} = \frac{2K}{(1-\tau)^3} \geq 0 \quad (5.4)$$

Thus even a small perceived benefit of the brand, can lead to the prescription of a relatively more expensive brand drug when insurance coverage is high.

We now consider what would happen if physicians receive some benefit from the brand company if they prescribe the brand drug. The benefit is set as a fixed percentage ρ of the brand price where $0 < \rho < 1$. This reward scheme is motivated by the observation in the data that more expensive brand drugs are generally associated with higher payments (see Table 5.1 in Section 5.3). This assumption is not crucial for the results of the model. Indeed, if we were to assume that physicians receive a lump sum transfer for prescribing the brand drug, all predictions go through.

Underlying this set-up is an assumption that payments are regular and linked with prescription volumes. The next section provides statistics that corroborate the claim that ‘paid’ physicians receive multiple payments from the same pharmaceutical company, and often also in relation to the same product. Sales representatives keenly monitor physicians prescription data and are known to sometimes closely associate prescribing volumes with an expectation. An example provided by Shahram Ahari, a former pharmaceutical sales representative for Eli Lilly, is “*So, doc, you’ll choose Drug X for the next 5 patients who are depressed and with low energy? Oh, and don’t forget dinner at Nobu next month. I’d love to meet your wife*” (Fugh-Berman and Ahari, 2007 p. 622).

The addition of a brand-related payment implies that the physicians now trades off the costs and benefits that the patient faces, with their own personal gain. The physician will prescribe the brand drug when

$$U^B - (1-\tau)P^B + \rho P^B \geq U^G + (1-\tau)P^G, \quad (5.5)$$

which can, employing the same notation as before, be rewritten as,

$$P^B \leq P'^* \equiv \frac{K}{1-\tau-\rho} \quad (5.6)$$

For any level of insurance coverage, for a positive value of ρ , $P'^* > P^*$. Thus, the probability of prescribing the brand drug is always higher for a physician who receives payments. In the case where $\tau + \rho \geq 1$, the brand will always be prescribed, regardless of it’s price. Thus at high levels of insurance coverage, even a small ρ can successfully incentivize brand prescription.

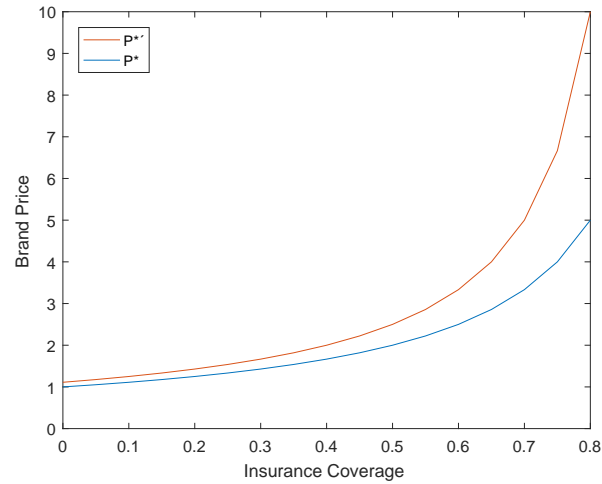
The model offers one final insight on how the probability of brand prescription differs between physicians who receive payments and those who do not. We define $\Delta := P'^* - P^*$. Taking the partial derivative Δ of with respect to τ ,

$$\frac{\partial \Delta}{\partial \tau} = \frac{K}{(1-\tau-\rho)^2} - \frac{K}{(1-\tau)^2} \geq 0, \quad (5.7)$$

it is evident that as insurance coverage increases, physicians receiving payments increase their brand prescriptions at a faster rate. Figure 5.7, where K has been set to 1 and ρ has been set to 0.1, illustrates this feature. The intuition behind this result is that as insurance coverage increases, the patient is more willing to accept higher brand prices which in turn allows for higher benefits

to the physician, encouraging brand prescriptions.²⁸

Figure 5.7: Price threshold for brand prescription



In sum, this simple model yields insights which are in line with the empirical findings in this paper. Firstly, physicians expecting a reward are more likely to prescribe the brand drug (for any level of insurance coverage). Secondly, as the patient's out-of-pocket expenditures decline, physicians receiving payments become *increasingly* more likely to prescribe the brand.

²⁸To see that this is also the case if we assume physicians receive a fixed lump sum Q , let $\Delta := \frac{K+Q}{(1-\tau)} - P^*$. Thus $\frac{\partial \Delta}{\partial \tau} = \frac{Q}{(1-\tau)^2} \geq 0$.

Bibliography

- Abrell, J., M. Kosch, and S. Rausch (2019). How effective was the UK carbon tax? A machine learning approach to policy evaluation. CER-ETH Center of Economic Research at ETH Zurich Working Paper 19/317.
- Ackerley, N., J. Eyraud, C. Franz, B. Kissel, and D. Metivier (2009). Financial conflict-of-interest disclosure and voting patterns at FDA advisory committee meetings. *Washington, DC: Eastern Research Group*.
- Adams, C. P. and V. V. Brantner (2006). Estimating the cost of new drug development: Is it really \$802 million? *Health Affairs* 25(2), 420–428.
- Agha, L. and D. Zeltzer (2019). Drug diffusion through peer networks: The influence of industry payments. National Bureau of Economic Research Working Paper.
- Aghion, P., J. Van Reenen, and L. Zingales (2013). Innovation and institutional ownership. *American Economic Review* 103(1), 277–304.
- Allman, E. S., C. Matias, and J. A. Rhodes (2009). Identifiability of parameters in latent structure models with many observed variables. *The Annals of Statistics*, 3099–3132.
- Amaral-Garcia, S. (2020). Medical device companies and doctors: Do their interactions affect medical treatments? ECARES Working Paper 2020-18.
- Anderson, L. R. and C. A. Holt (1997). Information cascades in the laboratory. *American Economic Review* 87(5), 847–862.
- Angrist, J. D. and J.-S. Pischke (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.
- Antón, M., F. Ederer, M. Gine, and M. C. Schmalz (2018). Innovation: The bright side of common ownership? Working Paper. Available at SSRN 3099578.

- Antón, M., F. Ederer, M. Giné, and M. C. Schmalz (2020). Common ownership, competition, and top management incentives. CESifo Working Paper Series No. 6178.
- Appel, I. R., T. A. Gormley, and D. B. Keim (2016). Passive investors, not passive owners. *Journal of Financial Economics* 121(1), 111–141.
- Appelt, S. (2015). Authorized generic entry prior to patent expiry: Reassessing incentives for independent generic entry. *Review of Economics and Statistics* 97(3), 654–666.
- Asch, S. E. (1951). Effects of group pressure upon the modification and distortion of judgments. In H. Guetzkow (Ed.), *Groups, Leadership, and Men*, pp. 222–236. Pittsburgh: Carnegie Press.
- Associated Press (2006). Contaminated with metal, a painkiller is recalled. *New York Times*. November 10. Online. Available at: <http://www.nytimes.com/2006/11/10/health/10recall.html> (Accessed: 11 November 2019).
- Athey, S. and G. W. Imbens (2019). Machine learning methods that economists should know about. *Annual Review of Economics* 11.
- Azar, J. (2017). Portfolio diversification, market power, and the theory of the firm. Working Paper, IESE Business School, Universidad de Navarra, 2017.
- Azar, J., S. Raina, and M. C. Schmalz (2019). Ultimate ownership and bank competition. Working Paper. Available at SSRN 2710252.
- Azar, J., M. C. Schmalz, and I. Tecu (2018). Anticompetitive effects of common ownership. *The Journal of Finance* 73(4), 1513–1565.
- Backus, M., C. Conlon, and M. Sinkinson (2019). The common ownership hypothesis: Theory and evidence. Brookings Working Paper.
- Banal-Estañol, A., J. Seldeslachts, and X. Vives (2018). Common ownership: Product market consequences of a shift from active to passive investors. Working Paper.
- Banal-Estañol, A., J. Seldeslachts, and X. Vives (2020). Diversification, common ownership, and strategic incentives. In *AEA Papers and Proceedings*, Volume 110, pp. 561–64.

- Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics* 107(3), 797–817.
- Bena, J., M. A. Ferreira, P. Matos, and P. Pires (2017). Are foreign investors locusts? The long-term effects of foreign institutional ownership. *Journal of Financial Economics* 126(1), 122–146.
- Berger, A. N., I. Hasan, and M. Zhou (2010). The effects of focus versus diversification on bank performance: Evidence from Chinese banks. *Journal of Banking & Finance* 34(7), 1417–1435.
- Bernheim, D. B. and C. L. Exley (2015). Understanding conformity: An experimental investigation. Harvard Business School Working Paper.
- Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy* 100(5), 992–1026.
- Boller, L. and F. S. Morton (2020). Testing the theory of common stock ownership. National Bureau of Economic Research Working Paper.
- Boone, A. L. and J. T. White (2015). The effect of institutional ownership on firm transparency and information production. *Journal of Financial Economics* 117(3), 508–533.
- Booraem, G. (2013). Passive investors, not passive owners: A look at Vanguard’s approach to proxy voting and corporate governance. Online. Available at: <https://corpgov.law.harvard.edu/2013/05/10/passive-investors-not-passive-owners/> (Accessed: 11 November 2019).
- Branstetter, L., C. Chatterjee, and M. J. Higgins (2016). Regulation and welfare: Evidence from paragraph IV generic entry in the pharmaceutical industry. *The RAND Journal of Economics* 47(4), 857–890.
- Brav, A., W. Jiang, S. Ma, and X. Tian (2018). How does hedge fund activism reshape corporate innovation? *Journal of Financial Economics* 130(2), 237–264.
- Brav, A., W. Jiang, F. Partnoy, and R. Thomas (2008). Hedge fund activism, corporate governance, and firm performance. *The Journal of Finance* 63(4), 1729–1775.
- Brennan, G. and P. Pettit (2000). The hidden economy of esteem. *Economics and Philosophy* 16, 77–98.

- Brito, D., L. Cabral, and H. Vasconcelos (2016). Competitive effects of partial control in an input supplier. CEPR Discussion Paper No. DP11397.
- Brito, D., R. Ribeiro, and H. Vasconcelos (2014). Measuring unilateral effects in partial horizontal acquisitions. *International Journal of Industrial Organization* 33, 22–36.
- Brown, M. B. (2009). Federal advisory committees in the United States: A survey of the political and administrative landscape. *Scientific Advice to Policy Making: International Comparison*, 17–39.
- Bursztyn, L., F. Ederer, B. Ferman, and N. Yuchtman (2014). Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions. *Econometrica* 82(4), 1273–1301.
- Cai, H., Y. Chen, and H. Fang (2009). Observational learning: Evidence from a randomized natural field experiment. *American Economic Review* 99(3), 864–82.
- Caliendo, M. and S. Kopeinig (2005). Some practical guidance for the implementation of propensity score matching.
- Callander, S. (2007). Bandwagons and momentum in sequential voting. *Review of Economic Studies* 74(3), 653–684.
- Camara, F. and N. Dupuis (2014). Structural estimation of expert strategic bias: The case of movie reviewers. Working Paper.
- Camara, F. and M. Kyle (2016). Experts and financial ties: Evidence from FDA advisory committees. Working Paper.
- Cameron, A. C. and P. K. Trivedi (2013). *Regression analysis of count data*, Volume 53. Cambridge university press.
- Carey, C. M., E. M. Lieber, and S. Miller (2020). Drug firms’ payments and physicians’ prescribing behavior in Medicare part D. National Bureau of Economic Research Working Paper.
- Carleton, W. T., J. M. Nelson, and M. S. Weisbach (1998). The influence of institutions on corporate governance through private negotiations: Evidence from TIAA-CREF. *The Journal of Finance* 53(4), 1335–1362.
- Carrera, M., D. P. Goldman, G. Joyce, and N. Sood (2018). Do physicians respond to the costs and cost-sensitivity of their patients? *American Economic Journal: Economic Policy* 10(1), 113–52.

- Caudill, S. B. et al. (1988). An advantage of the linear probability model over probit or logit. *Oxford Bulletin of Economics and Statistics* 50(4), 425–427.
- Centers for Disease Control and Prevention (2017). National diabetes statistics report, 2017. Atlanta, GA: Centers for Disease Control and Prevention, U.S. Dept of Health and Human Service.
- Centers for Medicare and Medicaid Services (2019). Medicare fee-for service provider utilization and payment data part D prescriber public use file: A methodological overview. The Centers for Medicare and Medicaid Services, Office of Enterprise Data and Analytics.
- Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal* 21, C1–C68.
- Cici, G., S. Gibson, and C. M. Rosenfeld (2015). Cross-company effects of common ownership: Dealings between borrowers and lenders with a common blockholder. Working Paper. Available at SSRN 2705856.
- Cipriani, M. and A. Guarino (2014). Estimating a structural model of herd behavior in financial markets. *American Economic Review* 104(1), 224–251.
- Coffee, J. C. (1991). Liquidity versus control: The institutional investor as corporate monitor. *Columbia law review* 91(6), 1277–1368.
- Cooper, J. C. and J. Golec (2017). Conflicts of interest on expert committees: The case of FDA drug advisory committees. Working Paper.
- Costa-Font, J., A. McGuire, and N. Varol (2014). Price regulation and relative delays in generic drug adoption. *Journal of Health Economics* 38, 1–9.
- Cubanski, J., A. Damico, and T. Neuman (2018). Medicare part D in 2018: The latest on enrollment, premiums, and cost sharing. Kaiser Family Foundation. May 17. Online. Available at: <https://www.kff.org/medicare/issue-brief/medicare-part-d-in-2018-the-latest-on-enrollment-premiums-and-cost-sharing/> (Accessed: 11 November 2019).
- Cubanski, J., T. Neuman, S. True, and A. Damico (2019). How much does Medicare spend on insulin? Kaiser Family Foundation. April 1. Online. Available at: <https://www.kff.org/medicare/issue-brief/how-much-does-medicare-spend-on-insulin/> (Accessed: 11 November 2019).

- Cubanski, J., M. Rae, K. Young, and A. Damico (2019). How does prescription drug spending and use compare across large employer plans, medicare part d, and medicaid. Kaiser Family Foundation. Online. May 20. Available at: <https://www.kff.org/medicare/issue-brief/how-does-prescription-drug-spending-and-use-compare-across-large-employer-plans-medicare-part-d-and-medicaid/> (Accessed: 6 January 2020).
- Cunningham, C., F. Ederer, and S. Ma (2019). Killer acquisitions. Working Paper. Available at SSRN 3241707.
- Danzon, P. M. and M. V. Pauly (2002). Health insurance and the growth in pharmaceutical expenditures. *The Journal of Law and Economics* 45(S2), 587–613.
- Datta, A. and D. Dave (2017). Effects of physician-directed pharmaceutical promotion on prescription behaviors: Longitudinal evidence. *Health Economics* 26(4), 450–468.
- DeJong, C., T. Aguilar, C.-W. Tseng, G. A. Lin, W. J. Boscardin, and R. A. Dudley (2016). Pharmaceutical industry–sponsored meals and physician prescribing patterns for Medicare beneficiaries. *JAMA Internal Medicine* 176(8), 1114–1122.
- Del Guercio, D. and J. Hawkins (1999). The motivation and impact of pension fund activism. *Journal of Financial Economics* 52(3), 293–340.
- Deutsch, M. and H. B. Gerard (1955). A study of normative and informational social influences upon individual judgment. *Journal of Abnormal and Social Psychology* 51(3), 629.
- Donohue, J. M., M. Cevalco, and M. B. Rosenthal (2007). A decade of direct-to-consumer advertising of prescription drugs. *New England Journal of Medicine* 357(7), 673–681.
- Dreiling, M. and D. Darves (2011). Corporate unity in american trade policy: A network analysis of corporate-dyad political action. *American Journal of Sociology* 116(5), 1514–63.
- Duggan, J. and C. Martinelli (2001). A bayesian model of voting in juries. *Games and Economic Behavior* 37(2), 259–294.
- Duggan, M. and F. Scott Morton (2010). The effect of medicare part d on pharmaceutical prices and utilization. *American Economic Review* 100(1), 590–607.

- Duggan, M. and F. M. Scott Morton (2006). The distortionary effects of government procurement: Evidence from Medicaid prescription drug purchasing. *The Quarterly Journal of Economics* 121(1), 1–30.
- Duso, T., L.-H. Röller, and J. Seldeslachts (2014). Collusion through joint R&D: An empirical assessment. *Review of Economics and Statistics* 96(2), 349–370.
- Eagly, A. H. and L. L. Carli (1981). Sex of researchers and sex-typed communications as determinants of sex differences in influenceability: A meta-analysis of social influence studies. *Psychological Bulletin* 90(1), 1.
- EFPIA (2019). The pharmaceutical industry in figures. European Federation of Pharmaceutical Industries and Associations.
- Elhauge, E. (2015). Horizontal shareholding. *Harvard Law Review* 129, 1267.
- Elhauge, E. (2020). How horizontal shareholding harms our economy-and why antitrust law can fix it. *Harvard Business Law Review* 10, 207.
- Elhauge, E. (2021). The causal mechanisms of horizontal shareholding. *Ohio State Law Journal*, forthcoming 82(2).
- Engelberg, J., C. A. Parsons, and N. Tefft (2014). Financial conflicts of interest in medicine. Working Paper. Available at SSRN 2297094.
- Eyster, E. and M. Rabin (2010). Naive herding in rich-information settings. *American Economic Journal: Microeconomics* 2(4), 221–43.
- FDA (2019). Statement from fda commissioner scott gottlieb, m.d., on new policy to improve access and foster price competition for drugs that face inadequate generic competition. Press release. February 19. Online. Available at: <https://www.fda.gov/NewsEvents/Newsroom/PressAnnouncements/ucm631456.htm> (Accessed: 20 February 2019).
- Fernandez, F. and D. Zejcirovic (2018). The role of pharmaceutical promotion to physicians in the opioid epidemic. Working Paper.
- Fichtner, J., E. M. Heemskerk, and J. Garcia-Bernardo (2017). Hidden power of the big three? passive index funds, re-concentration of corporate ownership, and new financial risk. *Business and Politics* 19(2), 298–326.
- Freeman, K. (2019). The effects of common ownership on customer-supplier relationships. *Kelley School of Business Research Paper* (16-84).

- Freeman, L. C. (1977). A set of measures of centrality based on betweenness. *Sociometry*, 35–41.
- Freeman, L. C. (1978). Centrality in social networks conceptual clarification. *Social Networks* 1(3), 215–239.
- Fugh-Berman, A. and S. Ahari (2007). Following the script: how drug reps make friends and influence doctors. *PLoS Medicine* 4(4), e150.
- Gagnon, M.-A. and J. Lexchin (2008). The cost of pushing pills: A new estimate of pharmaceutical promotion expenditures in the united states. *PLoS Medicine* 5(1), e1.
- Geng, H., H. Hau, and S. Lai (2017). Patent success, patent holdup, and the structure of property rights. Working Paper, Sch. Econ. Finance, Vic. Univ. Wellingt.
- Gerakos, J. and J. Xie (2019). Institutional horizontal shareholdings and generic entry in the pharmaceutical industry. (3285161). Tuck School of Business Working Paper.
- Ghosh, A. and H. Morita (2017). Knowledge transfer and partial equity ownership. *The RAND Journal of Economics* 48(4), 1044–1067.
- Gilje, E. P., T. A. Gormley, and D. Levit (2020). Who’s paying attention? measuring common ownership and its impact on managerial incentives. *Journal of Financial Economics* 137(1), 152–178.
- Gilo, D., Y. Moshe, and Y. Spiegel (2006). Partial cross ownership and tacit collusion. *The RAND Journal of Economics* 37(1), 81–99.
- Glowicka, E., S. Lorincz, E. Pesaresi, L. S. Romero, and V. Verouden (2009). Generic entry in prescription medicines in the EU: Main characteristics, determinants and effects. European Commission Report.
- Gorlach, I. and G. Pham-Kanter (2013). Brightening up: the effect of the physician payment sunshine act on existing regulation of pharmaceutical marketing. *The Journal of Law, Medicine & Ethics* 41(1), 315–322.
- Grabowski, H. G. and J. M. Vernon (1992). Brand loyalty, entry, and price competition in pharmaceuticals after the 1984 drug act. *The Journal of Law and Economics* 35(2), 331–350.

- Graham, D. J., D. Campen, R. Hui, M. Spence, C. Cheetham, G. Levy, S. Shoor, and W. A. Ray (2005). Risk of acute myocardial infarction and sudden cardiac death in patients treated with cyclo-oxygenase 2 selective and non-selective non-steroidal anti-inflammatory drugs: nested case-control study. *The Lancet* 365(9458), 475–481.
- Graham, J. R. (1999). Herding among investment newsletters: Theory and evidence. *Journal of Finance* 54(1), 237–268.
- Gramlich, J. and S. Grundl (2017). Testing for competitive effects of common ownership. Finance and Economics Discussion Series 29.
- Grennan, M., K. Myers, A. Swanson, and A. Chatterji (2018). Physician-industry interactions: Persuasion and welfare. National Bureau of Economic Research Working Paper.
- Guo, T., S. Sriram, and P. Manchanda (2018). The effect of information disclosure on industry payments to physicians. *Working Paper*. Available at SSRN 3064769.
- Hams, M. and W. Wilkinson (2013). Vermont shows the way on physician payment sunshine act enforcement. November 7. Online. Available at: <https://communitycatalyst.org/blog/vermont-shows-the-way-on-physician-payment-sunshine-act-pps-a-enforcement> (Accessed: 11 November 2019).
- Handelsblatt Global (2016). The supposed influence of asset managers. Online. May 10. Available at: <https://global.handelsblatt.com/opinion/the-supposed-influence-of-asset-managers-617976> (Accessed: 11 November 2019).
- Hansen, C. and D. Kozbur (2014). Instrumental variables estimation with many weak instruments using regularized jive. *Journal of Econometrics* 182(2), 290–308.
- Hansen, S., M. McMahon, and A. Prat (2018). Transparency and deliberation within the fomc: A computational linguistics approach. *The Quarterly Journal of Economics* 133(2), 801–870.
- Hansen, S., M. McMahon, and C. V. Rivera (2014). Preferences or private assessments on a monetary policy committee? *Journal of Monetary Economics* 67, 16–32.
- Harford, J., D. Jenter, and K. Li (2011). Institutional cross-holdings and their effect on acquisition decisions. *Journal of Financial Economics* 99(1), 27–39.

- He, J. J. and J. Huang (2017). Product market competition in a world of cross-ownership: Evidence from institutional blockholdings. *The Review of Financial Studies* 30(8), 2674–2718.
- Helland, E. and S. A. Seabury (2016). Are settlements in patent litigation collusive? evidence from paragraph IV challenges. National Bureau of Economic Research Working Paper.
- Helland, E. and M. Sovinsky (2019). Do research joint ventures serve a collusive function? CEPR Discussion Paper No. DP13533.
- Hemphill, C. S. and M. Kahan (2019). The strategies of anticompetitive common ownership. *Yale Law Journal* 129, 1392.
- Hoadley, J., J. Cubanski, and T. Neuman (2016). Medicare part D in 2016 and trends over time. Kaiser Family Foundation. September 16. Online. Available at: <https://www.kff.org/report-section/medicare-part-d-in-2016-and-trends-over-time-section-1-part-d-enrollment-and-plan-availability/> (Accessed: 11 November 2019).
- Hochberg, Y. V., A. Ljungqvist, and Y. Lu (2007). Whom you know matters: Venture capital networks and investment performance. *The Journal of Finance* 62(1), 251–301.
- Hoerl, A. E. and R. W. Kennard (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* 12(1), 55–67.
- Holland, P. W. (1986). Statistics and causal inference. *Journal of the American statistical Association* 81(396), 945–960.
- Hong, H., J. D. Kubik, and A. Solomon (2000). Security analysts’ career concerns and herding of earnings forecasts. *The RAND Journal of Economics*, 121–144.
- Hovenkamp, E. and J. Lemus (2018). Delayed entry settlements at the patent office. *International Review of Law and Economics* 54, 30–38.
- Hudson, J. (2000). Generic take-up in the pharmaceutical market following patent expiry: A multi-country study. *International Review of Law and Economics* 20(2), 205–221.
- Hurwitz, M. A. and R. E. Caves (1988). Persuasion or information? promotion and the shares of brand name and generic pharmaceuticals. *The Journal of Law and Economics* 31(2), 299–320.

- Iaryczower, M., X. Shi, and M. Shum (2018). Can words get in the way? the effect of deliberation in collective decision making. *Journal of Political Economy* 126(2), 688–734.
- Iaryczower, M. and M. Shum (2012). The value of information in the court: Get it right, keep it tight. *American Economic Review* 102(1), 202–37.
- Iizuka, T. (2012). Physician agency and adoption of generic pharmaceuticals. *American Economic Review* 102(6), 2826–58.
- Jacobo-Rubio, R., J. L. Turner, and J. W. Williams (2020). The distribution of surplus in the us pharmaceutical industry: Evidence from paragraph IV patent-litigation decisions. *The Journal of Law and Economics* 63(2), 203–238.
- Kennedy, P., D. P. O’Brien, M. Song, and K. Waehrer (2017). The competitive effects of common ownership: Economic foundations and empirical evidence. Working Paper. Available at SSRN 3008331.
- King, G. and R. Nielsen (2019). Why propensity scores should not be used for matching. *Political Analysis* 27(4).
- Knight, B. and N. Schiff (2010). Momentum and social learning in presidential primaries. *Journal of Political Economy* 118(6), 1110–1150.
- Kostovetsky, L. and A. Manconi (2020). Common institutional ownership and diffusion of innovation. Working Paper. Available at SSRN 2896372.
- Kyle, M. K. (2006). The role of firm characteristics in pharmaceutical product launches. *The RAND Journal of Economics* 37(3), 602–618.
- Lakdawalla, D. N. (2018). Economics of the pharmaceutical industry. *Journal of Economic Literature* 56(2), 397–449.
- Li, D. (2017). Expertise versus bias in evaluation: Evidence from the NIH. *American Economic Journal: Applied Economics* 9(2), 60–92.
- Li, D. and L. Agha (2015). Big names or big ideas: Do peer-review panels select the best science proposals? *Science* 348, 434–438.
- Li, X., T. Liu, and L. A. Taylor (2020). Do venture capitalists stifle competition? Jacobs Levy Equity Management Center for Quantitative Financial Research Paper. Available at SSRN 3479439.

- López, Á. L. and X. Vives (2019). Overlapping ownership, R&D spillovers, and antitrust policy. *Journal of Political Economy* 127(5), 2394–2437.
- Lu, F. (2014). Insurance coverage and agency problems in doctor prescriptions: evidence from a field experiment in china. *Journal of Development Economics* 106, 156–167.
- Lundin, D. (2000). Moral hazard in physician prescription behavior. *Journal of Health Economics* 19(5), 639–662.
- Lurie, P. and A. Zieve (2006). Sometimes the silence can be like the thunder: Access to pharmaceutical data at the FDA. *Law and Contemporary Problems* 69(3), 85–97.
- Majumdar, S. (2019). Consequences of oligopoly: Horizontal ownership concentration and telecommunications sector deterred entry. Working Paper.
- Malenko, N. and Y. Shen (2016). The role of proxy advisory firms: Evidence from a regression-discontinuity design. *The Review of Financial Studies* 29(12), 3394–3427.
- Martin, A. B., M. Hartman, B. Washington, A. Catlin, and N. H. E. A. Team (2016). National health spending: faster growth in 2015 as coverage expands and utilization increases. *Health Affairs* 36(1), 166–176.
- Martin, A. B., M. Hartman, B. Washington, A. Catlin, and N. H. E. A. Team (2017). National health spending: faster growth in 2015 as coverage expands and utilization increases. *Health Affairs* 36(1), 166–176.
- McCaffrey, D. F., G. Ridgeway, and A. R. Morral (2004). Propensity score estimation with boosted regression for evaluating causal effects in observational studies. *Psychological Methods* 9(4), 403.
- McCahery, J. A., Z. Sautner, and L. T. Starks (2016). Behind the scenes: The corporate governance preferences of institutional investors. *The Journal of Finance* 71(6), 2905–2932.
- McKelvey, W. and N. H. Kerr (1988). Differences in conformity among friends and strangers. *Psychological Reports* 62(3), 759–762.
- Mizruchi, M. S. and J. Galaskiewicz (1993). Networks of interorganizational relations. *Sociological Methods & Research* 22(1), 46–70.

- Moffitt, S. L. (2010). Promoting agency reputation through public advice: Advisory committee use in the FDA. *The Journal of Politics* 72(3), 880–893.
- Moreno-Torres, I., J. Puig-Junoy, and J.-R. Borrell (2009). Generic entry into the regulated spanish pharmaceutical market. *Review of Industrial Organization* 34(4), 373–388.
- Moretti, E. (2011). Social learning and peer effects in consumption: Evidence from movie sales. *Review of Economic Studies* 78(1), 356–393.
- Morton, F. M. S. (2000). Barriers to entry, brand advertising, and generic entry in the us pharmaceutical industry. *International Journal of Industrial Organization* 18(7), 1085–1104.
- Morton, F. M. S. (2002). Horizontal integration between brand and generic firms in the pharmaceutical industry. *Journal of economics & management strategy* 11(1), 135–168.
- Mullainathan, S. and J. Spiess (2017). Machine learning: an applied econometric approach. *Journal of Economic Perspectives* 31(2), 87–106.
- New York Times (2016). Rise of institutional investors raises questions of collusion. Online. April 12. Available at: <https://www.nytimes.com/2016/04/13/business/dealbook/rise-of-institutional-investors-raisesquestions-of-collusion.html> (Accessed: 11 November 2019).
- Newham, M., J. Seldeslachts, and A. Banal-Estañol (2018). Common ownership and market entry: Evidence from pharmaceutical industry. *DIW Berlin Discussion Paper No. 1738*.
- O’Brien, D. P. and S. C. Salop (1999). Competitive effects of partial ownership: Financial interest and corporate control. *Antitrust Law Journal* 67, 559.
- O’Brien, D. P. and K. Waehrer (2017). The competitive effects of common ownership: We know less than we think. *Antitrust Law Journal* 81, 729.
- OECD (2012). Unilateral disclosure of information with anticompetitive effects. *OECD Report*.
- OECD (2017). Policy roundtable: Common ownership by institutional investors and its impact on competition. Online. Available at: <http://www.oecd.org/daf/competition/common-ownership-and-its-impact-on-competition.htm> CESifo Working Paper Series No. 6178 (Accessed: 1 November 2020).

- Ottaviani, M. and P. Sørensen (2001). Information aggregation in debate: Who should speak first? *Journal of Public Economics* 81(3), 393–421.
- Patel, M. S. (2018). Common ownership, institutional investors, and antitrust. *Antitrust Law Journal* 82(1).
- Pauly, M. V. (1968). The economics of moral hazard: comment. *American Economic Review* 58(3), 531–537.
- Pawliczek, A., A. N. Skinner, and S. L. Zechman (2019). Facilitating tacit collusion: A new perspective on common ownership and voluntary disclosure. Working Paper. Available at SSRN 3382324.
- Perlis, R. H. and C. S. Perlis (2016). Physician payments from industry are associated with greater Medicare part D prescribing costs. *PloS one* 11(5), e0155474.
- Pham-Kanter, G. (2014). Revisiting financial conflicts of interest in FDA advisory committees. *The Milbank Quarterly* 92(3), 446–470.
- Posner, E. A., F. M. Scott Morton, and E. G. Weyl (2017). A proposal to limit the anti-competitive power of institutional investors. *Antitrust Law Journal* 81.
- Prasad, R. (2019). The human cost of insulin in america. *BBC News*. March 14. Online. Available at: <https://www.bbc.com/news/world-us-canada-47491964> (Accessed: 11 November 2019).
- Qian, J., R. A. Hansen, D. Surry, J. Howard, Z. Kiptanui, and I. Harris (2017). Disclosure of industry payments to prescribers: industry payments might be a factor impacting generic drug prescribing. *Pharmacoepidemiology and Drug Safety* 26(7), 819–826.
- Regan, T. L. (2008). Generic entry, price competition, and market segmentation in the prescription drug market. *International Journal of Industrial Organization* 26(4), 930–948.
- Reiffen, D. and M. R. Ward (2005). Generic drug industry dynamics. *Review of Economics and Statistics* 87(1), 37–49.
- Reynolds, R. J. and B. R. Snapp (1986). The competitive effects of partial equity interests and ventures. *International Journal of Industrial Organization* 4(2), 141–153.
- Rock, E. B. and D. L. Rubinfeld (2018). Common ownership and coordinated effects. NYU Law and Economics Research Paper.

- Rosenbaum, P. R. and D. B. Rubin (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika* 70(1), 41–55.
- Rotemberg, J. (1984). Financial transaction costs and industrial performance. Working Paper, Alfred P. Sloan School of Management.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* 66(5), 688.
- Rubin, D. B. (1990). Formal mode of statistical inference for causal effects. *Journal of Statistical Planning and Inference* 25(3), 279–292.
- Rubinstein, A. and M. E. Yaari (1983). The competitive stock market as cartel maker: Some examples. Suntory and Toyota International Centres for Economics and Related Disciplines, LSE.
- Saha, A., H. Grabowski, H. Birnbaum, P. Greenberg, and O. Bizan (2006). Generic competition in the us pharmaceutical industry. *International Journal of the Economics of Business* 13(1), 15–38.
- Samii, C., L. Paler, and S. Z. Daly (2016). Retrospective causal inference with machine learning ensembles: An application to anti-recidivism policies in colombia. *Political Analysis* 24(4), 434–456.
- Sarpawari, A., J. DiBello, M. Zakarian, M. Najafzadeh, and A. S. Kesselheim (2019). Competition and price among brand-name drugs in the same class: A systematic review of the evidence. *PLoS Medicine* 16(7), e1002872.
- Scharfstein, D. S., J. C. Stein, et al. (1990). Herd behavior and investment. *American Economic Review* 80(3), 465–479.
- Schmalz, M. (2015). How passive funds prevent competition. Online. Available at: <http://ericposner.com/martin-schmalz-how-passive-funds-prevent-competition/> (Accessed: 1 November 2020).
- Schmalz, M. C. (2018). Common-ownership concentration and corporate conduct. *Annual Review of Financial Economics* 10, 413–448.
- Schmidt, C. and R. Fahlenbrach (2017). Do exogenous changes in passive institutional ownership affect corporate governance and firm value? *Journal of Financial Economics* 124(2), 285–306.
- Scott Morton, F. M. (1999). Entry decisions in the generic pharmaceutical industry. *The Rand Journal of Economics* 30, 421–440.

- Seldeslachts, J., M. Newham, and A. Banal-Estañol (2017). Common ownership of german companies. *DIW Economic Bulletin* 30.
- Shiller, R. J. (1995). Conversation, information, and herd behavior. *American Economic Review* 85(2), 181–185.
- Shrank, W. H., E. R. Cox, M. A. Fischer, J. Mehta, and N. K. Choudhry (2009). Patients' perceptions of generic medications. *Health Affairs* 28(2), 546–556.
- Shrank, W. H., J. N. Liberman, M. A. Fischer, C. Girdish, T. A. Brennan, and N. K. Choudhry (2011). Physician perceptions about generic drugs. *Annals of Pharmacotherapy* 45(1), 31–38.
- Smith, L. and P. Sørensen (2000). Pathological outcomes of observational learning. *Econometrica* 68(2), 371–398.
- Suzuki, J. (2013). Land use regulation as a barrier to entry: evidence from the texas lodging industry. *International Economic Review* 54(2), 495–523.
- The Economist (2015). Mutual funds and airline competition: Who really owns the skies? Online. July 11. Available at: <https://www.economist.com/blogs/freeexchange/2015/06/mutual-funds-and-airline-competition> (Accessed: 11 November 2019).
- Thomas, K. (2012). An increase in scrutiny for generics. *New York Times*. October 1. Online. Available at: <http://www.nytimes.com/2012/11/01/business/fda-increases-scrutiny-of-some-generic-drugs.html> (Accessed: 11 November 2019).
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* 58(1), 267–288.
- Tikhonov, A. N. (1963). On the solution of ill-posed problems and the method of regularization. In *Doklady Akademii Nauk*, Volume 151, pp. 501–504. Russian Academy of Sciences.
- Urfalino, P. and P. Costa (2015). Secret-public voting in FDA advisory committees. In J. Elster (Ed.), *Secrecy and Publicity in Votes and Debates*, Chapter 8, pp. 165–195. Cambridge: Cambridge University Press.
- Varian, H. R. (2014). Big data: New tricks for econometrics. *Journal of Economic Perspectives* 28(2), 3–28.
- Vitali, S., J. B. Glattfelder, and S. Battiston (2011). The network of global corporate control. *PloS one* 6(10), e25995.

- Wager, S. and S. Athey (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* 113(523), 1228–1242.
- Wang, Y. I., J. Li, and R. Anupindi (2018). Manufacturing and regulatory barriers to generic drug competition: A structural model approach. Working Paper.
- Weizsäcker, G. (2010). Do we follow others when we should? A simple test of rational expectations. *American Economic Review* 100(5), 2340–60.
- Welch, I. (1992). Sequential sales, learning, and cascades. *Journal of Finance* 47(2), 695–732.
- Wiedman, E. (2014). Good information cascades. Working Paper.
- Yala, S. M., O. K. Duru, S. L. Ettner, N. Turk, C. M. Mangione, and A. F. Brown (2014). Patterns of prescription drug expenditures and medication adherence among medicare part D beneficiaries with and without the low-income supplement. *BMC Health Services Research* 14(1), 665.
- Zou, H. and T. Hastie (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Methodological)* 67(2), 301–320.
- Zweifel, P. and F. Breyer (1997). *Health Economics*. New York: Oxford Univ. Press.

**Department of Management,
Strategy & Innovation**
Naamsestraat 69
3000 LEUVEN, Belgium
melissa.newham@kuleuven.be
www.kuleuven.be

