Three Essays on the Economics of Merger Control

vorgelegt von
M.Sc.
Pauline Luise Affeldt

von der Fakultät VII - Wirtschaft und Management
der Technischen Universität Berlin
zur Erlangung des akademischen Grades

Doktorin der Wirtschaftswissenschaften
- Dr. rer. oec. -
genehmigte Dissertation

Promotionsausschuss:

Vorsitzender:	Prof. Dr. Axel Werwatz
Gutachter:	Prof. Dr. Tomaso Duso
Gutachterin:	Prof. Dr. Radosveta Ivanova-Stenzel

Tag der wissenschaftlichen Aussprache: 05. Juli 2019

Berlin, 2019

Abstract

Competition policy is the design and enforcement of competition rules ensuring that companies compete fairly with each other. It is one of the cornerstones of the European Union's program to enhance the European single market and foster growth. Competition policy covers four areas ranging from monitoring and blocking anti-competitive agreements, to abuses by dominant firms, to mergers and acquisitions, and state aid. Among these areas of antitrust enforcement, merger control plays a special role as it is the only area of *ex ante* enforcement. Since 1990, when the European Communities Merger Regulation came into force, all major concentrations must be notified and scrutinized by the Directorate-General for Competition to ensure that consumers are not harmed. This dissertation empirically analyzes the effectiveness of European merger control. First, we study the time-dynamics of the European Commission's merger decision procedure over the first 25 years of European merger control using a new relevant market level dataset containing all merger cases with an official decision documented between 1990 and 2014. Second, we evaluate the predictability of the European Commission's merger decision procedure before and after the 2004 merger policy reform using the highly flexible, non-parametric random forest algorithm to predict competitive concerns in markets affected by a merger. Finally, we focus on one particular market and empirically investigate the impact of multi-homing in two-sided markets using a dataset on the Italian daily newspaper market. Ignoring multi-homing behavior is likely to bias the conclusions of exercises such as market definition or merger evaluation in cases involving multi-sided platforms.

Keywords: competition policy, merger control, merger policy reform, European Union, DG Competition, prediction, machine learning, causal forests, random forests, two-sided markets, newspapers, network effects, platforms, multi-homing, AIDS, logit

Zusammenfassung

Wettbewerbspolitik beinhaltet die Entwicklung und Durchsetzung von Wettbewerbsregeln, um den fairen Wettbewerb zwischen Unternehmen zu gewährleisten. Sie ist zentraler Bestandteil des Programmes der Europäischen Union zur Stärkung des europäischen Binnenmarktes und zur Förderung des wirtschaftlichen Wachstums. Die Wettbewerbspolitik besteht aus vier Bereichen, der Unterbindung koordinierten Verhaltens und des Missbrauchs einer marktbeherrschenden Stellung, der Fusionskontrolle und der Prüfung staatlicher Beihilfen. Innerhalb der verschiedenen Bereiche der Wettbewerbspolitik spielt die Fusionskontrolle eine besondere Rolle, da es sich bei ihr als einzige um eine *ex ante* Durchsetzung handelt. Seit dem Inkrafttreten der Fusionskontrollverordnung im Jahr 1990, müssen alle größeren Unternehmenszusammenschlüsse, welche die Märkte mehrerer europäischer Länder betreffen, bei der Generaldirektion Wettbewerb angemeldet und von dieser geprüft werden, um sicherzustellen, dass Verbraucher durch die Fusion nicht schlechter gestellt werden. Die vorliegende Arbeit analysiert empirisch die Effektivität der europäischen Fusionskontrolle. Zunächst wird die Zeitdynamik der Fusionsverfahren der Europäischen Kommission über die Jahre 1990 bis 2014 auf Basis eines Datensatzes, der alle von der Fusion betroffenen Produkt- und geographischen Märkte erfasst, untersucht. Zweitens wird mit Hilfe von flexiblen, nicht-parametrischen Random Forest Algorithmen die Vorhersagbarkeit der Fusionsentscheidungen vor und nach der Reform der Fusionskontrollverordnung in 2004 analysiert. Mit dem Fokus auf einen konkreten Markt, werden abschließend die Auswirkungen von Multi-Homing in zweiseitigen Märkten unter Verwendung eines Datensatzes von italienischen Tageszeitungen untersucht. Die Ergebnisse zeigen, dass im Falle von mehrseitigen Plattformen eine Nicht-Berücksichtigung von Multi-Homing die wettbewerbliche Beurteilung bezüglich der Marktdefinition und der Wirkung von Fusionen verzerren kann.

Schlüsselwörter: Wettbewerbspolitik, Fusionskontrolle, Reform der Fusionskontrollvereinbarung, Europäische Union, Generaldirektion Wettbewerb, Vorhersage, Maschinelles Lernen, Causal Forest, Random Forest, zweiseitige Märkte, Zeitungen, Netzwerkeffekte, Plattformen, Multi-Homing, AIDS, Logit

Acknowledgements

This dissertation was written during my time as a research associate at Technische Universität Berlin (TU Berlin) and at the German Institute for Economic Research (DIW Berlin). I am lucky to have been part of these two institutions and the vibrant research community in Berlin.

First, and most of all, I thank my supervisors, Tomaso Duso and Radosveta Ivanova-Stenzel, who gave me the opportunity to do independent research, provided their support and guidance throughout, and made this time a unique, challenging, and rewarding experience. I thank Tomaso Duso for all his time, energy, encouragement, and valuable critiques. He constantly pushed me to explore and learn new things and to become more confident in presenting and discussing my work. I am also very grateful to my second supervisor, Radosveta Ivanova-Stenzel, who not only always managed to dispel my doubts and insecurities, but who also gave me the opportunity to test and improve my teaching skills.

I am thankful to my co-authors Elena Argentesi, Lapo Filistrucchi and Florian Szücs for countless discussions about our projects but also research and life in general and for making working on our papers less lonely and more fun. My sincere thanks also goes to Hannes Ullrich who always took the time to give valuable feedback and to discuss methodological or technical details. Further, I want to thank my office mates, Melissa Newham and Kevin Tran, for all the good times mostly in, but also out of, the office together. I thank my colleagues from the Firms and Markets department at DIW Berlin and from the TU Berlin, as well as my graduate center cohorts, both from the Berlin Doctoral Program in Economics and Management Science (BDPEMS) and the DIW Graduate Center.

Finally, and most importantly, I thank my family and friends, for their continued encouragement, love and support throughout my doctoral studies and life in general.

Rechtliche Erklärung

Hiermit versichere ich, dass ich die vorliegende Dissertation selbstständig und ohne unzulässige Hilfsmittel verfasst habe. Die verwendeten Quellen sind vollständig im Literaturverzeichnis angegeben. Die Arbeit wurde noch keiner Prüfungsbehörde in gleicher oder ähnlicher Form vorgelegt.

Berlin, den 26. April 2019

Pauline Luise Affeldt

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 General Introduction

Competition policy is the design and enforcement of competition rules ensuring that businesses and companies compete fairly with each other, i.e. on the basis of their products and prices, with no unfair advantages. Thus, the main goal of competition policy is to promote competition, as competition puts businesses under constant pressure to increase efficiency, offer a wide choice for consumers, reduce prices, and improve quality. When companies try to limit competition, the role of competition authorities is to prevent or correct anti-competitive behavior and preserve the well-functioning and competitiveness of markets to the benefit of consumers.

In light of the growing body of economic research reporting the global rise of concentration, profits, mark-ups, and market power across many markets and industries,[1] the importance and role of competition policy as one tool to prevent abusive behavior and protect competition is currently widely discussed. According to a recent special report on competition by *The Economist*, market concentration has increased in about two-thirds of 900 American industries between 1997 and 2012 and about 10% of the economy consists of industries in which more than two-thirds

[1]For example, Grullon, Larkin, and Michaely (2018) document the broad increases in concentration and profits in over 75% of U.S. industries since the late 1990s. Gutiérrez and Philippon (2017) analyze competition, measured by the Herfindahl-Hirschman-Index (HHI) of concentration, and investment in the U.S. They find that the increase in concentration is mainly driven by a decrease in domestic competition that, in turn, leads to a decrease in firm-level investments. Hartman-Glaser, Lustig, and Xiaolan (2018) and Autor, Dorn, Katz, Patterson, and Van Reenen (2017) focus on the role of large firms. In particular, Autor, Dorn, Katz, Patterson, and Van Reenen (2017) document the growing importance of large firms dominating the market, leading to higher concentration and a decrease in labor's share of GDP in the U.S. and European OECD countries. De Loecker, Eeckhout, and Unger (2018) estimate mark-ups using the ratio of costs of goods sold for the U.S. since 1995. They find that average mark-ups have increased from 21% above marginal costs to 61% since 1980, mostly within industries for all industries. They also discuss the macroeconomic implications of an increase in average market power, notably declining labor and capital shares.

of the market are controlled by only four firms.[2] The newspaper also documents increasing profits in the U.S. and similar though less pronounced trends in Europe, concluding that competition "can help spread wealth by making goods cheaper and reducing the monopsony power that firms can have over workers. It creates wealth by pushing firms to innovate."[3]

When companies try to limit competition, the role of competition policy is precisely to prevent or correct anti-competitive behavior and preserve the well-functioning and competitiveness of markets to the benefit of consumers. For example, Gutiérrez and Philippon (2018) claim that European markets have become more competitive than their U.S. counterpart since the 1990s due to the increased economic integration and the enactment of the European single market. They attribute a key role in this process to the tough enforcement of competition policy rules in the European Union (EU).

Competition policy covers the monitoring and, where necessary, blocking of anticompetitive agreements, abuses of market power by dominant firms, mergers and acquisitions, as well as state aid. Among the different areas of competition policy, this dissertation focuses on the European Commission's (EC) merger control. Merger control plays a particular role among the different areas of antitrust enforcement. First, it is the only area where there is *ex ante* enforcement. Secondly, it also has important implications for the other areas of antitrust: if anticompetitive mergers that reduce competition and strengthen the dominant position of the merging firms are not prohibited, it might make the *ex post* control of abusive behaviors more difficult. Finally, mergers are the area in antitrust where the largest consensus on good practices exists. Therefore, among competition policy tools, it attracts much policy interest and economic research.

There is a large body of both theoretical and empirical literature in the field of industrial organization focusing on questions such as firms' incentives to merge and merger policy effectiveness. Duso, Gugler, and Szücs (2013) identify three dimensions along which merger policy effectiveness can be evaluated: the predictability, correctness, and deterrence effects of a decision. A large part of the literature studying the effectiveness of merger control, looks at whether the competition authority made the correct decision in a particular case (ex-post evaluations of merger policy) (Duso, Neven, and Röller, 2007; Duso, Gugler, and Yurtoglu, 2011; Kwoka, 2013). A correct decision in this context is a decision that achieves the goals set up in the legal framework - in the EU as well as in most other jurisdictions the goal of

[2]The Economist. The Next Capitalist Revolution. November 17, 2018.

[3]The Economist. The Next Capitalist Revolution. November 17, 2018. Special report Competition, page 12.

competition policy is the protection of consumer surplus. A merger that decreases consumer surplus is considered to be anti-competitive. In order to judge whether a particular decision was correct, one, thus, must determine whether a given merger harmed consumer surplus.

The first part of this dissertation (Chapters 2 to 4) focusses instead on the first part of merger policy effectiveness: understanding the determinants of merger decisions and studying its predictability. The goal is to understand how the Directorate-General Competition (DG Comp) decides on interventions in merger cases and whether it is possible to predict DG Comp's decision based on *ex ante* merger and market characteristics. However, these predictions do not allow for judging whether DG Comp's decision was correct in the sense that it protected consumer surplus. While ultimately the correctness of a decision is one of the main aspects of effective merger control, the predictability of decisions based on *ex ante* observable merger characteristics is of interest in its own respect. In particular, prior to the notification of a merger, legal certainty and the predictability of the merger control procedure are important for judges, competition lawyers, and for firms' choices of which kind of mergers to propose. A transparent and predictable process allows firms to better understand the authority's merger review process and, ultimately, predict the outcome of a merger review to a certain extent. Therefore, it should encourage self-compliance: firms should be encouraged to propose pro-competitive mergers and discouraged from proposing anti-competitive mergers (McAfee, 2010).

Chapter 2 documents the database used in Chapters 3 and 4. Chapter 3 studies the time-dynamics of the EC's merger decision procedure over the first 25 years of European merger control (1990-2014) and finds that while concentration as well as the merging parties' market shares have become less important decision determinants over time, barriers to entry as well as the risk of foreclosure are increasingly important to DG Comp's merger assessment since the early 2000s. This is in line with the goals of the 2004 merger policy reform, which aimed at adopting a more economics based approach of merger assessment and at putting less weight on simple structural indicators. Chapter 4 studies the predictability of DG Comp's merger policy and assesses how it changed following this reform. It shows that, even though DG Comp seems to base its assessment on a more complex interaction of merger and market characteristics post-reform, the highly flexible random forest algorithm is able to detect these potentially complex interactions and, therefore, still allows for a high prediction precision also post-reform.

The second part of this dissertation (Chapter 5) leaves the macro perspective of evaluating EU merger control over the last 25 years at the aggregate and instead focuses on one example of a particular market. Specifically, the last chapter empir-

ically studies the impact of multi-homing on price elasticities in two-sided markets.

Two-sided markets are markets in which firms sell two products or services to two different types of consumers taking into account that the two demands are linked by indirect network effects (Evans, 2003; Rochet and Tirole, 2003). The classical example of such a two-sided market is the newspaper market, where the demand for advertising is related to the number of readers and readers might like, dislike, or be indifferent to advertising in newspapers. However, especially in the growing digital economy, many markets are two-sided or multi-sided platform markets, characterized by indirect network externalities between the different groups of consumers.

The correct assessment of own- and cross-price elasticities in these platform markets, taking into account indirect network effects, is relevant as they are important inputs into market definition, merger assessment, and the assessment of market power. Furthermore, multi-homing, i.e. the use of more than one platform, is widespread, especially in online markets. Chapter 5 shows that ignoring multi-homing consumer behavior is likely to bias the conclusions of exercises like market definition or merger evaluation in antitrust cases involving multi-sided platforms. Thus, this part of the dissertation contributes to the current discussion about whether, and if so how, the antitrust toolkit might need to be re-designed or re-interpreted in order to equip competition agencies with the analytical tools they require to analyze multi-sided markets (in the digital economy). The importance of this debate is reflected in the ongoing activities of competition authorities and organizations worldwide to identify key digital challenges and their implications for competition policy. At the EC, Margrethe Vestager, Commissioner for Competition, appointed three Special Advisers from outside the Commission, Professors Heike Schweitzer, Jacques Crémer, and Assistant Professor Yves-Alexandre de Montjoye.[4] On April 4, 2019, the EC published the report on the future challenges of digitization for competition policy that these advisers had worked on for over one year (Crémer, De Montjoye, and Schweitzer, 2019). Similar initiatives have been taken by the German Federal Cartel Office (Bundeskartellamt), the OECD, and the United States' Federal Trade Commission (FTC).[5] According to Chris Pike from the

[4]See `https://ec.europa.eu/commission/commissioners/2014-2019/vestager/announcements/commission-appoints-professors-heike-schweitzer-jacques-cremer-and-assistant-professor-yves_en`. Last accessed on March 15, 2019.

[5]For example, the German Federal Cartel Office established its "Task Force for Internet Platforms" in 2015 (See `https://www.bundeskartellamt.de/SharedDocs/Meldung/EN/Meldungen%20News%20Karussell/2015/21_12_2015_Jahresr%C3%BCckblick.html`. Last accessed March 15, 2019.). The German Monopolies Commission (Monopolkommission) published a special report entitled "Competition Policy: The challenge of digital markets" in 2015 (Monopolkommission, 2015). In June 2017, the OECD held a Competition Commission Hearing looking at whether the tools traditionally used to define markets, to assess market power and efficiencies, and to assess the effects of exclusionary conduct and vertical restraints, remain suf-

OECD, *[t]he speed and extent of growth in the digital economy ... has made this one of the most important, pressing, and analytical challenges that competition agencies now face* (OECD, 2018, p.9).

The exercise of market definition illustrates the analytical challenges competition authorities face when dealing with multi-sided platforms. In a merger review process, the first step is the definition of the relevant (product and geographic) market(s). Market definition is a tool to identify and define the boundaries of competition between firms. The goal of market definition is to *"identify in a systematic way the competitive constraints that the undertakings involved face"*.[6] Therefore, it is a way to think about consumer demand and the relevant competitors that constrain the merging parties' behavior. Of course, market definition is not an end in itself but a first step in order to assess competitive constraints, market power, and the effects of the behavior or the merger under review.

Elasticities are an important input into market definition as the cross-price elasticities of demand between the merging parties' products reflect the size of the competitive constraint that is lost due to the merger while the own-price elasticity of demand helps to assess the degree of market power a particular product holds. Consequently, it is crucial to get them right. Traditional tools for market definition, such as the SSNIP test (Small but Significant Non-Transitory Increase in Price),[7] are designed for single-sided markets and cannot easily be applied to two-sided or

ficient to address those questions in the context of multi-sided markets. Following its Hearing, the OECD invited and published practical methodological proposals from a range of expert economists (OECD, 2018). In the fall of 2018 and spring of 2019, the FTC hosted Hearings on Competition and Consumer Protection in the 21st Century, examining whether the changes in the economy might require adjustments to competition and consumer protection law (See `https://www.ftc.gov/policy/hearings-competition-consumer-protection` for an overview of the hearings. Last accessed March 15, 2019.). A new FTC Technology Task Force to monitor competition in technology markets was launched in 2019 (See the Press Release of February 26, 2019: `https://www.ftc.gov/news-events/press-releases/2019/02/ftcs-bureau-competition-launches-task-force-monitor-technology`. Last accessed March 15, 2019.).

[6]Commission Notice on the definition of the relevant market for the purposes of Community competition law (97/C 372/03) [Official Journal C 372 of 9 December 1997]. In particular, a relevant product market *"comprises all those products and/or services which are regarded as interchangeable or substitutable by the consumer, by reason of the products' characteristics, their prices and their intended use."* A relevant geographic market is defined as *"the area in which the undertakings concerned are involved in the supply and demand of products or services, in which the conditions of competition are sufficiently homogeneous and which can be distinguished from neighbouring areas because the conditions of competition are appreciably different in those areas."*

[7]In particular, the SSNIP test asks whether a hypothetical monopolist of the product under consideration would find it profitable to permanently increase the price above the current level (by 5% to 10%). If this is the case, then the product does not face significant competitive constraints from other products and the relevant product market includes only this one product. However, if the price increase is not profitable for the hypothetical monopolist, then the next closest substitute product is considered and the question is asked again. If a small but significant, non-transitory price increase is profitable for the hypothetical monopolist selling these two products, then there is a relevant market. See e.g. Motta (2004).

multi-sided markets (Noel and Evans, 2005; Filistrucchi, Geradin, Van Damme, and Affeldt, 2014). In particular, correct market definition in a two-sided or multi-sided market needs to account for the interdependencies between quantities and prices on all sides and all feedback effects. Failing to correctly account for the two-sidedness of the market can lead to an erroneous market definition.[8] For example, assume that readers actually like newspaper advertising, while a newspaper is also more valuable for advertisers as the number of readers it reaches increases. An increase in advertising rates of one newspaper will, initially, decrease the amount of advertising in that newspaper. However, keeping the cover price fixed, the newspaper is then also less valuable for readers (as they like advertising). Consequently, fewer readers will buy the newspaper, which subsequently results in fewer advertisers and so on. This implies that an increase in advertising rates is less profitable than what it would seem if the indirect network externalities between the two sides are ignored. Consequently, the relevant market might be defined too narrowly (overestimating the profitability of price increases).

Besides the indirect network elasticities, it is also important to take multi-homing behavior into account when assessing competition in multi-sided markets. If, for example, newspaper readers single-home, the competitive bottleneck problem of Armstrong (2006) arises, whereby each newspaper is a monopolist over providing access to its exclusive readers, meaning that advertisers must patronize all platforms in order to reach all readers. However, if a fraction of readers patronizes more than one newspaper, the model predictions change quite dramatically. Now, advertisers can reach multi-homing readers on more than one platform. Therefore, newspapers no longer only compete for consumers on the reader side of the market but now also compete for advertisers on the advertising side of the market. This has important implications for platforms' strategies in terms of pricing, reactions to mergers, and content provision. Of course, this also matters for market definition, as a high degree of multi-homing by one group of consumers may indicate a relatively low degree of competition for these consumers, while a high degree of single-homing might indicate that platforms compete intensely for these consumers.

For example, Wismer and Rasek (2018) discuss the relevance of multi-homing for market definition. In particular, multi-homing might be interpreted as evidence of user switching their demand between platforms thereby implying strong substitutability and close competition between platforms. On the other hand, multi-homing can also indicate that consumers use different platforms in parallel to satisfy

[8] According to Dewenter, Heimeshoff, and Löw (2017), there is no quantitative method available that is a suitable, practical (i.e. not too data demanding) tool for market definition in platform markets. Dewenter, Heimeshoff, and Löw (2017) try to fill this gap by identifying competitors in two-sided markets based on time-series methods and simple correlation analysis.

different needs - which would imply that the services or products offered by the platforms might be viewed as complements rather than substitutes on at least one side of the market. Therefore, single-homing and multi-homing behavior can be relevant for market definition. If the rationale for multi-homing is that products are viewed as complements rather than substitutes, multi-homing behavior might actually justify more narrowly defined markets. This is in line with the findings of Chapter 5.

1.2 Outline of the Dissertation

1.2.1 Chapter 2: EU Merger Control Database: 1990-2014

In Chapter 2, which is joint work with Tomaso Duso and Florian Szücs, we document the database that we constructed based on almost the complete population of DG Comp's merger decisions between 1990 and 2014.

In particular, we document the data collection, data cleaning, and quality control procedures. We further describe all the merger and market characteristics contained in the final dataset in detail. Specifically, next to the identity of the merging parties, the type of decision, the notification date, and the decision date, the database also contains information on the type of merger, the geographic market definition, the product market definition, competitors, market shares and concentration measures, the type of competitive concerns and remedies, as well as sector information. Rather than taking a particular merger case as the level of observation, we decided to collect data at a finer level, defining an observation as a particular product/geographic market combination concerned by a merger. In total, the final dataset contains 5,196 DG Comp merger decisions, with 31,451 relevant market level observations.

1.2.2 Chapter 3: 25 Years of European Merger Control

In Chapter 3, which is joint work with Tomaso Duso and Florian Szücs, we study the time-dynamics of the EC's merger decision procedure over the first 25 years of European merger control using the relevant market level dataset containing all merger cases with an official decision documented by DG Comp between 1990 and 2014 that is described in Chapter 2. Specifically, we evaluate how consistently different arguments related to the structural market parameters – market shares, concentration, likelihood of entry, and foreclosure – put forward to motivate a particular decision are applied over time.

In a first step, we estimate the probability of intervention as a function of merger characteristics at the merger level. We find that the existence of barriers to entry, the increase of concentration measures, and, in particular, the share of product markets with competitive concerns increase the likelihood of an intervention. In order to obtain a more fine-grained picture of the decision determinants, we extend our analysis to the specific product and geographic markets concerned by a merger. We find that more determinants significantly affect the Commission's competitive concerns at the market level than we see at the merger level. Again, barriers to entry, but also the risk of foreclosure, play an important role for the competitive analysis. Moreover, while tightly defined (national) markets increase the probability of concerns, the number of active competitors decreases it. Finally, structural indicators of market shares and concentration have the expected effects, which are more relevant than in the merger-level analysis.

After this static investigation, we investigate how the impact of these key determinants changes over time. We generally find that the importance of market shares and concentration seems to have declined over time. However, the parametric estimations are quite volatile and do not allow for uncovering clear patterns over time. In a final step, we use the non-parametric causal forest algorithm proposed by Athey and Imbens (2016), to more precisely explore how the correlation between the structural market parameters and competitive concerns varies with all other merger and market characteristics. We find that concentration as well as the merging parties' market shares have become less important decision determinants over time and are even insignificant in most recent years. On the other hand, the importance of barriers to entry as well as risk of foreclosure has increased over time in DG Comp's merger assessment since the early 2000s.

1.2.3 Chapter 4: EU Merger Policy Predictability Using Random Forests

In Chapter 4, I study the predictability of DG Comp's merger policy and assess how it changed following the 2004 merger reform based on the comprehensive dataset covering almost all mergers notified to the EC between 1990 and 2014 described in Chapter 2.

One goal of the 2004 EU merger reform was to bring merger control closer to economic principles. Another was to increase legal certainty and transparency of the merger review process as evidenced by the publication of merger guidelines and the institutional changes made. However, the effect of the reform on the predictability of DG Comp's decisions is ambiguous, as the use of a "more economic approach"

in the merger review implies a shift from simple general rules, such as concentration thresholds, toward a more in depth case-by-case economic analysis. Thus, the question is whether the merger reform increased the *ex ante* predictability of decisions based on market and merger characteristics as well as how the merger reform changed the decision criteria on which DG Comp bases its merger assessment.

Rather than assessing mergers at the aggregate level, I define an observation as a particular product and geographic market combination concerned by a merger, as in Chapter 3. This allows studying the factors that cause competitive concerns in specific sub-markets. In addition, and unlike the existing literature studying the determinants of DG Comp's merger intervention decisions and their predictability, I use non-parametric random forests to predict DG Comp's assessment of competitive concerns arising in affected markets due to the merger. This machine learning algorithm is designed to maximize predictive performance rather than estimating causal effects and allows for highly flexible, non-linear interactions between covariates.

Using the random forest algorithm to predict DG Comp's assessment of competitive concerns in markets affected by a merger, I find that the predictive performance of the random forests is much better than the performance of simple linear models. In particular, the random forests do much better in predicting the rare event of competitive concerns. Secondly, post-reform, DG Comp seems to base its assessment on a more complex interaction of merger and market characteristics than pre-reform. The highly flexible random forest algorithm is able to detect these potentially complex interactions and, therefore, still allows for high prediction precision.

1.2.4 Chapter 5: Estimating Demand with Multi-Homing in Two-Sided Markets

In Chapter 5, which is joint work with Elena Argentesi and Lapo Filistrucchi, we leave the macro perspective of evaluating EU merger control at the aggregate across decisions and zoom into one particular market. Here, we empirically investigate the impact of multi-homing in two-sided markets. We first build a micro-founded structural econometric model, which encompasses the demand for differentiated products on both sides of the market and allows for multi-homing on each side of the market. We then use an original dataset on the Italian daily newspaper market that includes information on double-readership of newspapers to estimate demand alternatively taking into account and not taking into account information on multi-homing by readers.

In particular, on the readers' side of the market, demand derives from random utility maximization by readers and is estimated using a nested logit model, as in

Berry (1994). When information on multi-homing by readers is ignored, readers choose the newspaper that maximizes their utility. When taking into account information on multi-homing by readers, readers are allowed to choose between all possible pairs of newspapers. On the advertisers' side of the market, demand derives from advertisers' choice to allocate a given advertising budget, which changes with the business cycle, across different newspapers. We use a linear approximation of the Almost Ideal Demand System by Deaton and Muellbauer (1980) to model newspaper level advertising demand. Product differentiation is interpreted in the spatial sense proposed by Pinkse, Slade, and Brett (2002). Distance metrics are derived from differences among newspapers in the demographic characteristics of readers.

The results show that an econometric model that does not allow for multi-homing is likely to produce biased estimates of own- and cross-price elasticities on both the reader side and the advertising side of the market. In particular, mean own-price elasticities on the reader side increase when readers' multi-homing behavior is taken into account. Furthermore, while newspapers are assumed to be substitutes in the single-homing model, they can be substitutes or complements when multi-homing by readers is taken into account. We find that, while newspapers of the same type (general interest, sports, business) are substitutes, newspapers of different types are complements. We also show that, on the advertising side of the market, own-price elasticities decrease with the number of captive readers while cross-price elasticities increase with the number of overlapping readers between newspapers.

The chapter contributes to the economic literature on two-sided markets, in which empirical work accounting for multi-homing is still quite scarce. Moreover, our contribution allows a better understanding of how multi-homing by users in platform markets matters and how it influences price elasticities on both sides of the market. This is likely to bias the conclusions of such exercises as market definition or merger evaluation in which both own- and cross-price elasticities and own- and cross-network effect elasticities play a crucial role. Although print newspapers are a classical example of an *offline* two-sided market, the empirical part of this chapter should be seen more as an application that allows for studying the role of multi-homing in platform markets. Especially in light of the prevalence and rising importance of multi-sided platforms in digital markets and the relevance of multi-homing by users, the results and conclusions from this chapter are also relevant in the context of competition policy cases involving *online* multi-sided platform markets.

Chapter 2

EU Merger Control Database: 1990-2014 [1]

2.1 Introduction

Competition policy, i.e. the design and enforcement of competition rules, is a cornerstone of European Union policy designed to enhance European integration and foster growth. Among the different areas of the European Commission's (EC) antitrust enforcement, i.e. collusion, merger, and abuse-of-dominance cases, this dataset focuses on EC merger policy. As common European merger control started in 1990, we can now look back at, and evaluate more than 25 years of EC merger control.

We collected data on almost the complete population of the Directorate-General Competition's (DG Comp) merger decisions, both across time and with regard to the scope of the decisions encompassed. We started data collection with the very first year of common European merger control, 1990, and included all years up to 2014. This amounts to 25 years of data on European merger control.

With regard to the scope of the decisions, we collected data in all cases where a legal decision document exists. This includes all cases settled in the first phase of an investigation (Art. 6(1)(a), 6(1)(b), 6(1)(c) and 6(2)) and all cases decided in the second phase of an investigation (Art. 8(1), 8(2), and 8(3)). Note that this also includes all cases settled under a "simplified procedure", provided that a legal decision document exists.

Furthermore, we also intended to collect data on cases that were either referred back to member states by DG Comp or aborted by the merging parties. While we

[1] This chapter is the accepted manuscript published in the DIW Data Documentation Series as: Affeldt, P., Duso, T. and F. Szücs (2018). EU Merger Control Database: 1990-2014. *DIW Data Documentation Series 95*. We thank Ivan Mitkov, Fabian Braesemann, David Heine, Juri Simons and Isabel Stockton for their precious research assistance.

have collected some data on such cases, data on these cases is not always available. Therefore, we cannot guarantee that the final dataset covers all of these cases.

Rather than taking a particular merger case as the level of observation, we decided to collect data at a more fine-grained level, defining an observation as a particular product/geographic market combination concerned by a merger.

In total, the final dataset contains 5,196 DG Comp merger decisions, where each decision occupies a number of rows equal to the number of product/geographic markets identified in the specific transaction. Hence, the total dataset contains 31,451 observations.

The remainder of the data documentation is structured as follows. In Section 2.2, we provide a short overview of DG Comp's merger review process. In Section 2.3, we describe how we collected and recorded the merger data, in Section 2.4, we describe our data cleaning and quality control procedure. Section 2.5 contains a description of all the variables included in the final database. Lastly, we explain the data collection procedure with the help of an example case in Section 2.6.

2.2 EU Merger Review Process

Mergers that affect the European market must be notified to the EC when involving an EU community-wide dimension.

DG Comp then has 25 working days (which can be extended to a maximum of 35 working days) for an initial assessment of the merger. This is the so-called "phase-1 investigation." Based on this initial assessment DG Comp can clear the proposed merger (phase-1 clearance), clear it subject to remedies proposed by the merging parties (phase-1 remedy), or initiate a more in-depth investigation (phase-2 investigation) depending on whether the proposed transaction raises competitive concerns and depending on whether these can be addressed by initial remedies or not. Furthermore, the merging parties might also withdraw the proposed merger during phase-1 (phase-1 withdrawal).

If DG Comp initiates a more in depth investigation, this phase-2 investigation can take up to 90 working days. Following this second investigation phase, DG Comp can again unconditionally clear the merger (phase-2 clearance), clear the merger subject to commitments by the merging parties (phase-2 remedy), or prohibit the merger (phase-2 prohibition). Again, the merging parties can also still withdraw the proposed merger during phase-2 (phase-2 withdrawal). It has been argued that withdrawing a merger during phase-2 of the investigation process is virtually equivalent to a prohibition as parties often withdraw a merger before an actual prohibition by DG Comp takes place. Hence, both a prohibition as well as a phase-2 with-

drawal suggest that DG Comp and the notifying parties were unable to find suitable remedies to address the anti-competitive concerns of the proposed transaction.

2.3 Data Collection Procedure

All decisions by DG Comp are available and publicly accessible on the EC's website.[2] We downloaded all available merger decision documents for merger cases notified to the EC between 1990 and end of 2014.

These decision documents were then partly read and scanned for the relevant information that we wanted to collect in the appropriate sections of the decisions. For example, the recording of a particular case will typically start with the basic case information (number, dates, decision etc.) contained on the first page(s) of the document. The typical structure of a decision document is as follows:

- **Introduction**: The case is summarized on the first pages of the document. The final decision as well as the relevant dates and parties involved are also stated here.

- **The Parties, The Operation, Concentration of Community Dimension**: This section of the decision discusses the merging parties as well as the nature of the merger proposal in detail. Under the heading "Concentration and Community Dimension" DG Comp justifies why the case has an EU-wide dimension.

- **Compatibility with the Common Market**: This section is the main part of the decision and contains most information that we collected. The sections "Relevant Product Markets" and "Relevant Geographical Markets" explain in detail which markets and products are affected by the merger. The next section (called "Assessment" or similar) typically contains the market shares of the merging parties as well as of competitors in each concerned product/geographic market. The section "Competitive Assessment" contains the discussion of the potential competitive concerns of the merger in all relevant product/geographic markets. We filter out some of the characteristics of the concerned markets (see Section 2.5 for a description of the included variables).

[2]The types of notified mergers, decisions taken, and reports for each of DG Comp's decisions are available at: `http://ec.europa.eu/competition/mergers/cases/`; `http://ec.europa.eu/competition/mergers/legislation/simplified_procedure.html`.

- **Undertakings proposed by the Parties or Parties proposed Remedy**: This section of the decision contains the description of the remedies that the merging parties proposed in order to address the competitive concerns raised by DG Comp, distinguishing between behavioral and structural remedies.

- **Assessment of the proposed Modifications** This section contains DG Comp's evaluation of the appropriateness of the proposed remedies in alleviating the competitive concerns raised previously.

- **Overall Conclusion**: This section contains the final decision of DG Comp. Hence, it states whether the proposed merger is compatible with the common market or whether it would significantly impede competition in the common market and, consequently, is going to be prohibited.

- **Appendix**: The final assessment by DG Comp is typically followed by numerous appendices containing tables and figures highlighting certain aspects of the decision. These are not typically relevant for the type of information we collected.

During the data collection process, we recorded all the information gathered from the decision documents in Microsoft Excel tables. The format of these tables was uniform across all research assistants involved in the data gathering process, thus facilitating merging them later.

We then merged the individual data tables into a single matrix using the statistical software package STATA. This facilitated various tasks of cross-checking the data, quality control (see Section 2.4) and will also be helpful in the creation of standardized classification schemes. The cleaned and standardized dataset can then be exported back into any data format desired.

To date, data on almost all merger cases decided by DG Comp from 1990 through 2014, inclusive, has been collected. However, there are about 500 decision documents between 1990 and 2014 for which data is not yet recorded, primarily because most of these documents are not in English.

Given that we consider all merger cases notified to the EC between 1990 and 2014, some of these cases (around 50) were decided only in 2015.

2.4 Data Cleaning & Quality Control

In order to ensure a high quality and consistency of the data collected, we essentially took two measures.

First, we established a uniform data collection procedure for all research assistants going through the decision documents and recording the data. Secondly, we controlled the quality of the data once we imported the raw data from the Excel tables into STATA.

The first step is particularly crucial: we developed an approach to analyzing DG Comp's decision documents that i) makes it clear to the individual research assistant what information is to be collected from the decisions; ii) where in the decision documents this information can be found (or is most likely to be found); and iii) how these tasks can best be streamlined. To this end, we developed a "manual" that explains in detail how the data are to be collected. Furthermore, at the beginning of the data collection stage, we asked each research assistant to re-collect data on a few mergers that were already reliably recorded. This allowed us to compare the "canonical" data to the results delivered by the research assistant. Any discrepancies between the two were discussed with the research assistant, such that human mistakes or ambiguities in the data collection procedure could be ruled out to the largest extent possible.

Of course, human error cannot entirely be ruled out. That is why we conducted a second stage of quality control. While typos and other human errors are hard to spot in tables with thousands of rows and dozens of columns, the statistical evaluation of the resulting tables once imported into STATA made this consistency check easily possible. Thus, in the second stage of quality control we checked for typos in the data, unreasonably large or small values in specific variables, and missing data problems.

We corrected, for example, typos, coding errors, and missing values in the basic information about the decision (see Section 2.5 for a detailed description of the variables). Some case numbers and country information were corrected. Furthermore, we checked whether the notification date was always prior to the decision date, which allowed for spotting typos in the date variables. At times the outcome of a decision was also wrongly coded in the Excel files. We further corrected coding errors or missing values in the indicator variables describing the type of the merger as well as the geographic market concerned. Lastly, we harmonized merging party names across markets and imputed some missing market share information. In cases where the correct values of variables were not obvious, we went back to the respective decision documents in order to correct the data.

Following the data cleaning, the final dataset contains 31,451 observations belonging to 5,196 merger cases.

2.5 Database Content

This section describes in detail the information contained in the final merger database. As explained above, the unit of observation is not a particular merger case but rather a particular product/geographic market combination affected by the merger. Hence, some of the variables collected vary at the merger level while others vary at the level of the concerned product/geographic market combination. The overview table in Appendix 2.7.1 lists all variables contained in the database and specifies whether they vary at the merger or the product/geographic market level.

2.5.1 Basic Information about the Decision

The dataset contains first some basic information about the decision. The variable *casen* contains the case number as reported in the decision document. This variable uniquely identifies each merger case. The variables *notdate* and *decdate* contain the date of the notification to, and the date of the decision of DG Comp, respectively. We also included the variables *notyear* and *decyear* containing the year in which the notification respectively the decision took place.

We also collected information on acquiring and target firms. In some merger cases more than one acquiring and/or more than one target firm are involved. This is why the dataset contains information on up to three acquiring and up to two target firms. The string variables *acquirer1*, *acquirer2*, *acquirer3*, *target1* and *target2* contain the names of the acquiring firms as well as of the target firms. Tables 2.17 and 2.18 in Appendix 2.7.2 and 2.7.3 list the top 20 primary acquiring and target firms respectively. Note however that this is a preliminary assessment of acquiring and target firms before complete name harmonization.

The variables *countryacq1*, *countryacq2*, *countryacq3*, *countrytar1*, and *countrytar2* record the nationality of the acquiring and the target firms respectively. Table 2.19 in Appendix 2.7.4 lists the top 20 acquiring and target firms' countries based on the primary acquiring and target firm respectively. If the notified merger is a joint venture, the parties are ordered into acquirer and target according to the order the companies appear in the title of the decision.

The variable *outcome* contains the type of decision made by DG Comp distinguishing phase-1 clearances (outcome 1 "ph1 clear"), phase-1 clearances subject to remedies (outcome 2 "ph1 rem"), phase-2 clearances (outcome 3 "ph2 clear"), phase-

2 clearances subject to remedies (outcome 4 "ph2 rem"), prohibitions (outcome 5 "prohibition"), phase-1 withdrawals (outcome 6 "ph1 withdrawal"), phase-2 withdrawals (outcome 7 "ph2 withdrawal"), referrals back to the competition authority of the respective member state (outcome 8 "referral to MS"), as well as other types of decision documents (outcome 9 "other").

Phase-1 cases are decided under Art.6(1)(a), Art.6(1)(b), or Art.6(2) of the EC Merger Regulation. While phase-1 clearances are cases that are decided under Art.6(1)(a) or Art.6(1)(b) without imposing remedies, phase-1 clearances subject to remedies are cases decided under Art.6(1)(b) or Art.6(2) with imposition of remedies.

Phase-2 cases are decided under Art.8(1), Art.8(2), or Art.8(3) of the EC Merger Regulation. While phase-2 clearances are decided under Art.8(1) or Art.8(2) without imposing remedies, phase-2 clearances subject to remedies are decided under Art.8(2) with imposition of remedies. Prohibitions are decided under Art.8(3).

Cases that are referred back to national competition authorities are decided either under Art.4(4) or Art.9(3). Lastly, all other cases were included in the outcome category "other." These cases contain, for example, cases decided under Art.14 (fines for supplying incorrect or incomplete information or for putting into effect a concentration), Art.7(3) (derogation from suspension obligation imposed under 7(1)), or Art.22 (where a member state asks the EC to treat a specific merger case).

Table 2.1 reports the number of phase-1 clearances, phase-1 remedies, phase-2 clearances, phase-2 remedies, prohibitions, withdrawals, referrals to member states, and other decisions. Out of the 5,196 merger cases included in the database, about 95% of the cases are either cleared or cleared subject to remedies in phase-1. Only in about 3.5% of the merger cases does DG Comp initiate an in depth phase-2 investigation. The table also shows that once a phase-2 investigation is initiated, an unconditional clearance is rather unlikely. In five merger cases, the merging parties withdrew the transaction during the phase-2 investigation. As discussed in Section 2.2, withdrawing a merger in phase-2 of the investigation process could be regarded as equivalent to a prohibition since parties often withdraw a merger before an actual prohibition by DG Comp takes place.

In 69 merger cases (which corresponds to 406 product/geographic market observations in the dataset), the case is referred back to the national competition authority of the member state. "Other" comprises 16 decision documents, as discussed above.

Lastly, the database also contains the variable *simplified*. This indicator variable is equal to one if the case was settled under a "simplified procedure". Since 2000, the EC has introduced "simplified procedures" for those merger notifications that are very likely to be pro-competitive in nature, i.e. that do not raise competitive

Table 2.1: Type of Decisions, 1990-2014

Type of decision	frequency	percent
Phase-1 clearance	4,691	90.28
Phase-1 remedy	239	4.60
Phase-2 clearance	51	0.98
Phase-2 remedy	104	2.00
Prohibition	19	0.37
Phase-1 withdrawal	2	0.04
Phase-2 withdrawal	5	0.10
Referral to MS	69	1.33
Other	16	0.31
Total	5,196	100.00

concerns. In particular, conglomerate mergers, horizontal mergers with joint market shares below 20% and vertical mergers where the notifying parties have less than 30% market share in upstream and downstream markets are notified under these procedures. Information on whether a particular case was settled under simplified procedures can be downloaded from the EC's website and combined with our dataset via the case number.

Table 2.2 summarizes this variable by type of decision for the years 2000-2014. Since its introduction, 52% of the merger cases have been notified under simplified procedures. All of these cases have been decided in phase-1, almost entirely as phase-1 clearances.

Table 2.2: Indicator Variable for Simplified Procedure by Decision Type, 2000-2014

Type of decision	0	1	mean	standard deviation
Phase-1 clearance	1,628	2,221	0.58	0.494
Phase-1 remedy	189	1	0.01	0.073
Phase-2 clearance	36	0	0.00	0.000
Phase-2 remedy	74	0	0.00	0.000
Prohibition	10	0	0.00	0.000
Phase-1 withdrawal	0	2	1.00	0.000
Phase-2 withdrawal	5	0	0.00	0.000
Referral to MS	63	0	0.00	0.000
Other	13	1	0.07	0.267
Total	2,018	2,225	0.52	0.499

All of the variables containing basic information about the decision vary at the merger level.

Figure 2.1 shows the yearly number of merger notifications, phase-1 merger cases, mergers cleared subject to remedies (phase-1 and phase-2) and prohibitions between 1990 and 2014. Overall, merger notifications show an increasing trend with a big drop around 2002. Most of the notified mergers are decided in phase-1: Phase-1 mergers track the number of notifications very closely. The number of mergers cleared subject to remedies increased dramatically after 1996 and oscillates between 20 and 30 per year in more recent years. The number of prohibitions vary between zero and three prohibitions per year. Table 2.20 in Appendix 2.7.5 shows the number of notifications and decisions per year.

Figure 2.1: Enforcement History of DG Comp Merger Cases, 1990-2014

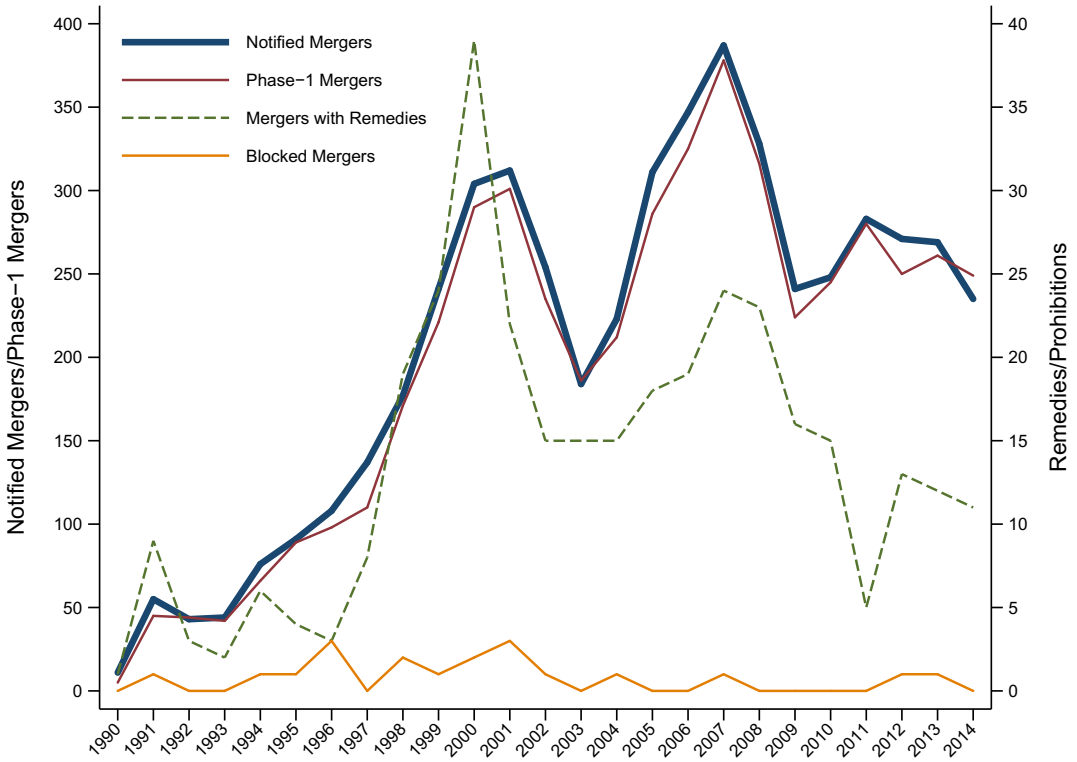

We report notified cases per notification year and phase-1 cases per decision year (left axis) as well as remedies (phase-1 and phase-2) and prohibitions per decision year (right axis). We exclude all cases where the decision type is "other".

2.5.2 Type of Merger

The dataset additionally contains some information about the nature of the merger.

The variable *vertical* is a dummy variable equal to one if product/geographic markets are vertically affected by the merger and zero otherwise. The variable *conglomerate* is a dummy variable that is equal to one if the merger is conglomerate in nature. In addition, we recorded whether DG Comp considered the merger to be a full merger and/or a joint venture. This information is stored in the dummy variables *fullmerger* and *jv* respectively.

While the variables *vertical* and *conglomerate* are market specific (and hence can vary within a particular merger case), the variables *fullmerger* and *jv* vary at the merger level.

While 8,421 product/geographic markets were affected vertically by the respective merger (corresponding to 27% of observations), mergers had conglomerate aspects in only 525 (about 2% of observations) of the affected markets (see Table 2.3).

Table 2.3: Indicator Variables for Vertical and Conglomerate Merger, 1990-2014

	0	1	mean	standard deviation
Conglomerate	30,926	525	0.017	0.128
Vertical	23,030	8,421	0.268	0.443

Out of the 5,196 mergers, 2,872 (55%) are full mergers and 1,908 (37%) are joint ventures (see Table 2.4).

Note also that the variables *fullmerger* and *jv* are not mutually exclusive. If DG Comp considers the merger to be a full merger, the firms merge in such a way that the target is completely controlled by the acquiring firm. If the merger is a joint venture, the two firms merge only for a particular purpose e.g. by founding a R&D joint-venture. If both variables are equal to zero, the firms merge but the acquiring firm does not fully control the target firm. These cases are partial mergers, in most cases acquisitions of shares.

2.5.3 Market Definition

As previously explained, the unit of observation in the merger database is a particular market concerned by the decision. A market is defined as a combination of a product and a geographic market. We recorded a number of variables that describe the particular market.

Table 2.4: Indicator Variables for Full Merger and Joint Venture, 1990-2014

	0	1	mean	standard deviation
Full merger	2,324	2,872	0.55	0.497
Joint Venture	3,288	1,908	0.37	0.482

The variable *broadmarket* is a variable that we created in order to make different product markets comparable across decisions. It provides a more standardized description of the product market and contains about 460 broad product markets. We further harmonized these broad product markets into 86 product market categories. Table 2.21 in Appendix 2.7.6 reports the number of notifications, phase-1 and phase-2 observations for these 86 product market categories. Many observations concern air transport and travel, banking, financial services and insurance, chemicals, communication services, energy supply, food and beverages, as well as pharmaceuticals.

The variable *prodmarket* is a string variable that contains the exact product market as specified in the decision document.

The variables *national*, *euwide*, *ww*, and *open* are dummy variables referring to the geographic market definition of DG Comp. The variables *national*, *euwide*, and *ww* are equal to one whenever the geographic market is considered to be national, EU wide, or worldwide, respectively. If DG Comp considered an exact definition of the geographic market unnecessary, the variable *open* is equal to one. The string variable *geogmarket* contains the actual verbal description DG Comp used to indicate the geographic market in the decision document.

Table 2.5 shows that DG Comp considers the market to be national in almost 60%, EU wide in about 20%, and worldwide in about 9% of the product/geographic markets. In 12% of the cases, DG Comp left the geographic market definition open.

Table 2.5: Geographic Market Definition, 1990-2014

	0	1	mean	standard deviation
National	13,004	18,447	0.59	0.492
EU wide	25,194	6,257	0.20	0.399
Worldwide	28,490	2,961	0.09	0.292
Left open	27,666	3,785	0.12	0.325

Table 2.6 reports the geographic market definition by type of decision.[3] While in phase-1 clearance cases the geographic market definition is often left open, mergers that are either prohibited or only cleared subject to remedies tend to affect narrow (i.e. national) geographic markets. Also note that in cases that were referred back to national competition authorities (outcome "Referral to MS"), the geographic market is evidently either defined as national or the geographic market definition is left open.

Table 2.6: Mean Geographic Market Definition by Decision Type, 1990-2014

Type of decision	National	EU wide	Worldwide	Left open
Phase-1 clearance	0.33	0.17	0.07	0.43
Phase-1 remedy	0.64	0.24	0.08	0.04
Phase-2 clearance	0.35	0.31	0.27	0.07
Phase-2 remedy	0.58	0.31	0.09	0.02
Prohibition	0.56	0.11	0.24	0.09
Phase-1 withdrawal	0.00	0.00	0.00	1.00
Phase-2 withdrawal	0.00	0.00	0.00	1.00
Referral to MS	0.99	0.00	0.00	0.01
Other	0.44	0.13	0.13	0.31

We take the mean of the geographic market definition indicator variables to collapse the information from market level to merger level.

The geographic market definition can also vary across product/geographic markets within a given merger case. This is the case in 1,064 of the merger cases (about 20% of the cases contained in the database).

2.5.4 Classification of Remedies

The dataset also includes some information about the nature of remedies proposed by the merging parties.

While the dummy variable *remedies* is equal to one whenever the merging parties proposed any remedies to address DG Comp's competitive concerns, the dummy variables *structural* and *behavioral* are indicator variables for whether structural (i.e. divestitures) and/or behavioral remedies were proposed. We do not distinguish whether a remedy affects only a particular market or not, hence the variables related to proposed remedies all vary at the merger level. As it is often difficult to assess whether a particular measure, for example a certain divestiture, affects one or several concerned markets, we prefer to define the remedy variables at the merger level.

[3] We first collapse the dataset from market to merger level by taking the mean of the geographic market indicator variables by merger case. We then report the mean market definition across all mergers included in the database.

In about 7% of the merger cases, remedies were proposed by the notifying parties. As DG Comp prefers structural to behavioral remedies, it is not surprising that in 5% of the cases structural remedies were proposed while behavioral remedies were proposed in only 3.5% of the merger cases (see Table 2.7).

Note also that the variables *remedies*, *structural*, and *behavioral* are equal to one whenever the decision document contains information about remedies proposed by the merging parties. This implies that even for a merger that was prohibited by DG Comp, the variable *remedies* can be equal to one. This is the case whenever the merging parties proposed remedies but DG Comp considered these insufficient to address the competitive concerns and, thus, ultimately prohibited the merger.

Table 2.7: Indicator Variables for Proposed Remedies, 1990-2014

	0	1	mean	standard deviation
Remedies	4,845	351	0.068	0.251
Behavioural remedies	5,016	180	0.035	0.183
Structural remedies	4,931	265	0.051	0.220

2.5.5 Competitive Concerns

Related to proposed remedies, we also included an indicator variable *concern* in the dataset that is a dummy variable indicating which specific product/geographic market affected by the merger (granted that the merger concerned multiple product markets) raised concerns on part of DG Comp.

The indicator variable *barriers* is equal to one if DG Comp considered barriers to entry to exist in the concerned market (hence, this variable varies at the market level). Similarly, *foreclosure* is an indicator for whether DG Comp raised concerns that the merger would foreclose other firms in a particular market.

Table 2.8: Indicator Variables for Competitive Concerns, 1990-2014

	0	1	mean	standard deviation
Concerns	27,769	3,682	0.117	0.322
Entry barriers	27,830	3,621	0.115	0.319
Risk of foreclosure	30,614	837	0.027	0.161

Table 2.8 summarized the information regarding competitive concerns. While DG Comp raised competitive concerns and considered entry barriers to exist in about 12% of the affected markets, it found a risk that the merger would foreclose competitors in only about 3% of the markets.

2.5.6 Competitors

In addition to the names of the acquiring and the target firm, we also included the names of competitors of the merging parties identified by DG Comp, in so far as such information is contained in the decision document. The identity and number of competitors varies by product/geographic market concerned. We hence record the identity of between 0 and 15 competitors (stored in the variables *rival1* to *rival15*).

In a few cases, DG Comp identifies more than 15 competitors of the merging parties. Given that this is the case for very few mergers and that competitors are typically very small in these cases, we considered the informational gain from keeping the identity of more than 15 competitors small compared to the increased unhandiness of a dataset containing many string variables.

Table 2.9: Number of Competitors, 1990-2014

Number of competitors	frequency	percent
0	17,671	56.19
1	1,909	6.07
2	2,746	8.73
3	3,514	11.17
4	2,183	6.94
5	1,468	4.67
6	732	2.33
7	461	1.47
8	286	0.91
9	136	0.43
10	117	0.37
>10	228	0.72
Total	31,451	100.00

Zero competitors means that there is no information on competitors in the decision document. This is either the case if the merger is a merger to monopoly or DG Comp does not mention competitor names in the decision document.

The database also contains the variables *compcount*, which is a count of the number of competitors in the concerned market, and *misscomp*, an indicator variable equal to one if no information on competitors is available. We coded the variable *compcount* as equal to zero whenever we have no information on competitors. In

these cases, the indicator *misscomp* is equal to one. Both variables vary at the market level. Missing information on competitors can have two reasons, either the merging parties have 100% market share in a given market or there is just no information on competitors in the decision document.

As Table 2.9 shows, there is no information on competitors in about 56% of the markets. In about 38% of the product/geographic market observations, we have information on between one and five competitors. Information on more than five competitors is very scarce.

Table 2.10: Indicator Variable for Missing Competitor Information by Decision Type, 1990-2014

Type of decision	No information available	Information available	mean	standard deviation
Phase-1 decision	16,124	11,546	0.58	0.493
Phase-2 decision	1,140	2,187	0.34	0.475
Referral to MS	363	43	0.89	0.308
Other	44	4	0.92	0.279
Total	17,671	13,780	0.56	0.496

Table 2.10 reports the number of product/geographic markets without information on competitors (variable *misscomp* is equal to one) by type of decisions. Phase-1 cases comprise phase-1 clearances, phase-1 remedies, and phase-1 withdrawals, while phase-2 cases are phase-2 clearances, phase-2 remedies, prohibitions, and phase-2 withdrawals. The table highlights that information on competitors is mostly missing in phase-1 case documents: in 58% of the phase-1 case observations no information on competitors is available while this is only the case for 34% of the phase-2 product/geographic market observations.

Table 2.11 reports instead the mean number of competitors for notifications, phase-1, and phase-2 decisions.[4] There is no information on the number of competitors in about 63% of notified mergers and 64% of phase-1 decisions. However, it is much more likely that DG Comp investigates the competitors in detail in a phase-2 investigation. Thus, there is no information on competitors in only about 10% of phase-2 decisions, while in about 85% of phase-2 decisions there is information on between one and five competitors.

[4]We collapse the dataset from market to merger level by taking the mean number of competitors rounded to the nearest integer by merger case.

Table 2.11: Mean Number of Competitors, 1990-2014

Number of competitors	Number of notifications	Percent of notifications	Number of phase-1 decisions	Percent of phase-1 decisions	Number of phase-2 decisions	Percent of phase-2 decisions
0	3,259	62.7	3,169	64.3	17	9.5
1	539	10.4	505	10.2	30	16.8
2	475	9.1	428	8.7	43	24.0
3	405	7.8	356	7.2	48	26.8
4	258	5.0	240	4.9	18	10.1
5	121	2.3	109	2.2	11	6.1
6	60	1.2	57	1.2	2	1.1
7	38	0.7	33	0.7	5	2.8
8	19	0.4	17	0.3	2	1.1
9	6	0.1	5	0.1	1	0.6
10	11	0.2	8	0.2	2	1.1
>10	5	0.1	5	0.1	0	0.0
Total	5,196	100.0	4,932	100.1	179	100.0

We take the mean number of competitors rounded to the nearest integer to collapse the information from market level to merger level. Note that phase-1 and phase-2 decisions do not add up to the number of notifications due to the 69 referrals to Member States and the 16 cases classified as "other".

2.5.7 Market Shares

We collected data on the market shares of the merging parties as well as the competitors, where available. This information was collected from DG Comp's competitive assessment in the decision document. Thus, data availability is constrained by the extent of DG Comp's analysis.

Given that DG Comp generally reports only the range of the market shares in the publicly available documents, we defined the market shares to be equal to the central value of the interval (see Section 2.6 for an illustration).[5]

Market share information is collected at the level of the relevant product/geographic market combination, hence, in cases concerning multiple product/geographic markets, we collected market shares of the merging parties and the competitors for each individual market concerned whenever this information is available.

The market shares of the merging parties are stored in the variables *acq1ms*, *acq2ms*, *acq3ms*, *tar1ms*, and *tar2ms* for acquiring firms 1 to 3 and target firms 1 and 2, respectively, while the variable *Sum* contains the sum of the market shares of the merging parties in percent. In some cases, the decision document only contains

[5]If, for example, the market share range indicated is [0-10] percent, we record a market share of 5 percent. However, if the interval given in the decision is only 5 percentage points wide, we report the conservative lower market share bound. If for example the market share interval is [15-20] percent, we report 15 percent market share.

information on the sum of the merging parties' market shares but not on individual market shares. Competitors' market shares (in percent) are contained in the variables *riv1ms* to *riv15ms* if available.

Table 2.12 shows summary statistics for the market shares of the merging firms as well as competitors. The average market share of the primary acquiring firm is about 20%, the average market share of the primary target is about 18%, and the average joint market share of the merging parties is about 33%. However, there is large variability in the data as the high standard deviations show. The table also reports the market shares of the second and third acquiring firm as well as of the second target firm. These secondary merging parties are in general much smaller: the mean market shares of these firms lie only between 5% and 8%. The mean market share of the first competitor is relatively high, at an average of 25%. Competitors' market shares decrease as the number of competitors increases: The average market share of the second competitor is about 14%, while the average market share of competitor 15 is only about 2%.

Table 2.12 also reports the number of non-missing observations in the column labelled "observations." As this column shows, market share information is relatively scarce: While information on the joint market share of the merging parties is available in 23,136 out of 31,451 markets (hence in about 74% of the markets), information on at least one competitor's market share is available in only about 33% of the markets. The last column labelled "cases" counts the number of merger cases for which the respective market share information is available in at least one of the concerned product/geographic market combinations. Information on primary acquirer's and primary target's market shares is available in about 1,600 out of the 5,196 merger cases.

2.5.8 Concentration Measures

We calculated the level of the post-merger Herfindahl-Hirschman-Index (HHI) in case that data on the market shares of competitors was available (variables *hhi_low* and *hhi_high* ranging from 0 to 10,000).

The variable *hhi_low* is a lower bound of the post-merger HHI: it is calculated as the square of the merging parties joint markets share plus the sum of squared market shares of competitors whenever information on competitors' market shares is available. This assumes that competitors are very small, whenever market share information of competitors is not available but market shares do not add up to 100%. The variable *hhi_high*, on the other hand, is an upper bound for the post-merger HHI: it adds the square of all missing market shares (100% minus all available

Table 2.12: Summary Statistics Market Shares and HHI

	mean	sd	min	max	observations	cases
Acquirer 1 market share	19.7	20.84	0	100	13,683	1,576
Acquirer 2 market share	8.2	15.17	0	100	893	181
Acquirer 3 market share	5.3	8.81	0	30	11	6
Target 1 market share	17.5	21.04	0	100	13,701	1,585
Target 2 market share	7.8	15.10	0	100	385	76
Joint market share	32.6	23.65	0	100	23,136	2,468
Competitor 1 market share	24.8	16.34	0	100	10,354	1,645
Competitor 2 market share	14.1	9.76	0	100	8,468	1,532
Competitor 3 market share	9.7	7.55	0	95	5,988	1,323
Competitor 4 market share	7.5	6.14	0	93	3,210	949
Competitor 5 market share	6.4	5.81	0	65	1,798	605
Competitor 6 market share	5.7	6.22	0	85	957	348
Competitor 7 market share	4.9	6.15	0	95	551	191
Competitor 8 market share	5.4	6.12	0	45	330	111
Competitor 9 market share	4.6	5.26	0	45	202	70
Competitor 10 market share	4.7	5.62	0	35	139	49
Competitor 11 market share	4.1	5.91	0	45	102	34
Competitor 12 market share	3.6	3.97	0	20	78	21
Competitor 13 market share	4.2	6.64	0	35	64	17
Competitor 14 market share	2.4	3.03	0	15	45	13
Competitor 15 market share	2.0	4.34	0	25	42	11
Post-merger HHI (lower bound)	2,156.2	2,371.89	0	10,000	23,136	2,468
Post-merger HHI (upper bound)	5,643.0	2,242.93	650	10,000	23,136	2,468
Delta HHI	443.9	778.83	0	8,450	12,957	1,467

market share information) to *hhi_low*. This hence treats all missing market share information as one missing competitor.

From the merging parties' market shares, we also calculated the increase in HHI due to the merger in the specific markets, stored in the variable *deltahhi*. In case of one acquiring and one target firm, it is calculated as $2 \cdot acq1ms \cdot tar1ms$.[6] As the market share information is specific to a certain product/geographic market combination, the concentration measures also vary at the market level.

Summary statistics for *hhi_low*, *hhi_high*, and *deltahhi* are also contained in Table 2.12. The mean post-merger HHI is between 2,156 (lower bound) and 5,643 (upper bound), while the mean increase in HHI due to the merger is about 440.

2.5.9 Complexity

The variable *complexity* contains a count of the relevant product/geographic markets concerned by the merger. Hence, it varies at the merger level.

[6]We distinguish cases with one acquirer and one target, two acquirers and one target, three acquirers and one target, one acquirer and two targets, two acquirers and two targets, and three acquirers and two targets. In a case involving, for example, two acquiring and one target firm, the change in HHI is calculated as $2 \cdot acq1ms \cdot acq2ms + 2 \cdot acq1ms \cdot tar1ms + 2 \cdot acq2ms \cdot tar1ms$. The change for the other cases is calculated accordingly.

The merger cases included in the database concern on average 6 product/geographic market combinations, varying between a minimum of 1 and 245 concerned markets (see Table 2.13).

Table 2.13: Summary Statistics Complexity

	mean	sd	min	max
Number of markets	6.05	13.37	1	245
Observations	5,196			

2.5.10 Sector Information

Lastly, we include information on which NACE sector(s) are concerned by the proposed merger. NACE codes are an industry classification system used by the European Union to classify different economic activities.[7] Information on the main NACE sectors concerned by the mergers can be downloaded from the EC's website and combined with the dataset via the case number.

Merger cases can concern multiple NACE sectors. The dataset contains all NACE codes reported on the EC's website (dropping duplicate NACE codes).[8] They are stored in the variables *nace1* to *nace15*. Table 2.14 reports the number of merger cases with information on no up to 15 NACE codes, distinguishing phase-1 and phase-2 cases as well as referrals to member states and other decision documents. For 3,894 out of the 5,196 cases, one NACE code is reported. Note that for 140 cases there is no information on the NACE code. Most of these cases are phase-1 cases. Only in a few cases are more than three NACE codes reported.

Table 2.14: Number of NACE Codes by Decision Type, 1990-2014

Type of decision	No NACE code	1	2	3	4	5	6	7	8	9	11	15
Phase-1 decision	107	3,715	742	235	76	30	19	3	1	2	1	1
Phase-2 decision	2	138	25	6	5	2	0	0	1	0	0	0
Referral to MS	21	35	8	2	1	1	1	0	0	0	0	0
Other	10	6	0	0	0	0	0	0	0	0	0	0
Total	140	3,894	775	243	82	33	20	3	2	2	1	1

[7] See `http://ec.europa.eu/competition/mergers/cases/index/nace_all.html` for a list of NACE codes.

[8] Following our question on whether an allocation of NACE codes to the merging parties is possible, the merger registry informed us, that the order in which NACE codes are reported is random and that NACE codes cannot be allocated to acquiring and target firms.

Table 2.15 reports the number of notifications, phase-1, and phase-2 decisions by primary NACE section (the most aggregate classification level). By far the most merger cases with 2,257 out of 5,196 cases concern mergers in the manufacturing industry, followed by wholesale and retail trade (487 cases), information and communication (478 cases), and financial and insurance activities (477 cases).

Note that phase-1 and phase-2 decisions do not always add up to the number of notifications within a given NACE section due to the 69 referrals to member states and the 16 cases classified as "other".

Table 2.15: Decisions by Primary NACE Section, 1990-2014

NACE section	Description	Notifications	Phase-1 decisions	Phase-2 decisions
A	Agriculture, forestry and fishing	38	34	3
B	Mining and quarrying	135	125	8
C	Manufacturing	2,257	2,143	103
D	Electricity, gas, steam and air conditioning supply	281	265	10
E	Water supply; sewerage; waste managment and remediation activities	63	62	0
F	Construction	90	87	1
G	Wholesale and retail trade; repair of motor vehicles and motorcycles	487	470	7
H	Transporting and storage	326	313	7
I	Accommodation and food service activities	65	63	1
J	Information and communication	478	444	27
K	Financial and insurance activities	477	475	2
L	Real estate activities	87	87	0
M	Professional, scientific and technical activities	60	58	2
N	Administrative and support service activities	105	100	4
O	Public administration and defence; compulsory social security	22	22	0
P	Education	4	4	0
Q	Human health and social work activities	27	21	0
R	Arts, entertainment and recreation	38	36	2
S	Other services activities	14	14	0
T	Activities of households as employers; undifferentiated goods - and services - producing activities of households for own use	2	2	0
Missing		140	107	2
Total		5,196	4,932	179

Note that phase-1 and phase-2 decisions do not add up to the number of notifications due to the 69 referrals to Member States and the 16 cases classified as "other".

2.6 Case Example

In the following, the assessment of different characteristics concerning EU-merger decisions is explained with the help of one sample case, illustrating many of the different core and non-core elements that are potentially relevant for all (non-simplified) cases. The case example is the case number 623 Kimberley-Clark/Scott, an Art. 8(2) decision.

Most of the variables described above are collected by skimming the merger decisions and transcribing the main information concerning the characteristics of the merger firstly into an Excel spreadsheet. In the following, the collection is hence explained in a step-by-step procedure. Note, again, that the level of observation are product/geographic market combinations, thus for each case, the database contains as many observations (rows) as analyzed markets. This implies that some general information about the merger (e.g., the notification date) is the same for each product market involved by the merger and, therefore, it appears in all rows of a decision. In the case of the merger between Kimberley-Clark and Scott, three product markets were concerned by the transaction, hence there are three observations for this merger case.

Figure 2.2 shows the basic information for the merger decision. Besides the case number *casen* that serves as an identifier, the type of decision and the notification and decision dates are collected. The type of decision is assigned either to the variable *decision* - if it is decided according to Article 6(1)(b) or 6(1)(c) during phase-1 - or to *decision2* - if the case under investigation is decided according to article 8(1), 8(2) or 8(3) during phase-2. The variable *notifdat* captures the notification date and *phase1dat* and *phase2dat* the decision dates of phase-1 and phase-2 cases, respectively.

Figure 2.2: Basic Case Information - 1

B	C	D	E	F	G	H
casen	Decision		Decision Date			
casen	decision	decision2	notifdat	phase1dat	phase2dat	announcement
623		8,2	08.08.95		16.01.96	
623		8,2	08.08.95		16.01.96	
623		8,2	08.08.95		16.01.96	

The information about the merging companies is captured by means of three variables for each of the parties as illustrated in Figure 2.3. While the variables *acquirer1* and *countryacq1* report the acquirer's name and country, the variable *acq1ms* indi-

cates its market share in the respective market. Similarly, the information on the company to be acquired is stored in the variables *target1*, *countrytar1*, and *tar1ms*. In some cases, more than two parties are involved (mostly in the case of joint ventures); for these cases additional columns are provided. The variable *Sum* displays the sum of the acquirer's and target's market share after the merger in the specific product market.

Figure 2.3: Basic Case Information - 2

	R	S	T	U	V	W	X	Y	Z	AA
1	Parties									
2	acquirer 1	country acq 1	acq 1 ms	acquirer 2	country acq 2	acq 2 ms	target 1	country tar1	tar1 ms	Sum
3	KIMBERLY-CLARK (KC)	USA	15				SCOTT PAPER	USA	35	50
4	KIMBERLY-CLARK (KC)	USA	5				SCOTT PAPER	USA	15	20
5	KIMBERLY-CLARK (KC)	USA	35				SCOTT PAPER	USA	5	40

Next, data on the outcome of DG Comp's investigation is collected. The variables shown in Figure 2.4 deal with the implemented remedies, the theory of harm, and the type of merger proposed. The three variables *remedies*, *structural*, and *behavioral* capture the remedies proposed and discussed by DG Comp. In this case, both structural and behavioral remedies were proposed by the merging parties; hence, all tree variables are equal to one.

The variables on the theory of harm include the indicators for barriers of entry, foreclosure, conglomerate concerns, or whether the merger includes a vertical component. In the merger between Kimberley-Clark and Scott, DG Comp raised concerns about barriers of entry.

Figure 2.4: Merger Characteristics

Merger Characteristics								
remedies	structural	behavioral	conglomerate	foreclosure	vertical	barriers	fullmerger	jv
1	1	1	0	0	0	1	1	0
1	1	1	0	0	0	1	1	0
1	1	1	0	0	0	1	1	0

Lastly, the announced concentration between the parties can either be described as a full merger between the companies (*fullmerger* = 1), a joint venture (*jv* = 1) or a non-full merger (i.e. the acquirer buys only parts of the target: *fullmerger* and (*jv* = 0). In this particular case, the transaction between Kimberley-Clark and Scott is a full merger.

Figure 2.5 illustrates the systematic assessment of the product and geographic market for the case in the Excel spreadsheet. In the decision document, a detailed description of the relevant product and geographic market is provided. Further, the

decision contains a competitive assessment in which the relevant market shares of the merging parties and the main competitors are provided for each product market.

In order to make the different product markets comparable across decisions, the variable broad market provides a more standardized description of the product market. In case 623, Kimberley-Clark/Scott, the product markets "toilet paper," "kitchen paper," and "handkerchiefs" can all be summarized under the broader term "paper products." This broader definition allows identifying connections to other cases of the same industry or value chain.

Figure 2.5: Market Definition

	AE	AF	AG	AH	AI	AJ	AK
1	Market Dimension						
2	broad market	prod.market	national	eu-wide	ww	open	geog.market
3	Paper products	toilet paper	1	0	0		United Kingdom and Ireland
4	Paper products	kitchen paper	1	0	0		United Kingdom and Ireland
5	Paper products	handkerchiefs	1	0	0		United Kingdom and Ireland

In addition to the product market, the geographic market is captured by a number of variables. The indicator variables *national*, *eu-wide*, *ww*, and *open* indicate whether the geographic scope of a product market is national, EU-wide, worldwide, or whether there is no geographic market definition provided in the decision. To allow for a more precise geographic market definition, the variable *geog.market* names the precise geographic market definition used in the decision. In case 623, Kimberley-Clark/Scott, the market of UK and Ireland is perceived as one interrelated market. Thus, the market definition is national but comprises two countries. Hence, using the detailed description of the market in *geog.market*, one could also classify this market as cross-border/regional.

Lastly, Figure 2.6 reports the information on competitors in case 623. In this particular case, the decision document contains information on three competitors, including market shares.

Figure 2.6: Competitors

rival1	riv1 ms	rival2	riv2 ms	rival3	riv3 ms	rival4	riv4 ms
Jamont	5	Fort Sterling	5	SCA/PWA	5		
Jamont	5	Fort Sterling	5	SCA/PWA	5		
Jamont	5	Fort Sterling	5	SCA/PWA			

2.7 Appendix

2.7.1 List of Variables

Table 2.16: List of Variables Contained in Database

Variable	Type	Description	Varies at level
casen	numeric	Case number as reported in decision document.	merger
notdate	date	Date of notification to DG COMP.	merger
decdate	date	Date of decision by DG COMP.	merger
notyear	numeric	Year of notification to DG COMP.	merger
decyear	numeric	Year of decision by DG COMP.	merger
outcome	string	Type of decision.	merger
simplified	dummy	1 if decided under simplified procedure.	merger
acquirer1	string	Name of primary acquiring firm.	merger
countryacq1	string	Nationality of primary acquiring firm.	merger
acquirer2	string	Name of second acquiring firm.	merger
countryacq2	string	Nationality of second acquiring firm.	merger
acquirer3	string	Name of third acquiring firm.	merger
countryacq3	string	Nationality of third acquiring firm.	merger
target1	string	Name of primary target firm.	merger
countrytar1	string	Nationality of primary target firm.	merger
target2	string	Name of second target firm.	merger
countrytar2	string	Nationality of second target firm.	merger
vertical	dummy	1 if market is vertically affected.	market
conglomerate	dummy	1 if merger has conglomerate aspects in market.	market
fullmerger	dummy	1 if merger is a full merger.	merger
jv	dummy	1 if merger is a joint venture.	merger
broadmarket	string	Standardized description of product market.	market
prodmarket	string	Product market as specified in the decision document.	market
national	dummy	1 if geographic market is defined as national.	market
euwide	dummy	1 if geographic market is defined as EU.	market
ww	dummy	1 if geogaphic market is defined as worldwide.	market
open	dummy	1 if geographic market definition is left open.	market
geogmarket	string	Geographic market as specified in decision document.	market
remedies	dummy	1 if remedies were proposed by the parties.	merger
structural	dummy	1 if structural remedies were proposed by the parties.	merger
behavioral	dummy	1 if behavioral remedies were proposed by the parties.	merger
concern	dummy	1 if DG COMP raised competitive concerns in market.	market
barriers	dummy	1 if DG COMP finds entry barriers in market.	market
foreclosure	dummy	1 if DG COMP finds risk of foreclosure in market.	market
rival1-rival15	string	Names of competitors in market.	market
compcount	numeric	Number of competitors in market.	market
misscomp	dummy	1 if no information on competitors contained in decision document.	market
acq1ms-acq3ms	numeric	Market shares of acquiring firms in percent.	market
tar1ms-tar2ms	numeric	Market shares of target firms in percent.	market
Sum	numeric	Joint market share of merging parties in percent.	market
riv1ms-riv15ms	numeric	Market shares of competitors in market.	market
hhi_low	numeric	Post-merger HHI in market (lower bound).	market
hhi_high	numeric	Post-merger HHI in market (upper bound).	market
deltahhi	numeric	Delta HHI in market due to merger.	market
complexity	numeric	Count of the number of concerned markets.	merger
nace1-nace10	string	NACE codes reported in decision documents.	merger

2.7.2 Top 20 Primary Acquiring Firms

Table 2.17: Top 20 Primary Acquiring Firms, 1990-2014

Primary acquirer	Number of cases
ADVENT INTERNATIONAL CORPORATION	24
GENERAL ELECTRIC COMPANY	21
DEUTSCHE BANK AG	17
GOLDMAN SACHS GROUP, INC.	14
VOLKSWAGEN AG	13
ELECTRICITé DE FRANCE	12
GENERAL ELECTRIC	12
UNITED TECHNOLOGIES CORPORATION	12
3I GROUP PLC	11
CVC CAPITAL PARTNERS SICAV-FIS S.A.	11
PAI PARTNERS S.A.S.	11
SIEMENS AG	11
THE CARLYLE GROUP	11
BERTELSMANN AG	10
DEUTSCHE BANK	10
DEUTSCHE POST AG	10
KKR	
& CO. L.P.	10
MITSUBISHI CORPORATION	10
SIEMENS	10
THOMSON-CSF	10

2.7.3 Top 20 Primary Target Firms

Table 2.18: Top 20 Primary Target Firms, 1990-2014

Primary target	Number of cases
MITSUBISHI	6
SIEMENS	6
ENDESA	5
SOLVAY S.A.	5
ABB	4
ALSTOM	4
DEGUSSA	4
DELPHI CORPORATION	4
HOECHST AG	4
IMPERIAL CHEMICAL INDUSTRIES	4
SHELL	4
ABN AMRO HOLDING N.V.	3
BANCA NAZIONALE DEL LAVORO S.P.A.	3
BASF	3
BTR	3
DEUTSCHE TELEKOM	3
EDISON	3
GUIDANT	3
HOWALDTSWERKE-DEUTSCHE WERFT AG	3
MANNESMANN AG	3

2.7.4 Top 20 Primary Acquiring and Target Firm Countries

Table 2.19: Top 20 Primary Acquiring and Target Firms' Countries, 1990-2014

Country acquiring firm	Country acquirer	Country target
USA	1,011	578
Germany	865	953
UK	651	692
France	493	407
Netherlands	329	395
Italy	157	275
Japan	145	85
Sweden	140	204
Switzerland	138	106
Spain	126	193
Austria	113	117
Left open	107	181
Luxembourg	106	59
Belgium	82	116
Denmark	77	83
Finland	67	76
Canada	60	43
Norway	56	64
Missing	36	41
Jersey	31	11

We display primary acquiring and target firms' countries for the top 20 primary acquiring firms' countries.

2.7.5 Number of Notifications and Decisions over Time

Table 2.20: Number of Notifications and Decisions by Year, 1990-2014

Year	Notifications	Decisions
1990	11	5
1991	55	49
1992	43	49
1993	44	46
1994	76	71
1995	91	95
1996	108	107
1997	137	119
1998	178	180
1999	243	232
2000	304	311
2001	314	319
2002	254	247
2003	184	194
2004	226	224
2005	313	301
2006	349	348
2007	388	393
2008	329	336
2009	241	233
2010	249	254
2011	283	293
2012	272	262
2013	269	266
2014	235	257
2015	.	5
Total	5,196	5,196

We count notifications by notification year and decisions by decision year.

2.7.6 Decisions by Broad Product Market

Table 2.21: Decisions by Broad Product Market, 1990-2014

Broad product market	Notifications	Phase-1 decisions	Phase-2 decisions
IT and services	66	66	0
agricultural products	690	382	304
air transport and travel	1,589	1,294	282
aircraft avionic equipment	6	6	0
aircraft supplies	61	3	58
aircrafts	164	141	23
airport services	7	7	0
automation	32	16	16
automotive industry	670	639	30
banking, financial services and insurance	1,835	1,823	11
betting and gambling	9	9	0
building materials	685	530	58
car components	974	946	28
care and justice services	5	0	0
catering and restaurants	42	28	9
chemicals	2,074	1,883	187
childcare products and toys	5	5	0
communication devices	97	86	11
communication services	1,663	1,396	247
computers (hardware and software)	827	801	26
construction	281	264	0
consulting	29	5	24
cosmetics	469	319	150
defense industry	110	110	0
electrical appliances	1,075	976	99
electricity devices (batteries etc.)	399	381	18
electricity supply	44	38	6
electronic components	239	239	0
electronic devices	43	43	0
energy plants	15	3	12
energy supply	2,435	2,171	170
engines	8	8	0
entertainment	36	36	0
explosives and weapons	115	115	0
fire fighting equipment	15	15	0
food and beverages	2,266	1,946	246
furniture	79	79	0
glass	4	4	0
healthcare	72	60	0
heating systems	11	11	0
industrial engineering	127	69	58

Table 2.21: Continued

Broad product market	Notifications	Phase-1 decisions	Phase-2 decisions
left open	265	243	3
luxury goods	17	17	0
machinery and equipment	864	796	68
management services	17	17	0
media	1,318	1,038	263
medical devices	911	647	264
medical services	72	70	0
medical supplies and products	51	51	0
metal products	623	594	29
metals and minerals	244	223	21
office supplies	51	51	0
optics	15	15	0
packaging	359	357	0
paints and colours	89	89	0
paper	279	134	145
paper products	415	345	70
passenger transport	4	4	0
personal services	2	2	0
personnel services	234	234	0
pet food	62	62	0
pharmaceuticals	2,431	2,326	77
photography	19	10	9
plastics	18	18	0
printing	25	25	0
protective equipment	60	60	0
railway industry	233	137	96
raw materials	699	653	46
real estate	151	151	0
retail	233	232	1
sanitary	157	148	9
security	6	6	0
ships and port services	106	99	7
sports industry	59	59	0
steel industry	26	26	0
storage	15	15	0
textile and clothing	129	124	5
tobacco	99	99	0
tourism industry	411	347	51
traffic management	41	38	3
transport and logistics	838	771	52
utilities	49	32	9
various	315	294	21
waste management	30	27	0

Table 2.21: Continued

Broad product market	Notifications	Phase-1 decisions	Phase-2 decisions
water supply	16	16	0
wood and wood products	20	15	5
Total	31,451	27,670	3,327

Note that phase-1 and phase-2 decisions do not add up to the number of notifications due to the 69 referrals to Member States and the 16 cases classified as "other".

Chapter 3

25 Years of European Merger Control [1]

3.1 Introduction

Competition policy, that is, the design and enforcement of competition rules, is a cornerstone of the European Union (EU)'s program to enhance the European single market and foster growth.[2] The European Commission's (EC) Directorate General for Competition (DG Comp) ensures the application of EU competition rules and retains jurisdiction over community-wide competition matters, representing the lead antitrust agency in the European context. Competition policy covers several areas ranging from monitoring and blocking anticompetitive agreements – in particular hardcore cartels – to abuses by dominant firms, to mergers and acquisitions as well as to state aid. Among these areas of antitrust enforcement, merger control plays a peculiar role. First, it is the only area where there is *ex-ante* enforcement. Second, it has important implications for the other areas of antitrust: if anticompetitive mergers that reduce competition and strengthen the dominant position of the merging firms are not prevented, it might make the *ex-post* control of abusive behaviors more difficult. Finally, mergers are the area of antitrust where the largest consensus on best practices exists. Therefore, among competition policy tools, it is an area that attracted much policy interest and economic research.

[1]This chapter is the accepted manuscript published in the DIW Discussion Paper Series as: Affeldt, P., Duso, T. and F. Szücs (2019). 25 Years of European Merger Control. *DIW Discussion Paper No. 1797.* We thank Ivan Mitkov, Fabian Braesemann, David Heine, Juri Simons and Isabel Stockton for their help with data collection.

[2]Gutiérrez and Philippon (2018) claim that since the 1990s, European markets have become more competitive than their US counterparts because of the increased economic integration and the enactment of the European single market. They attribute a key role in this process to the tough enforcement of competition policy rules.

The European Communities Merger Regulation (ECMR), the legal basis for common European merger control, came into force in 1990. Over the course of the next 25 years, European merger control saw significant changes. While in the early 1990s there were approximately 50 notified cases per year, the annual workload increased significantly in the late 1990s and has averaged around 280 cases in the 2000s. DG Comp's enforcement activity reflects these changes. Procedurally, many novelties were implemented in the 2004 amendment to the ECMR: not only were new horizontal merger guidelines and the office of the chief economist introduced, but also, more importantly, a new substantive test, the so called "significant impediment of effective competition" (SIEC) test and an efficiency defense were introduced. These amendments marked a substantial change in the legal basis for merger control enforcement in Europe. Yet, the pressure for these changes began much earlier with the increasing belief that a mere form-based assessment of mergers could often result in wrong decisions. The three overturned prohibitions by the Court of First Instance at the beginning of the 2000s marked the peak of this process.

In this paper, we employ a new dataset containing all merger cases with an official decision documented by DG Comp (more than 5000 individual decisions) to evaluate the time dynamics of the EC's decision procedures (see Affeldt, Duso, and Szücs (2018)). Specifically, we assess how consistently different arguments related to the so called structural market parameters – market shares, concentration, likelihood of entry, and foreclosure – put forward to motivate a particular decision were applied over time. In order to obtain a more fine-grained picture of the decision determinants, we extend our analysis to the specific relevant product and geographic markets concerned by a merger. Thus, instead of only looking at the determinants of a merger decision in the aggregate, we also investigate the factors that caused competitive concerns in specific sub-markets and how they have changed over time. This step is particularly important because larger mergers typically affect many different product markets in many different geographic regions. For example, the mergers in our data affect an average of six markets. Therefore, by analyzing individual markets, thus conducting a more disaggregate analysis, we better model the process that lead to a specific merger decision. Thus, the scope and depth of our data allow us to go beyond the existing literature by i) not relying on a sample of decisions but instead reporting patterns for the whole population of merger cases examined by DG Comp; and ii) allowing for heterogeneity within merger cases by examining the individual product and geographic markets concerned.

In a first step, and in line with the existing literature, we start by estimating the probability of intervention as a function of merger characteristics at the merger level. We find that the existence of barriers to entry, the increase of concentration

measures and, in particular, the share of product markets with competitive concerns are positively associated with the likelihood of an intervention. This approach naturally extends to the level of the individual markets: instead of estimating the overall probability of an intervention, we estimate the likelihood that competitive concerns are found in that specific product/geographical market under consideration. We find that, again, barriers to entry, but also the risk of foreclosure play a role. While tightly defined (national) markets increase the probability of concerns, the number of active competitors decreases it. Structural indicators of market shares and concentration show the expected positive and significant correlation with the likelihood of competitive concerns. After this static investigation, we then study the dynamics of the impact of a number of key determinants over time. We find that the importance of 'structural' indicators of market power has declined over the years, though we observe a large volatility in the estimates over time.

In a second step, we bring well-developed non-parametric prediction methods to the analysis of competition policy outcomes: supervised machine learning techniques. In particular, we implement the causal forest algorithm proposed by Athey and Imbens (2016). This step allows a more flexible approach to model the heterogeneity in merger control decisions. Specifically, the association between structural indicators and the Commission's decisions is made a function of all other covariates. Especially after the reform of 2004, a so-called effects-based approach centered on a clearly stated theory of harm was made a cornerstone of EU merger control. In such an approach, the reliance on structural parameters was expected to decrease, leaving space for the use of counterfactual analysis where the interactions of different elements might play a crucial role to substantiate the theory of harm. Using this model, we find that the importance of market share and concentration measures has declined over time while the importance of barriers to entry and the risk of foreclosure has increased in DG Comp's decision making. Yet, the impact of structural indicators appears to be much less volatile than in the simple linear probability model. Thus, the arguments put forward by the EC to substantiate its decisions appear to be more consistently applied once the process underlying these decisions is modelled in a flexible way.

The paper is structured as follows. In Section 3.2, we discuss the institutional details of European merger control and review studies that empirically investigate the determinants of merger intervention. In Section 3.3, we describe the dataset used. We present the parametric model as well as estimation results for the determinants of EC merger interventions in Section 3.4, while Section 3.5 presents the model and results for non-parametric estimation of heterogeneous correlations between merger characteristics and intervention by the EC. We conclude in Section 3.6.

3.2 Literature & Institutional Details

3.2.1 Institutional Details

The European Communities Merger Regulation (ECMR) was passed in 1989 and came into force in September 1990.[3] It specifies the scope of intervention and juridical competence of the European Commission in merger cases with a "community dimension." In article 1.2 of regulation 4064/89, a combination is defined to have community dimension by meeting the following conditions:

(a) the aggregate worldwide turnover of all the undertakings concerned is more than ECU[4] 5 000 million, and

(b) the aggregate Community-wide turnover of each of at least two of the undertakings concerned is more than ECU 250 million, unless each of the undertakings concerned achieves more than two-thirds of its aggregate Community-wide turnover within one and the same Member State.

That means that from 1990 onwards, all major combinations affecting EU markets have been scrutinized by the EC, whereas national competition authorities have been focusing solely on mergers affecting one single Member State. In 1997, the above definition was significantly widened by the passing of regulation 1310/97, which made the definition of a community dimension less stringent.[5]

Notice that these definitions also include companies that are located, produce, and sell outside of Europe, as long as their sales to European markets are sufficiently high. Thus, a merger can be subject to the jurisdiction of more than one competition authority. This resulted in diplomatic strife, for instance, when the merger of the two U.S. companies *General Electric* and *Honeywell* was ratified by American authorities, but prohibited by the European Commission.

Once it is established that a combination is subject to EC jurisdiction, the merging parties are required to notify the Commission prior to the implementation of the

[3] Council Regulation (EEC) No 4064/89 of 21 December 1989 on the control of concentrations between undertakings [Official Journal L 395 of 30 December 1989].

[4] ECU was replaced by Euro in 1998.

[5] Council Regulation (EC) No 1310/97 of 30 June 1997 [Official Journal L 180 of 9 July 1997] defines a community dimension when i) the combined aggregate worldwide turnover of all the undertakings concerned is more than EUR 2 500 million; ii) in each of at least three Member States, the combined aggregate turnover of all the undertakings concerned is more than EUR 100 million; iii) in each of at least three Member States included for the purpose of point (b), the aggregate turnover of each of at least two of the undertakings concerned is more than EUR 25 million; and iv) the aggregate Community-wide turnover of each of at least two of the undertakings concerned is more than EUR 100 million, unless each of the undertakings concerned achieves more than two-thirds of its aggregate Community-wide turnover within one and the same Member State.

concentration. On receipt of the notification, the Commission publishes a note in the Official Journal of the European Communities, where third parties can comment on the proposed transaction.

After the notification of the Commission (and the receipt of all necessary information), phase-1 proceedings are initiated. The EC then has 25 working days (which can be extended to a maximum of 35 working days) for an initial assessment of the merger. Based on this initial assessment the EC can clear the proposed merger (phase-1 clearance), clear it subject to remedies proposed by the merging parties (phase-1 remedy), or initiate a more in-depth investigation (phase-2 investigation) depending on whether the proposed transaction raises competitive concerns and depending on whether these can be addressed by initial remedies or not. Furthermore, the merging parties can also withdraw the proposed merger during phase-1 (phase-1 withdrawal).

If the EC initiates an in-depth investigation, the phase-2 investigation may take up to 90 working days. Following this second investigation phase, the EC can again unconditionally clear the merger (phase-2 clearance), clear the merger subject to commitments by the merging parties (phase-2 remedy) or prohibit the merger (phase-2 prohibition). Again, the merging parties can also withdraw the proposed merger in phase-2 (phase-2 withdrawal). It is argued that withdrawing a merger in phase-2 of the investigation process is virtually equivalent to a prohibition as parties often withdraw a merger before an actual prohibition by the EC can take place (Bergman, Jakobsson, and Razo, 2005). Hence, both a prohibition as well as a phase-2 withdrawal suggest that the EC and the notifying parties were unable to find suitable remedies to address the anti-competitive concerns of the proposed transaction. Thus, we thus consider prohibitions, phase-2 remedies, phase-2 withdrawals, and phase-1 remedies as an intervention in our empirical analysis.

Significant changes to European merger control were introduced in 2004 through an amendment to ECMR with the aim of bringing merger control closer to economic principles: the concept of an efficiency defense was introduced, a chief economist was appointed, the timetable for remedies was improved and horizontal merger guidelines were issued. The reception of the new merger regulation was generally favorable (Lyons, 2004). One of the most significant changes was the change from the "dominance test" (DT) for market power in favor of a "significant impediment of effective competition test" (SIEC).

The pre-2004 dominance test required the creation or strengthening of a dominant position as a necessary condition for the prohibition of a merger. It is argued that the dominance test was deficient in cases of collective dominance and tacit collusion, and that the "substantial lessening of competition" test employed by the

United States' Federal Trade Commission (FTC) would be preferable. After the 2004 reform, the test used by the European Commission can be most accurately described as a significant impediment of effective competition (SIEC) test, which is more closely aligned with U.S. practice (Bergman, Coate, Jakobsson, and Ulrick, 2007; Szücs, 2012).

3.2.2 Previous Literature

Mergers are studied extensively, with a large body of both theoretical and empirical literature on questions such as firms' incentives to merge and merger policy effectiveness. In the present paper, we evaluate the time dynamics of the EC's decision procedures and how the importance of structural market parameters in motivating a particular merger decision evolved over time. Thus, this paper most closely relates to the literature that empirically studies the determinants of merger policy intervention decisions by competition authorities.

Most of the related literature – with the prominent exceptions of Bradford, Jackson, and Zytnick (2018) and Mini (2018) – investigate the determinants of merger intervention decisions *at the merger level* and for a *sample of merger cases* only. The scope and depth of our data (see Section 3.3) allow us to go beyond the existing literature by, firstly, not relying on a sample of decisions but instead reporting patterns for the entire population of merger cases examined by DG Comp and, secondly, allowing for heterogeneity within merger cases by examining the individual product and geographic markets concerned. Furthermore, all of the existing literature uses parametric models to empirically study the determinants of merger intervention decisions. We instead go one step further and use flexible, non-parametric machine learning techniques to study the heterogeneity in the association between the structural market parameters and the intervention decision.

Bergman, Jakobsson, and Razo (2005) are the first to study the determinants of EU merger control. They employ a logit model for a sample of 96 EU merger cases to estimate the likelihood of going to phase-2 or prohibition decisions as a function of market-relevant and political variables. They find that decisions of the European Commission are only influenced by variables that directly affect welfare. In both estimated models (likelihood of phase-2 and likelihood of prohibition), the probability of intervention increases with the market share of the companies involved in the merger. Dummy variables indicating the possibility of post-merger joint dominance and the existence of entry barriers are also relevant determinants of the intervention decision while political/institutional variables are not significant. Bergman, Coate, Jakobsson, and Ulrick (2010) examine instead similarities between

EU and U.S. merger decisions using a sample of horizontal phase-2 mergers between 1990-2004 for both the EU (109 cases) and the U.S. (166 cases). They estimate a probit model for each regime to evaluate enforcement policy, where the dependent variable is an indicator for intervention (one for prohibition, approval subject to substantial remedies or withdrawal by the parties at least one month into the phase-2 investigation). They find that market shares, the Herfindahl-Hirschman-Index (HHI),[6] and entry barriers matter for the intervention decision. In a second step, they then apply the model of the EU authority to the U.S. case sample and vice versa to predict the challenge probabilities for dominant firm unilateral effect cases if the other regime had decided the case. For dominance mergers, the study finds that the EU is tougher than the U.S. on average, in particular for mergers with moderate market shares of the notifying parties. The U.S., on the other hand, seem to be more aggressive for coordinated interaction and non-dominance unilateral effects cases. In the most recent study, Bergman, Coate, Mai, and Ulrick (2016) update the dataset of Bergman, Coate, Jakobsson, and Ulrick (2010) by adding observations both to the EU as well as the U.S. dataset for the time period after the 2004 EU merger policy reform. The final dataset, covering 1993-2013, used in the analysis contains a sample of 151 EU phase-2 cases and 260 U.S. cases. Separate logit models on an intervention indicator variable are estimated for the EU cases (distinguishing pre- and post-reform) and U.S. cases. Market shares and entry barriers are found to have a significant positive effect on the probability of intervention. As the EU merger reform increases the likelihood that the EC challenges a merger under a coordinated effects theory of harm and reduces the likelihood that a merger case will raise concerns under the dominance standard, it should affect the difference between EU and U.S. policy. Predictions of interventions using the model of respectively the other jurisdiction (and distinguishing pre- and post-reform cases) show evidence of convergence between U.S. and EU case decisions in unilateral effects mergers, where EU policy seems to be less aggressive post-reform.

Similar to this study, Szücs (2012) investigates the convergence between U.S. and EU merger policy following the 2004 EU merger policy reform. In particular, he uses a sample of 309 EU and 286 U.S. merger cases scrutinized by DG Comp and the FTC, respectively, between 1991 and 2008. For each of the pre-reform EU, post-reform EU and U.S. merger samples, he estimates a logit model on the decision to intervene and then uses the estimated models to predict the probability of intervention for each merger case from the point of view of both competition authorities. Based on the decreasing differences in the predicted intervention probabilities between the EU and the U.S. authorities over time, he concludes that EU and U.S. merger policy

[6]The HHI is defined as the sum of squared market shares of all firms active in the market.

are converging in the era following the 2004 EU merger policy reform. Both pre- and post-reform, barriers to entry as well as the existence of a dominant player in the market increase the likelihood of intervention. Post-reform, also the HHI has a positive and significant effect on intervention.

Duso, Gugler, and Szücs (2013) evaluate European merger policy effectiveness along three dimensions: the predictability, correctness, and deterrence effects of a decision. Regarding predictability of European merger policy, Duso, Gugler, and Szücs (2013) estimate two probit models (one pre-reform, one post-reform) for a sample of 368 EU merger cases where the intervention decision of DG Comp (remedies or prohibition) is a function of *ex ante* observable merger characteristics. Unlike the existing literature, they do not use characteristics derived from the decision itself but constructed by matching the merger data to firm-level data from Datastream and Compustat. Prior to the 2004 merger policy reform full mergers, conglomerate mergers, and mergers, where the parties have high market value, increase the probability of intervention while mergers involving US firms are less likely to be challenged. Post-reform, mergers between U.S. firms, full mergers, and cross-border mergers, decrease the probability of intervention while conglomerate mergers are more likely to be challenged.

Mai (2016) studies the effect of the EU merger policy reform on the probability of a merger being challenged by DG Comp based on a sample of 341 phase-1 and phase-2 horizontal mergers between 1990 and 2012. The probability of a challenge in a probit model pooling pre- and post-reform cases is driven by the market shares of the merging parties, entry barriers, and some other factors. Political factors, measured as the country of the merging firms, are found to be insignificant. The merger reform reduces the probability of challenge by between 8 and 16 percentage points. Mai (2016) also estimates separate pre- and post-reform models and applies the methodology used by Bergman, Coate, Jakobsson, and Ulrick (2010), Szücs (2012), and Bergman, Coate, Mai, and Ulrick (2016) by predicting the probability of challenge for pre-reform mergers using the post-reform model and vice versa. The author finds that the EU merger policy seems to have slightly softened post-reform and that market shares and entry barriers are important predictors of challenge both pre- and post-reform. However, the importance of market shares is lower post-reform.

Two recent papers differentiate from the previous literature by significantly expanding the sample of mergers analyzed. Bradford, Jackson, and Zytnick (2018) empirically investigate whether European merger control is used for protectionism. Similar to our data, they collect information on all merger cases scrutinized by DG Comp between 1990 and 2014. However, their analysis is still conducted at the

level of the merger rather than the concerned product and geographic market. Furthermore, they do not collect information on the structural parameters of market shares, concentration, likelihood of entry, and foreclosure from the case documents. While the authors use control variables measuring relative market size and market concentration, both HHI as well as market size are based on European-wide industry sales data[7] rather than on the market shares of merging parties and competitors as reported in the case documents. The authors find that DG Comp did not intervene more frequently or extensively in transactions involving non-EU or U.S.-based firms. While transaction value, HHI, hostile takeovers, and whether the merger is horizontal increase the likelihood of intervention, mergers involving a financial sponsor, taking place in large markets, and being stock acquisitions are less likely to be challenged.

The paper that is most closely related to this study in terms of data is the study by Mini (2018). Similar to this paper and unlike all other studies, Mini (2018) also collected information on the universe of EU merger decisions from the publicly available case documents between 1990 and 2013, recording each market concerned by the transaction as a separate observation. Thus, for each merger, he records potentially many observations and collects similar merger and market level characteristics from the case documents as we do. He then estimates probit models at this concerned market level for horizontal overlap markets, interacting all explanatory variables with a post-reform indicator variable. In the first model, the main variables of interest are the merging parties' market shares and the change in market shares, while in the second he focuses on post-merger HHI as well as the change in HHI due to the merger. Similarly to Bergman, Coate, Jakobsson, and Ulrick (2010), Szücs (2012), Bergman, Coate, Mai, and Ulrick (2016) and Mai (2016), he uses the models to predict how the estimated pre-reform model would have handled post-reform cases, decomposing observed differences into policy and case mix effects. He concludes that while the EC changed neither its stance towards mergers to quasi-monopoly or monopoly nor towards mergers in unconcentrated markets, it has challenged fewer mergers due to unilateral concerns for mid ranges of market shares and HHI post-reform. Unlike previous studies (and also this paper), rather than using the midpoints of the market share ranges reported in the case documents, Mini (2018) constructs the expected market shares and expected HHI from the reported market share ranges. Thus, the author highlights the issue of measurement error in market shares and HHI and how to explicitly account for it in estimation.

[7]The HHI and market size variables are constructed based on European-wide sales at the two-digit NACE code industry level from the Amadeus database. Clearly, these measures are quite different from those calculated by the Commission itself in well-defined product and geographic markets.

Thus, Mini (2018) is the only paper that studies the determinants of merger policy interventions at the relevant product and geographic market level based on the population of European merger decisions as we do. However, we focus on a different aspect in our analysis by studying the heterogeneity in the association between structural market parameters and other merger and market characteristics and the intervention decision by DG Comp. To this end, we use flexible, non-parametric machine learning techniques and, in particular, show how the association between structural market parameters and the intervention decision has evolved over time. Unlike the existing literature, we let the data determine time patterns rather than imposing different pre- and post-reform models.

3.3 Data and Descriptives

The data contain almost the entire population of DG Comp's merger decisions, both in the dimension of time and with regard to the scope of the decisions encompassed. The data were obtained from the publicly accessible cases published by DG Comp on the EC's webpage.[8] We started data collection with the very first year of common European merger control, 1990, and included all years up to 2014. This amounts to data on the first 25 years of European merger control.

Rather than taking a particular merger case as the level of observation, we collected data at a more fine-grained level and defined an observation as a particular product and geographic market combination concerned by a merger.

For the analysis in this study, we dropped cases that were referred back to member states as well as phase-1 withdrawals.[9] The final dataset used in the estimation contains 5,109 DG Comp merger decisions, where each decision includes a number of observations equal to the number of product/geographic markets affected in the specific transaction. The dataset contains a total of 30,995 market level observations. For further details on the merger database as well as the data collection procedure, we refer the reader to the data documentation (Affeldt, Duso, and Szücs, 2018).

The dataset contains information on the name and country of the merging parties (acquirer and target), the date of the notification, the date of the decision[10] and the type of decision eventually taken by DG Comp (clearance, remedy, and prohibition)

[8]The types of notified mergers, decisions taken and reports for each of the EC's decisions can be downloaded from: `http://ec.europa.eu/competition/mergers/cases/` and `http://ec.europa.eu/competition/mergers/legislation/simplified_procedure.html`.

[9]We only have information on two phase-1 withdrawals in the data.

[10]Note that the notification of a merger and the decision do not necessarily take place in the same year. We calculate the number of notifications based on the notification year and the number of decisions of a certain type based on the decision year.

or whether the proposing parties withdrew the notification. The data also allow us to distinguish between a policy action taking place in the initial (phase-1) or second phase (phase-2) of the merger review process.

Figure 3.1 shows the number of yearly merger notifications, phase-1 merger cases, mergers cleared subject to remedies (phase-1 and phase-2) and prohibitions between 1990 and 2014. Overall, merger notifications show an increasing trend with a big drop around 2002. Most of the notified mergers are decided in phase-1: Phase-1 mergers track the number of notifications very closely. The number of mergers cleared subject to remedies increased dramatically after 1996 and oscillates between 10 and 25 per year in more recent years. The number of prohibitions varies between zero and three prohibitions per year.

Figure 3.1: Enforcement History of DG Comp Merger Cases, 1990-2014

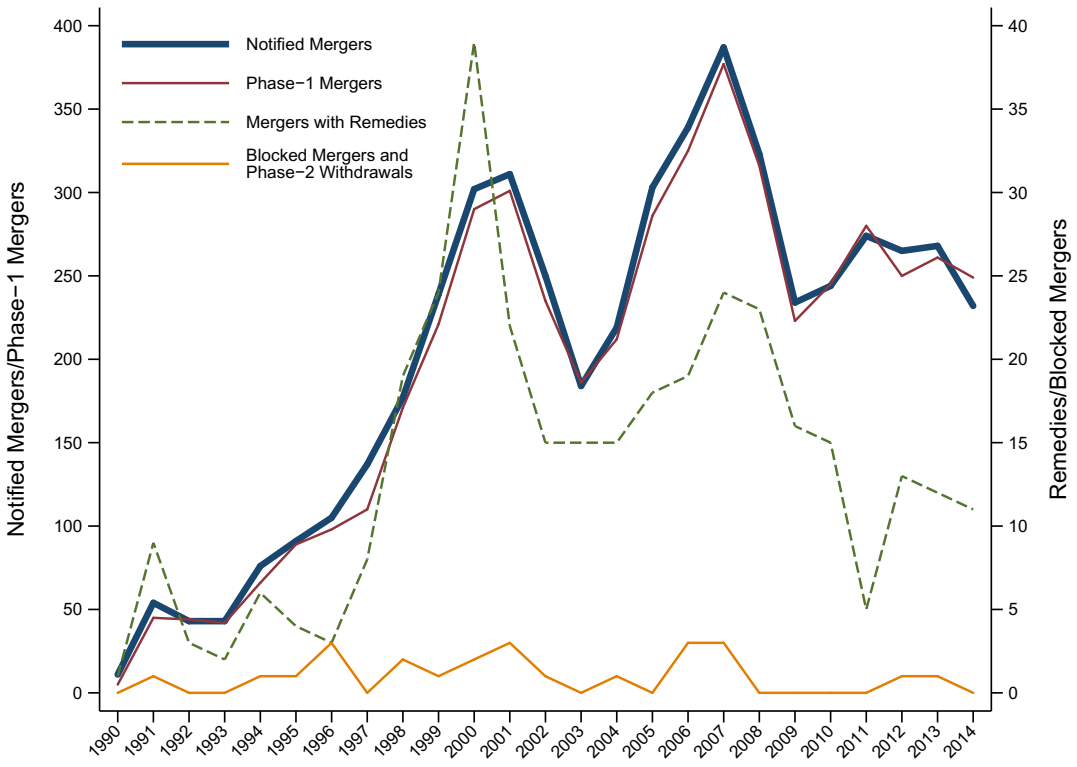

We report notified cases per notification year and phase-1 cases per decision year (left axis) as well as remedies (phase-1 and phase-2) and prohibitions per decision year (right axis). We exclude phase-1 withdrawals from the count of phase-1 mergers and include phase-2 withdrawals in the count of prohibitions. We exclude all cases where the decision type is "other."

The dataset further contains information on the nature of mergers. Variables for full mergers and joint ventures indicate whether DG Comp considered the case to be a full merger (55% of the notified mergers) and/or a joint venture (37% of the mergers); these are reported in Table 3.1.

Further indicator variables for vertical and conglomerate transactions indicate whether a product/geographic market is vertically affected by the merger (26% of the concerned markets) and whether the merger is conglomerate in nature in the particular concerned market (2% of the concerned markets), see Table 3.2.

Table 3.1: Summary Statistics Indicator Variables at Merger Level, 1990-2014

	0	1	mean	sd
Intervention	4,742	367	0.07	0.258
Full merger	2,293	2,816	0.55	0.497
Joint Venture	3,228	1,881	0.37	0.482

Furthermore, the dataset contains information on the geographic market definition adopted in each market by DG Comp. In about 58% of the concerned markets the geographic market is defined as national, in about 20% it is considered to be EU wide, in only 10% it is defined as a worldwide market while in about 12% of the cases the geographic market definition is left open (see Table 3.2).

We also observe which markets DG Comp considered to be problematic. The variable *concern* indicates the geographic and product markets affected by the merger, in which competitive concerns arose. This is the case in about 11% of markets. Further indicator variables record whether DG Comp considered barriers to entry to exist and whether DG Comp raised concerns that the merger would foreclose other firms in a particular market. As Table 3.2 shows, DG Comp considered entry barriers to exist in about 12% of the concerned markets, while risk of foreclosure was present in about 3% of markets.

Table 3.2: Summary Statistics Indicator Variables at Market Level, 1990-2014

	0	1	mean	sd
Concerns	27,675	3,320	0.11	0.309
Vertical merger	22,802	8,193	0.26	0.441
Conglomerate merger	30,472	523	0.02	0.129
National market	12,990	18,005	0.58	0.493
EU wide market	24,741	6,254	0.20	0.401
Worldwide market	28,037	2,958	0.10	0.294
Left open market	27,218	3,777	0.12	0.327
Entry barriers	27,423	3,572	0.12	0.319
Risk of foreclosure	30,184	811	0.03	0.160
No competitor information	13,733	17,262	0.56	0.497

The database also contains a count of the number of competitors in the concerned market and an indicator variable equal to one if no information on competitors is available. Merging parties face, on average, 1.6 competitors, with the number of competitors varying between 0 and 34. However, information on competitors is missing in about 56% of the markets - these are mainly mergers that were cleared in phase-1. We also include a variable indicating the complexity of a particular merger case, measured as the count of product/geographic markets concerned by the merger. A merger affects on average 6 geographic/product markets, ranging between one and 245 concerned markets.

Where available, data on the market shares of the merging parties were collected from DG Comp's competitive assessment in the decision document. Data availability is thus constrained by the extent of DG Comp's analysis. Market share information is collected at the level of the relevant product/geographic market combination. This information allows the calculation of the merging parties' combined market shares, the HHI and the change in HHI.[11]

Table 3.3 shows summary statistics for the market share related variables. The merging parties' average joint market share is 33%, with average post-merger HHI between 2,148 and 5,639 depending on the calculation method.[12] The mean change in HHI due to the merger is about 445, ranging from 0 to 8,450. As Table 3.3 shows, market share information is not available for all observations: while joint market share and HHI information is available for about 23,000 out of the 31,000 observations, the change in HHI due to the merger can be calculated for only about 13,000 observations.

Lastly, the data include information on the main industry in which a merger took

[11]Since DG Comp generally reports only a range of market shares in the publicly available documents, we defined the market shares to be equal to the central value of the interval. If for example the market share range indicated is [0-10] percent, we record a market share of 5 percent. If however the interval given in the decision is only 5 percentage points wide, we report the conservative lower market share bound. If for example the market share interval is [15-20] percent, we report 15 percent market share. Therefore, we cannot avoid that market shares contain measurement error; however this is an issue that this study shares with the existing literature. To our knowledge, Mini (2018) is the only one who, rather than using the midpoints of the market share ranges reported in the case documents, constructs the expected market shares and expected HHI from the reported market share ranges. Thus, he highlights the issue of measurement error in market shares and HHI, explicitly accounting for it in estimation.

[12]We calculate two different HHI measures. The variable *Post-merger HHI (low)* is a lower bound of the post-merger HHI: it is calculated as the square of the merging parties' joint market share plus the sum of squared market shares of competitors, whenever information on competitors' market shares is available. This assumes that competitors are very small whenever market share information of competitors is not available but market shares do not add up to 100%. The variable *Post-merger HHI (high)*, on the other hand, is an upper bound for the post-merger HHI: it adds the square of all missing market shares (100% minus all available market share information) to *Post-merger HHI (low)*. This hence treats all missing market share information as one missing competitor. In our empirical analysis, we use *Post-merger HHI (high)*.

Table 3.3: Summary Statistics Continuous Variables at Market Level

	mean	sd	min	max	observations
Joint market share	32.5	23.6	0	100	22,812
Post-merger HHI (low)	2,147.7	2,368.3	0	10,000	22,812
Post-merger HHI (high)	5,639.0	2,251.1	650	10,000	22,812
Delta HHI	444.7	779.1	0	8,450	12,875
Number of competitors	1.6	2.3	0	34	30,995

place. The industry is identified by NACE codes, which is the industry classification system used by the European Union to classify different economic activities. For the empirical analysis, we group the industries into 25 groups, as shown in Table 3.4, where some NACE codes are grouped together but, primarily, the manufacturing industry has been further divided into smaller subgroups. In 150 merger cases, the industry code was missing. For these cases, we went back to the decision documents and manually classified the mergers into the 25 industry groups according to our best judgement.

Table 3.4: Industry Groups, 1990-2014

Industry group	obs	cases
Accomodation and food service	192	64
Agriculture, forestry, fishing, mining	1,106	173
Arts, other services, households as employers	392	55
Electricity, gas, steam	1,381	280
Financial service activities	960	249
Information and communication	1,304	259
Insurance and pensions	925	237
Manufacturing (coke, petroleum, chemicals)	3,827	401
Manufacturing (computer, electronics, optical products)	1,702	247
Manufacturing (food, beverages, tobacco)	1,845	230
Manufacturing (furnitures , other manufacturing)	669	52
Manufacturing (machinery and equipment)	865	173
Manufacturing (metals and metallic products)	1,113	219
Manufacturing (motor vehicles, trailers, transport equipment)	1,539	302
Manufacturing (pharmaceuticals)	2,068	106
Manufacturing (rubber, plastic, non-metallic)	1,086	165
Manufacturing (textiles, clothes, leather)	169	31
Manufacturing (wood, paper, printing)	1,031	152
Public administration, education, human health, social work	169	47
Real estate, professional activities, administrative service activities	1,162	254
Repair, installation of machinery and equipment	1,046	200
Telecommuications	1,090	224
Transporting and storage	2,729	329
Water supply, waste management, construction	520	152
Wholesale and retail trade	2,105	508
Total	30,995	5,109

Note that all of these merger and market characteristics are characteristics, *as stated in DG Comp's decision documents.* As such, they reflect, to some extent, the assessment, subjective views, and potential mistakes of DG Comp. However, this issue is present in all papers in the empirical literature on the determinants of merger decisions.

The final merger sample contains information on 5,109 merger cases concerning 30,995 markets. For the analysis at the merger level, we take the mean value across concerned markets for those variables that vary at the market level.

3.4 Linear Probability Model

In this section, we explore the association between merger characteristics and the intervention decision by DG Comp within a parametric approach. We first replicate the results of the existing literature, which explain a competition authority's decision as a function of merger characteristics *at the merger level.* In contrast to previous studies, we explicitly estimate different models in various sub-samples to assess the issue of sample selection, which could arise because some important indicators – prominently market share and concentration measures – are only observable for ca. 60% of the mergers. Second, as a merger often affects many different markets, while its characteristics and effects on competition can be heterogeneous across these affected markets, we investigate in a second step the correlation between merger characteristics and DG Comp's intervention decision *at the market level.* Lastly, in order to allow for heterogeneity in the correlation between merger characteristics and intervention decisions, we look at the evolution of these relationships over time.

3.4.1 Methodology

We employ a linear probability model to estimate the relationship between merger characteristics and the intervention decisions of DG Comp.[13]

The dependent variable is an indicator variable for whether DG Comp intervened following a merger notification. We define the indicator variable *intervention* to be equal to one if DG Comp prohibited the merger, cleared the merger subject to remedies in phase-1, cleared the merger subject to remedies in phase-2, or the merging parties withdrew the merger proposal in phase-2. As Table 3.1 shows, DG Comp intervened in 367 out of the 5,109 merger cases in the estimation dataset (i.e. 7% of mergers).

[13]We decided to use a linear probability model rather than a probit or logit specification for easy interpretability of the estimated coefficients as well as the possibility to include industry fixed effects.

The estimation equation for the probability of intervention at the merger level is:

$$P_j(Y_j = 1|X_j, \overline{X}_{ij}, \eta_{m_j}, \eta_{t_j}) = \beta_0 + \beta_1 X_j + \beta_2 \overline{X}_{ij} + \eta_{m_j} + \eta_{t_j} + \epsilon_j \tag{3.1}$$

where i refers to a particular concerned market, j refers to a merger, m_j refers to an industry group, and t_j refers to the year when merger j took place. The merger characteristics X_j vary at the merger level, while X_{ij} are market-specific characteristics within merger j. In the merger-level regressions, we use the average of market-level variables ($\overline{X}_{ij}$).

This approach naturally extends to the level of the individual markets. Thus, in a second step, we estimate the correlation between market and merger characteristics and DG Comp's assessment at the level of the concerned product/geographic market. Instead of estimating the overall probability of intervention, the dependent variable used in the estimation at the market level is *concern*, which is a dummy variable indicating that a specific product/geographic market i affected by merger j raised competitive concerns according to DG Comp. As Table 3.2 shows, DG Comp raised competitive concerns in about 11% of the concerned markets.

The estimation equation for the probability of competitive concerns at the market level is:

$$P_{ij}(Y_{ij} = 1|X_j, X_{ij}, \eta_{m_j}, \eta_{t_j}) = \beta_0 + \beta_1 X_j + \beta_2 X_{ij} + \eta_{m_j} + \eta_{t_j} + \epsilon_{ij} \tag{3.2}$$

where the unit of observation is now the concerned market i in merger j rather than the merger j itself, X_j are the characteristics varying at the merger level, while X_{ij} are the characteristics varying at the market level.

Lastly, we explore the heterogeneity in the correlation between merger characteristics and competitive concerns by DG Comp over time. We run separate OLS regressions at the market level dividing the dataset into sub-samples based on the notification year.

The explanatory variables of primary interest are four determinants of competitive concerns that are expected to drive DG Comp's intervention decision. The so called structural market parameters - market shares, concentration, the likelihood of entry, and the likelihood of foreclosure - are measured as follows:

- Indicator variable for *high post-merger concentration*: equal to one if post-merger HHI is above 2000 and the change in HHI is larger than 150.[14]

[14] We used the variable *Post-merger HHI (high)* for the construction of the indicator variable. Results obtained with *Post-merger HHI (low)* are qualitatively similar.

- Indicator variable for *joint market share*: equal to one if the merging firms' joint market share is above 50% in the concerned market.[15]
- Indicator variable *barriers to entry*: equal to one if DG Comp considered barriers to entry to exist in the concerned market.
- Indicator variable *risk of foreclosure*: equal to one if DG Comp raised concerns that the merger would foreclose other firms in a particular market.

In addition to these four determinants of competitive concerns of a merger, we control for further merger characteristics. We include the market definition indicator variables for national, EU wide, and worldwide geographic markets as well as all information on the type of merger available in the data. Specifically, we use indicator variables for vertical mergers, conglomerate mergers, full mergers, and joint ventures; the count of the number of competitors in concerned markets; an indicator variable for whether information on competitors is missing in the data as well as a measure of the complexity of the merger measured by a count of the concerned markets.

Lastly, we include different industry and year fixed effects, depending on the specification. Industry dummy variables are defined for the 25 different industry groups as presented in Table 3.4. For the OLS regressions at the merger and market level, we include a set of industry-year fixed effects, controlling for unobserved time-varying industry specific factors.[16] For the regressions that explore the heterogeneity in the correlation between merger characteristics and competitive concerns over time, we regrouped the years 1990-1994 into one group for the sample splits, as there are relatively few merger cases in these early years of European merger control. In each of the year-specific OLS regressions, we include industry fixed effects. We corrected the error term by clustering standard errors at the industry group level.

3.4.2 Estimation Results

3.4.2.1 Determinants of Intervention - Merger Level

We present four specifications run at both the merger and market levels. Specification 1 is run on the full dataset without including the market share variables. Hence,

[15] We also run models where we use the level of the market shares rather than the dummy variable for high market shares. Results are similar. We decided to use the dummy for comparability with the approach based on machine learning discussed in Section 3.5.

[16] As a robustness check, we use industry and year fixed effects separately and include a set of time-varying control variables at the industry based on Worldscope data (e.g., mean size, mean total assets, mean Tobit's q, mean R&D...) as suggested by Clougherty and Seldeslachts (2013) and Clougherty, Duso, Lee, and Seldeslachts (2016). However, this does not qualitatively change the results.

this specification basically includes all mergers decided by DG Comp. Market share and concentration information is not available for all cases. If we include the market share variables in the regression, the sample size decreases significantly. However, the change in the estimated coefficients could be driven by selection (market share information is most frequently missing for phase-1 clearances) rather than just by the inclusion of the additional explanatory variables. Hence, specifications 2 and 3 present the results for the same specification as 1 split into those cases without information on market shares (specification 2) and those with information on market shares (specification 3). Lastly, specification 4 adds the indicator variables for joint market share above 50% and high concentration to specification 3.

Table 3.5 contains the regressions at the merger level. Reassuringly, we find that the EC's decision determinants are rather similar across all four sub-samples considered: the share of markets where entry barriers exist, the number of markets rising concerns, as well as the total number of markets affected by the merger increase the probability of a challenge. While the size of the effects is relatively constant for the number of markets affected, the impact of barriers to entry is almost 50% larger in cases where no market share information was gathered.

Neither merger characteristics (full mergers and joint ventures) nor the variables indicating alternative theories of harm (foreclosure concerns, vertical mergers, conglomerate mergers) significantly affect the Commission's decisions. Interestingly, the size of the concerned markets (national, EU wide, worldwide) also has no effect. In the full sample (column 1), we find some evidence for more challenges after the 2004 reform, but the coefficient is not precisely estimated in the other samples. Finally, in the sample including market share information (column 4), the indicator for a joint market share above 50% has no effect whereas the indicator pertaining to HHIs strongly and significantly increases the probability of challenge. Mergers in markets with HHIs above 2000 that entail an HHI increase of at least 150 are almost 9% more likely to be remedied or blocked.

3.4.2.2 Determinants of Concern - Market Level

Table 3.6 contains the same sets of regressions at the concerned market level. In general, more covariates appear to be significantly associated with competitive concerns at the market level than what is observed at the merger level. While this might be a statistical results due to the larger number of observations in these regressions, it is likely that the aggregation to the merger level hides some of the EC's more fine-grained considerations concerning specific markets.

In line with the merger level regressions, we find that barriers to entry increase the

Table 3.5: Linear Probability Model for Intervention (Merger Level)

	(1) Full sample	(2) Selected sample no market share info	(3) Selected sample market share info	(4) Selected sample market share info
Mean barriers to entry	0.2673***	0.3793***	0.2278**	0.2127**
	(0.0560)	(0.0786)	(0.0899)	(0.0857)
Mean risk of foreclosure	0.0145	-0.0289	0.0016	0.0040
	(0.0691)	(0.0878)	(0.1115)	(0.1087)
Fullmerger	-0.0019	0.0170	-0.0079	-0.0044
	(0.0194)	(0.0116)	(0.0483)	(0.0472)
Joint Venture	-0.0150	0.0147	-0.0321	-0.0283
	(0.0159)	(0.0105)	(0.0464)	(0.0449)
Mean conglomerate merger	-0.0051	0.0404	-0.0222	-0.0238
	(0.0471)	(0.0770)	(0.0735)	(0.0740)
Mean vertical merger	-0.0024	0.0155	-0.0269	-0.0067
	(0.0107)	(0.0145)	(0.0240)	(0.0241)
Mean market definition national	0.0103	-0.0059	0.0171	0.0143
	(0.0075)	(0.0047)	(0.0646)	(0.0621)
Mean market definition EU wide	0.0202	0.0079	0.0068	0.0066
	(0.0137)	(0.0111)	(0.0589)	(0.0578)
Mean market definition worldwide	-0.0158	-0.0069	-0.0343	-0.0382
	(0.0120)	(0.0113)	(0.0781)	(0.0767)
Number of concerned markets	0.0036***	0.0030***	0.0030***	0.0031***
	(0.0005)	(0.0011)	(0.0009)	(0.0008)
Percentage of markets with concerns	0.9375***	0.7312***	0.9681***	0.9340***
	(0.0623)	(0.1094)	(0.1107)	(0.1117)
Total number of competitors in all product markets	0.0004	0.0003	0.0008	0.0006
	(0.0004)	(0.0008)	(0.0005)	(0.0005)
Post reform indicator	0.0333**	0.0042	0.1169	0.1384*
	(0.0147)	(0.0069)	(0.0824)	(0.0768)
Joint market share above 50%				-0.0009
				(0.0481)
HHI $\geq$ 2000 & Delta HHI $\geq$ 150				0.0881***
				(0.0169)
Constant	-0.0541***	-0.0211**	-0.1110	-0.2210**
	(0.0177)	(0.0090)	(0.0913)	(0.0924)
Industry Group Year FE	Yes	Yes	Yes	Yes
R2	0.609	0.557	0.682	0.689
Observations	5,109	3,665	1,444	1,444

We report heteroskedasticity robust standard errors clustered at the industry group level. Significance at the 1%, 5%, and 10% levels is represented by ***,** and * respectively.

likelihood of competitive concerns at the market level as well. In addition, the risk of foreclosure also has a positive and significant, though smaller, effect. Joint ventures appear to be treated more leniently. Market size now plays a more decisive role, with national markets increasing the probability of concerns in all specifications except (2). While the total number of competitors (across all markets) was insignificant at the merger level, the number of competitors in a specific market decreases the probability of competitive concerns in all four specifications. When the EC does not collect information on competitors, i.e. it does not spend too much time and effort to define the relevant market, the likelihood of concerns is expectedly lower.

Finally, in the sub-sample with market share information, both market power indicators now significantly raise the chance of concerns: a joint market share in excess of 50% increases it by almost a quarter, while the HHI indicator increases it by 10%.

3.4.2.3 Determinants of Concern - Market Level - Split Sample over Time

We explore the heterogeneity in the correlation between merger characteristics and competitive concerns by DG Comp over time by running separate OLS regressions splitting the market-level dataset over years (regrouping notification years 1990-1994).[17] For each of the sub-samples, we run specification 4 of the previous regressions - hence, the indicator variables for high concentration and joint market share above 50% are included as explanatory variables in all regressions. Although this decreases the sample size, we consider market share and concentration to be important determinants of merger decisions, thus these are included in the analysis. As discussed in the previous section, while the estimated coefficients might differ across samples, the relevant determinants of intervention or competitive concerns are the same across the different subsamples.

In this section, we only present regression coefficient plots for our four main explanatory variables of interest. The underlying regression results are found in Appendix 3.7.1. Note that we have relatively few observations from 2014 that include market share information. For this subsample, the barriers to entry indicator perfectly predicts the outcome variable of competitive concerns. We therefore show coefficient plots only up to and including the year 2013.

Figure 3.2 shows the impact of the HHI indicator. With few exceptions, coefficient

[17] We also explore whether the correlation between the main variables of interest and concerns identified by DG Comp differs across industries. We ran analogous specifications splitting the sample over industries rather than time. OLS regression results, as well as coefficient plots equivalent to the ones shown here, are found in Appendix 3.7.2.

Table 3.6: Linear Probability Model for Concern (Market Level)

	(1) Full sample	(2) Selected sample no market share info	(3) Selected sample market share info	(4) Selected sample market share info
Barriers to entry in submarket	0.3856***	0.3408***	0.4067***	0.3160***
	(0.0558)	(0.0856)	(0.0485)	(0.0406)
Risk of foreclosure in submarket	0.2066**	0.2958**	0.1849*	0.1777*
	(0.0956)	(0.1248)	(0.0921)	(0.0951)
Fullmerger	-0.0375	-0.0071	-0.0615	-0.0586
	(0.0250)	(0.0263)	(0.0373)	(0.0347)
Joint Venture	-0.0656**	-0.0218	-0.1192***	-0.1061***
	(0.0244)	(0.0285)	(0.0323)	(0.0301)
Conglomerate merger in submarket	0.0201	0.0302	0.0259	0.0140
	(0.0372)	(0.0469)	(0.0355)	(0.0353)
Vertical merger in submarket	-0.0024	0.0240	-0.0410***	-0.0135
	(0.0100)	(0.0180)	(0.0128)	(0.0125)
Market definition national	0.0182***	0.0042	0.0690***	0.0634***
	(0.0049)	(0.0076)	(0.0239)	(0.0213)
Market definition EU wide	-0.0108	0.0007	0.0039	0.0264
	(0.0087)	(0.0129)	(0.0246)	(0.0248)
Market definition worldwide	0.0076	0.0176	0.0245	0.0496**
	(0.0163)	(0.0224)	(0.0252)	(0.0224)
Number of concerned markets	0.0001	-0.0001	0.0002	0.0000
	(0.0003)	(0.0005)	(0.0004)	(0.0003)
Number of competitors	-0.0099***	-0.0066***	-0.0116***	-0.0080**
	(0.0030)	(0.0020)	(0.0040)	(0.0036)
Indicator no info on competitors	-0.0652***	-0.0358***	-0.0792***	-0.0502**
	(0.0152)	(0.0124)	(0.0230)	(0.0202)
Post reform indicator	-0.1916	-0.0332	-0.3779	-0.3113
	(0.1300)	(0.0305)	(0.2222)	(0.2339)
Joint market share above 50%				0.2313***
				(0.0226)
HHI $\geq$ 2000 & Delta HHI $\geq$ 150				0.1043***
				(0.0134)
Constant	0.2355*	0.0640**	0.4508*	0.2658
	(0.1360)	(0.0279)	(0.2417)	(0.2557)
Industry Group Year FE	Yes	Yes	Yes	Yes
R2	0.377	0.410	0.401	0.473
Observations	30,995	18,185	12,810	12,810

We report heteroskedasticity robust standard errors clustered at the industry group level. Significance at the 1%, 5%, and 10% levels is represented by ***,** and * respectively.

estimates are positive but only significantly during the years 1999-2001, as well as in 2003, 2005, and 2007. Thus, in the last six years of the data, 2008 - 2013, high concentration was not a significant determinant of competitive concerns.

Figure 3.2: OLS Regression Coefficient on High Concentration over Time

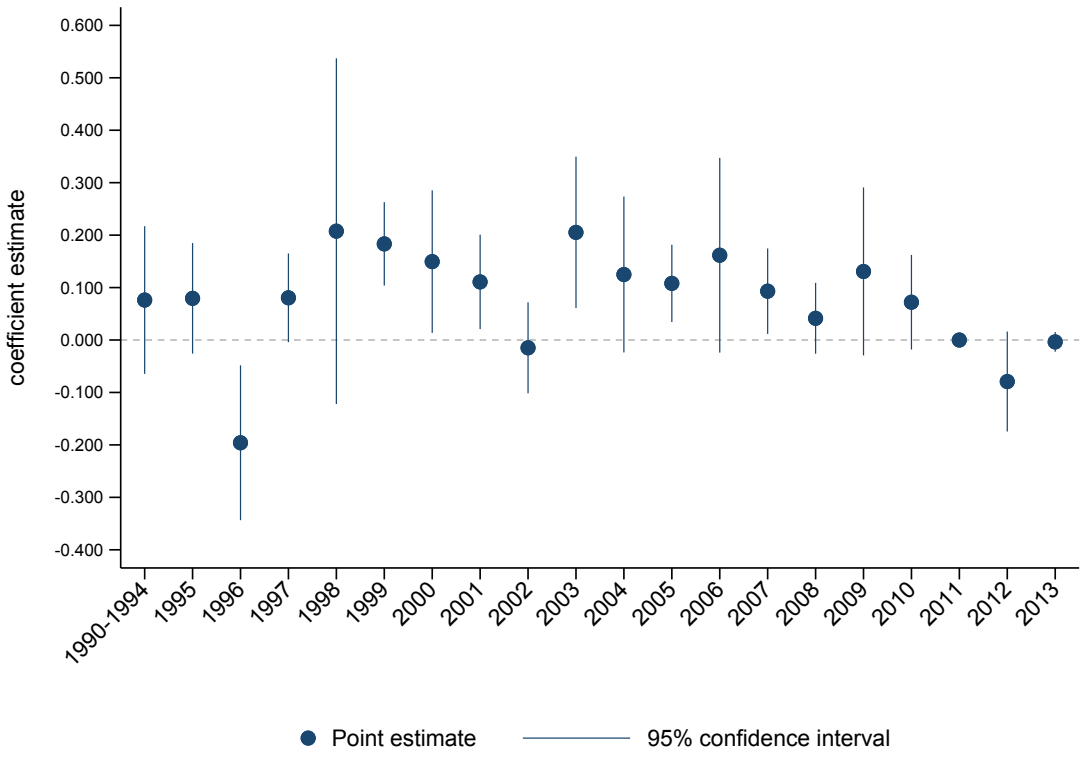

Regression coefficient on indicator variable for post-merger HHI above 2000 and change in HHI due to the merger larger than 150 in OLS regression on concerns. Each reported coefficient stems from a separate regression for the respective time period. Confidence intervals are based on heteroskedasticity robust standard errors clustered at the industry group level.

In Figure 3.3, we repeat the exercise focusing on the time dynamics of the joint market share of the merging parties. The impact of market share on competitive concerns was - with the exception of 2006 - consistently significant and positive from 1996 to 2009. The coefficient estimates are roughly twice the size of those associated with the concentration indicator presented above, suggesting that a high market share of the merging parties carries more weight in DG Comp's assessment than overall high concentration. However, similarly to the concentration measure, the importance of market shares seems to have declined after 2009.

Figure 3.3: OLS Regression Coefficient on Joint Market Share over Time

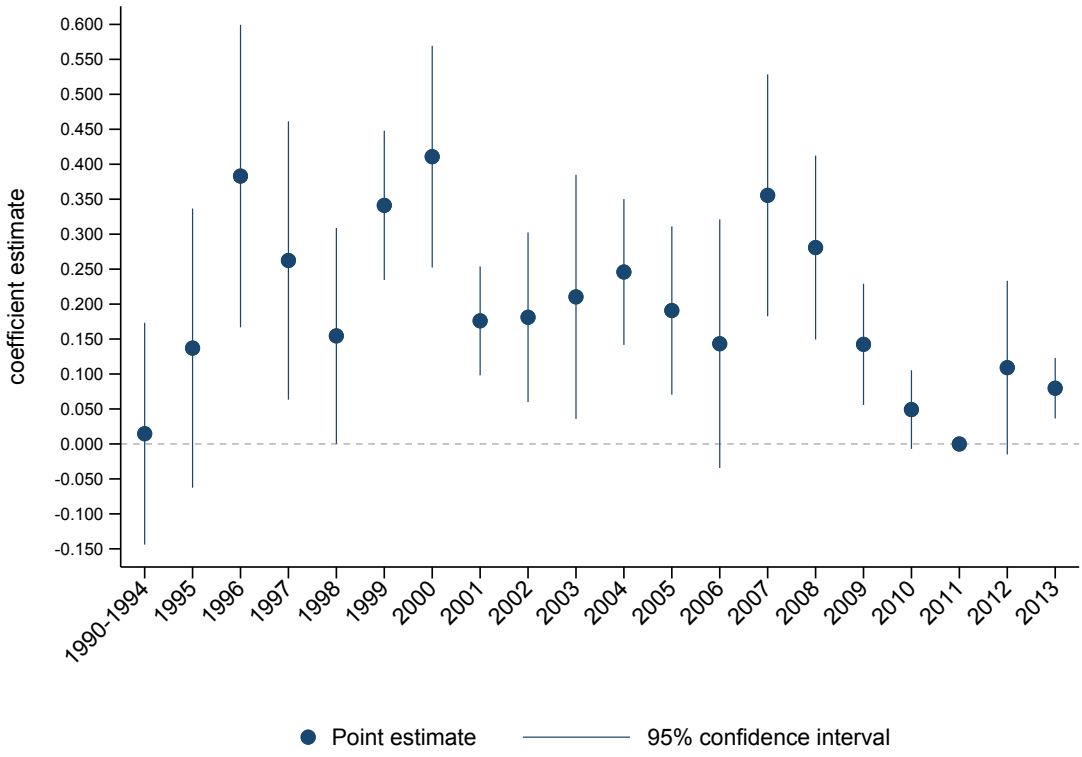

Regression coefficient on indicator variable for joint market share above 50% in OLS regression on concerns. Each reported coefficient stems from a separate regression for the respective time period. Confidence intervals are based on heteroskedasticity robust standard errors clustered at the industry group level.

Figure 3.4 reports the coefficient estimates for barriers to entry in different time periods. Similar to market shares, barriers to entry were consistently associated with a higher probability of intervention for a long period of time (1998 to 2009, with the exception of 2007). The size of the effect is, on average, even larger than that of market shares. As with market shares and high concentration, the importance of barriers to entry seems to have declined in the last years of the data.

Figure 3.4: OLS Regression Coefficient on Barriers to Entry over Time

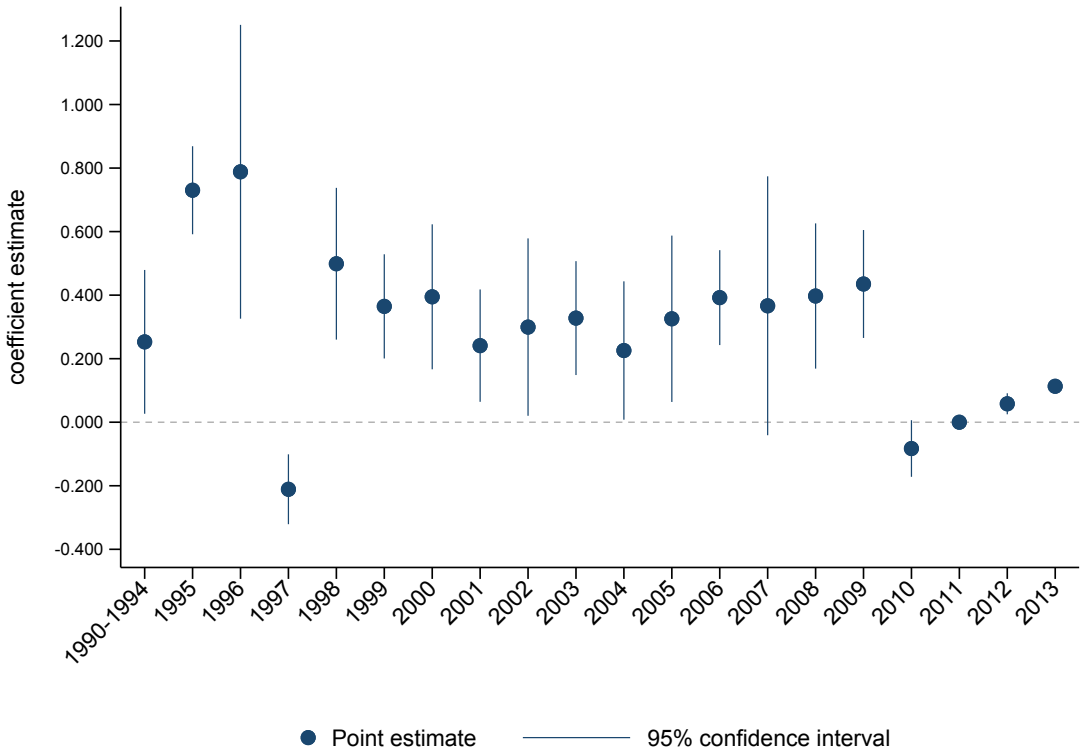

Regression coefficient on barriers to entry in OLS regression on concerns. Each reported coefficient stems from a separate regression for the respective time period. Confidence intervals are based on heteroskedasticity robust standard errors clustered at the industry group level.

Finally, in Figure 3.5 we report the period-specific coefficients associated with foreclosure concerns. While the coefficients are positive and, in a few periods, significant, no clear pattern seems to emerge. Note that the coefficients reported as zero without confidence intervals indicate years, in which no cases with foreclosure concerns were handled.

Figure 3.5: OLS Regression Coefficient on Risk of Foreclosure over Time

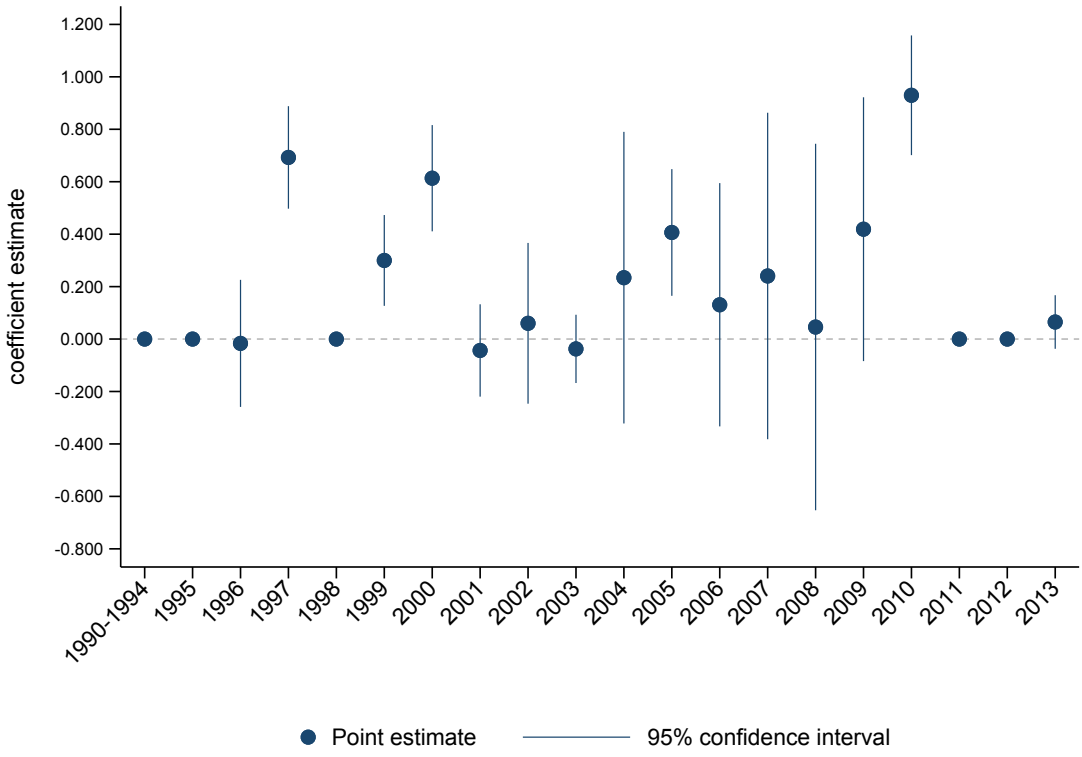

Regression coefficient on risk of foreclosure in OLS regression on concerns. Each reported coefficient stems from a separate regression for the respective time period. Confidence intervals are based on heteroskedasticity robust standard errors clustered at the industry group level.

3.5 Machine Learning/Causal Forests

In Section 3.4, we explore the association between concentration, market shares, entry barriers, and the risk of foreclosure with the intervention decision by DG Comp parametrically. However, the correlation between these variables might differ for different types of mergers. We try to investigate this heterogeneity by running separate regressions over time and industries. In this section, we take the idea of heterogeneous effects one step further by employing machine learning techniques. Specifically, we use the causal forest algorithm developed by Athey and Imbens (2016), Wager and Athey (2017), and Athey, Tibshirani, and Wager (2017) to explore the heterogeneity in these correlations non-parametrically. Causal forests are a flexible tool to uncover heterogeneous effects, in particular when there are many covariates and potentially complex interactions between them. They allow getting the richest possible specification supported by the data. This has three main advantages.

First, this approach allows a much better modelling of the process that leads to a particular decision by taking into account the specificities of each merger. As an example, consider that we want to measure the impact of high market shares on the likelihood that a market is considered problematic. In a facts-based approach, the Commission would surely consider that high market shares have a different impact if the market is narrowly defined or whether it is global in nature. Further, it is likely that industry specific information might also play a role: in national telecom markets, the role of high market shares is likely to be different than in a global manufacturing market. The strength of machine learning tools is that they allow determining the relevant interactions among covariates based on the observed data.

Second, by generating a more "saturated" model through the many interactions, this approach makes omitted variable bias less relevant than in the standard simple additive linear probability model discussed in the previous sections and used in the literature. While we still should be careful to interpret the coefficient estimates in a causal way, the potential bias in the coefficient estimates should be reduced. Put differently, the correlations that we retrieve are less spurious than in the OLS model.

Third, this approach makes the exact definition of the considered variables less relevant. When building the database, we face the trade-off between defining simple and general variables comparable across thousands of different mergers and the need to better measure single aspects of a decision. Therefore, some of our key concepts are measured by means of simple dichotomous dummy variables rather than more complex metrics. While this might be more problematic in the model discussed in the previous sections, it is less relevant in the context of this model, where the covariates become complex interactions among all indicator variables.

3.5.1 Methodology

3.5.1.1 Background on Heterogeneous Treatment Effects

The main goal of our analysis is to understand how the effect of one explanatory variable (in the present application, concentration, market shares, entry barriers, and risk of foreclosure) on an outcome variable (here, the competitive concerns raised by DG Comp) varies with the nature of the merger, where the nature of the merger is described by all other merger and market characteristics included in the dataset. Hence, we want to explore the heterogeneity in the effect of a key parameter of interest. This question relates to the literature on heterogeneous treatment effects, where one major problem is the fear that researchers might iteratively search for subgroups with high treatment effects and only report results for these subgroups. The reported heterogeneity in treatment effects might then be purely spurious.

The causal tree and causal forest algorithms address this problem as they non-parametrically identify subgroups that have different treatment effects. The methodology lets the data discover the relevant subgroups without invalidating the confidence intervals constructed on the treatment effects within the subgroups (Athey and Imbens, 2016).

In the context of heterogeneous treatment effect estimation, the model to be estimated is:

$$Y_{ij} = \tau(X_{ij})W_{ij} + \mu(X_{ij}) + \epsilon_{ij} \tag{3.3}$$

where Y_{ij} is the outcome variable (binary in the present case) for market i in merger j, W_{ij} is a binary treatment variable (i.e. our structural indicators), $\tau(X_{ij})$ is the effect of W_{ij} on Y_{ij} at point X_{ij} in covariate space, and ϵ_{ij} is an error term that may be correlated with W_{ij}. Using the notation of the potential outcomes framework by Rubin (1974), the treatment effect can be written as:

$$\tau(x) = \mathbf{E}\left[Y_{ij}^1 - Y_{ij}^0 | X_{ij} = x\right] \tag{3.4}$$

where Y_{ij}^1 is the potential outcome for unit ij under treatment – i.e. whether the EC identifies a concern when market shares are high – and Y_{ij}^0 is the potential outcome for unit ij absent treatment – i.e. whether the EC identifies a concern when market shares are low – where one of the two is not observed. The aim is to estimate how the function $\tau(x)$ varies with the covariates X. As Athey, Tibshirani, and Wager (2017) highlight, this is different from estimating a single parameter such as an average treatment effect while controlling for a large set of covariates, X.

The so-called unconfoundedness assumption implies that the treatment assignment W_{ij} is independent of potential outcomes Y_{ij} conditional on X_{ij}. This means that observations that are "close" in X-space can be treated as having come from a randomized experiment. Untreated observations that are close to the treated observation i under consideration can then be used to predict the outcome Y_{ij}^0 absent the treatment. In these instances, methods such as nearest-neighbor matching or other local methods allow for consistently estimating $\tau(x)$.

Notice that this is essentially the same identification assumption used in the OLS model discussed above. Thus, exactly as in that model, the causal interpretation of $\tau(x)$ should be careful, as the structural indicators could be correlated to the error term because of omitted factors. However, as discussed above, the causal forest model might be expected to outperform the simple OLS model since it contains a larger sets of covariates. Nonetheless, we cannot claim that we estimate any causal effect of these variables on DG Comp's intervention decision. We rather estimate the correlation between these treatment variables W_{ij} and the intervention decision

Y_{ij} and how this correlation varies with merger characteristics X_{ij}.

3.5.1.2 Estimation using Causal Forests

We use the causal forest algorithm by Athey, Tibshirani, and Wager (2017) implemented in the generalized random forest (grf) package in R to investigate how the correlation between the treatment variables and DG Comp's intervention decision varies with merger characteristics. Causal forests are based on the random forest methodology by Breiman (2001). They were developed by Athey and co-authors in a series of papers (see Athey and Imbens (2016), Wager and Athey (2017), and Athey, Tibshirani, and Wager (2017)), extending the regression tree and random forest algorithms so as to estimate average treatment effects for different subgroups, rather than predicting outcomes as is the case for regression trees and random forests.

In a standard regression tree, the aim is to predict individual outcomes Y_{ij} using the mean outcome Y of observations that are "close" in X-space. To determine which observations are "close," the algorithm starts to recursively split the covariate space (binary splits) until it is partitioned into a set of so-called leaves L that contain only a few observations. The algorithm automatically decides on the splitting variables and split points based on an in-sample goodness-of-fit criterion such as a mean squared error (i.e. how close the predicted outcomes are to the actual outcomes). The outcome Y_{ij} for observation ij is then predicted by identifying the leaf containing observation ij based on its characteristics X_{ij} and setting the prediction to the mean outcome within that leaf. A random forest is essentially an ensemble of trees, where the predictions of outcomes Y_{ij} are averaged across all trees in the forest to reduce variance and produce more robust predictions.

In case of a causal forest, we are not interested in predicting individual outcomes Y_{ij} but individual treatment effects $Y_{ij}^1 - Y_{ij}^0$ to study how treatment effects vary by subgroup. This implies that standard fit measures used in regression trees and random forests, such as the mean squared error, are not available since one of the potential outcomes and hence the actual treatment effect is never observed. However, the causal forest methodology builds on regression tree methods in that it also applies a "goodness-of-fit" criterion in treatment effects to decide on splits. Athey and Imbens (2016) show that the mean squared error function of a causal tree can be estimated and is a function of the variance of the estimated treatment effect. Basically, the goodness-of-fit measure to be minimized rewards a partition of the data for finding strong heterogeneity in treatment effects and penalizes a partition for high variance in leaf estimates. Minimizing the expected mean squared error of predicted treatment effects (rather than the infeasible mean squared error), is

shown to be equivalent to maximizing the variance of the predicted treatment effects across leaves with a penalty for within-leaf variance (variance of treatment and control group mean outcomes within leaves).

Within a causal tree, the conditional average treatment effects are then simply estimated as the difference of mean outcomes between treated and control observations within a leaf. Thus, causal trees are similar to nearest-neighbor methods as they also rely on the unconfoundedness assumption and use "close" observations to predict treatment effects. However, rather than defining closeness based on some pre-specified distance measure (such as Euclidean distance in k-nearest-neighbor matching), closeness is defined with respect to a decision tree and the closest control observations to ij are those that fall in the same leaf.

A causal forest, is then essentially an ensemble of causal trees, which only uses a random subset of the full dataset to grow each individual causal tree. The causal forest algorithm by Athey, Tibshirani, and Wager (2017) then weights nearby control observations according to the fraction of trees in which a control observation appears in the same leaf as the treated observation ij (Athey, Tibshirani, and Wager, 2017). This implies that for each observation an individual treatment effect τ_{ij} can be estimated while in a causal tree all units assigned to a given leaf have the same estimated treatment effect (Wager and Athey, 2017).

Athey and Imbens (2016) further introduce so-called "honesty" in causal trees to ensure correct inference: the data is divided in half, where one-half of the data is used to build the tree (i.e. determine the splits in covariate space) and the other half is used to predict treatment effects. Wager and Athey (2017) extend this idea to causal forests and develop theory for inference in causal forests. Thus, the causal forest algorithm by Athey, Tibshirani, and Wager (2017) does not only allow for predicting treatment effects but also for predicting confidence intervals.

The big advantage of causal trees and forests is that they allow the data to determine the relevant subgroups in a flexible, data-driven way without invalidating confidence intervals. This is particularly important in applications with many covariates and potentially complex interactions between these covariates that matter for measuring the effects. Wager and Athey (2017) also highlight that an advantage of trees is that the leaves can be narrower along some dimensions and wider along others, depending on how fast the signal is changing. For further technical background on the causal forest methodology and the implementation using the grf package, see Appendix 3.7.3.

As for the regressions presented in Section 3.4, we run the causal forests at the market (ij) rather than merger level (j). The outcome variable is therefore the *concern* dummy variable that indicates which specific product/geographic market

affected by the merger raised competitive concerns according to DG Comp. We run four different causal forests, each including one of the four determinants of competitive concerns that should influence DG Comp's intervention decision (the treatment variable in causal forest terminology). These are the same four indicator variables as those used in the previous regressions: *high post-merger concentration*, *joint market share above 50%*, *barriers to entry*, and *risk of foreclosure*.

In addition to the treatment variable, each of the causal forests includes a set of covariates X over which the correlation between the variable of interest and the outcome is allowed to vary. These are essentially the same as in the regression analyses of Section 3.4. Different from the regression analyses, we include the notification year as a continuous variable from 1990 to 2014 rather than year fixed effects, which allows the algorithm to determine the relevant binary splits over time. We include the market definition indicator variables for national, EU wide, and worldwide geographic markets as well as all information on the type of merger available in the data – vertical mergers, conglomerate mergers, full mergers, joint ventures, a count of the number of competitors in the concerned market as well as an indicator variable for whether information on competitors is missing in the data, and the complexity of the merger measured by a count of the concerned markets. Lastly, we include a set of industry fixed effects which are industry dummy variables for the 25 different industry groups defined as presented in Table 3.4.

Each of the causal forests is grown with a minimum node size of 10 and consists of 5000 trees.[18] Also note that the dataset used for the estimation of the causal forests for barriers to entry and risk of foreclosure differs from the dataset used for the estimation of the causal forests for the high concentration and joint market share measures. The dataset where the treatment variable is based on market share information has fewer observations because market shares are not available for all mergers. See the discussion of the issue in Section 3.4.2.1.

3.5.2 Estimation Results

In this section, we present the results of the correlation analysis between the four main variables of interest and the competitive concerns by DG Comp using causal forests. While a causal forest allows for predicting conditional average treatment effects, we are not primarily interested in the average correlation between a variable

[18]The term "minimum node size" is a bit misleading. The minimum node size in a causal forest is rather the minimum number of observations that must be part of a node in order for a split to be attempted. We ran causal forests for the entry barrier treatment using minimum node sizes of 5, 10, 15, 20, 30, and 40. The estimated conditional average treatment effect did not change much using these different node sizes.

of interest and the outcome variable, rather, we want to explore and visualize how this correlation varies over the covariate space X. We look in particular at how the correlation between high concentration, market shares, entry barriers, risk of foreclosure, and concerns identified by DG Comp varies over time and industry. We only show and discuss results for the variation over time here, predicted correlations across industries are shown in Appendix 3.7.5 as variation across industries is relatively small.

In order to explore how the correlation between the treatment variable and the outcome varies with one dimension included in the covariates X, we need to hold all other variables included in X constant and vary only the covariate of interest.[19]

The prediction plots below are obtained as follows: We generate a prediction dataset that contains the range of one X variable of interest (here notification year), for which we want to explore the heterogeneity in the association between the treatment variable and the outcome variable. We set all the other covariates included in X to their mean respectively median sample value.[20] We then predict the treatment effects at the data points of this prediction dataset using the causal forest grown and plot the treatment effect along with the point-wise 95% confidence intervals. In short, we take the mean/median merger in terms of all covariates, except time, and look at how the predicted correlation between for example the presence of entry barriers and competitive concerns varies if that mean merger had been notified in different years.[21]

Once again, given that our treatment variables might be correlated with the error term, we interpret the predicted treatment as the correlation between this variable and the probability that DG Comp found competitive concerns in the affected market. Further, we discuss how this correlation varies over time.

[19]See also the example of the effect of child rearing on labor-force participation provided in Athey, Tibshirani, and Wager (2017), where the mother's age at first birth and the father's income are varied while all other covariates are set to their median values.

[20]This also implies that indicator variables are set to their mean sample value; for example, the mean value of an industry dummy variable. This also explains the sometimes large difference in predictions setting all other covariates to mean or median values, since the median of a dummy variable will be either zero or one.

[21]Rather than taking the mean merger over the entire sample, we also created a prediction dataset based on the mean merger for which we have information on the market shares and concentration variables. We then used this prediction dataset to create alternative predictions based on the causal forests for high concentration and joint market share. As the predicted "treatment" effects did not change by much, we only report the predictions based on the mean merger over the entire sample.

3.5.2.1 Treatment - High Concentration

Figure 3.6 shows the predicted correlation between the high concentration indicator variable and competitive concerns of DG Comp over time setting all other covariates to their mean (dark blue), respectively median (light blue), value. The conditional average treatment effect predicted by the causal forest is 0.14, which is slightly higher than the coefficient on the high concentration indicator in specification 4 in Table 3.6. Compared to the patterns obtained based on the OLS estimates reported in Figure 3.2, the estimated effect of high concentration obtained with the causal forest is much smoother over time. This indicates that, once we use a richer model that better describes the process behind DG Comp's decisions, the impact of this structural indicator is less volatile and much more consistent over time.

Figure 3.6: Effect of High Concentration on Concerns over Time

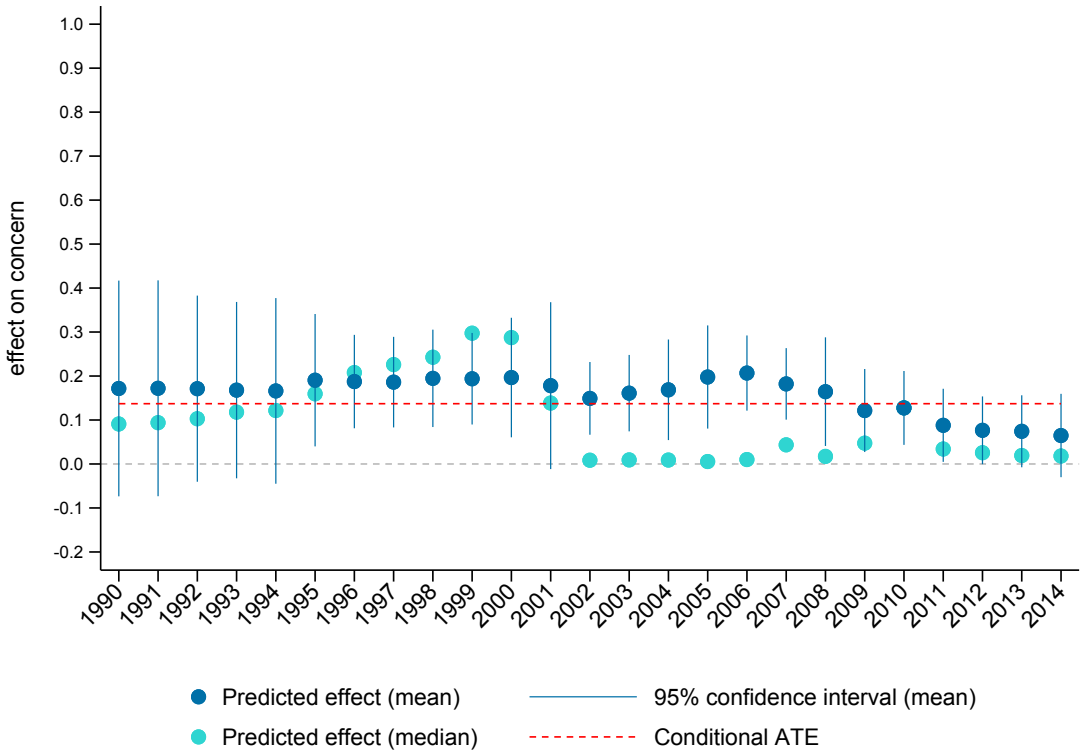

Predicted effect of indicator variable for post-merger HHI above 2000 and change in HHI larger than 150 on concerns over time, setting all other included explanatory variables equal to the sample mean/median.

Nonetheless, the importance of concentration appears to follow a downward trend over the years. The correlation between concentration and concerns is positive and mostly significant up to 2001, it seems to decrease since then and becomes insignificant in 2011. For the predicted correlation setting all other covariates to

median rather than mean values, the drop in correlation in 2001/2002 is even more pronounced and insignificant as of 2001.

3.5.2.2 Treatment - Joint Market Share above 50%

Figure 3.7 shows the predicted correlation between the indicator variable for merging parties' market shares above 50% and competitive concerns of DG Comp over time, as before setting all other covariates to their mean (dark blue), respectively median (light blue), value. The conditional average treatment effect predicted by the causal forest is 0.22, which is similar to the coefficient on the joint market share indicator in specification 4 in Table 3.6.

Figure 3.7: Effect of Joint Market Share on Concerns over Time

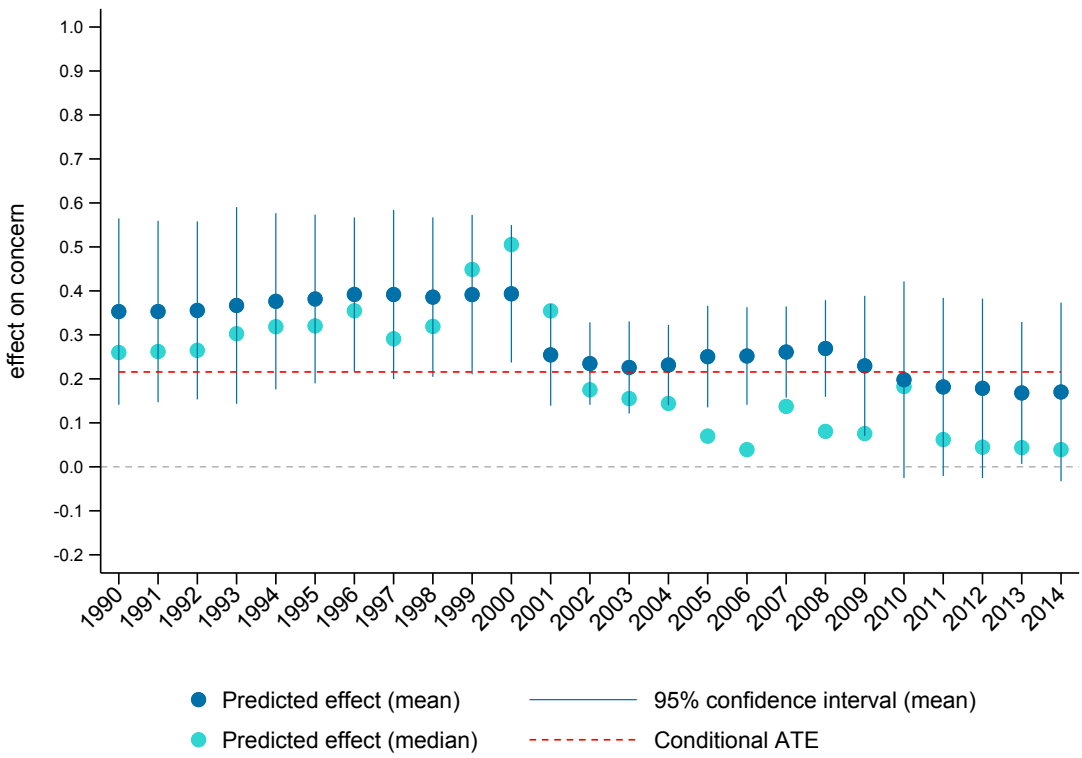

Predicted effect of indicator variable for joint market share above 50% on concerns over time, setting all other included explanatory variables equal to the sample mean/median.

Again, we find considerable heterogeneity in the predicted correlation between the market share indicator and concerns over time. While the predicted correlation is positive and significant up until 2010 (at least setting all other covariates to their mean), market shares seem to become a less important intervention decision criterion since the early 2000s and even become insignificant as of 2011. For

the predicted correlation setting all other covariates to median rather than mean values, the predicted correlation is even lower and mostly insignificant since 2002. Notice again that, as for concentration, the correlations estimated by means of the causal forest seem to be much less volatile and more consistent over time than those estimated based on the simple linear probability model.

Putting the developments of the correlation between concentration and market share measures with the intervention decision by DG Comp together highlights the shift away from evaluating mergers based on structural indicators towards a more economics based approach.

3.5.2.3 Treatment - Barriers to Entry

Figure 3.8 shows the predicted correlation between the presence of entry barriers in the concerned market and competitive concerns of DG Comp over time, again setting all other covariates to their mean (dark blue), respectively median (light blue), value. The conditional average treatment effect predicted by the causal forest is 0.46, which is higher than the coefficient on the entry barrier indicator in any specification in Table 3.6.

Furthermore, there is considerable heterogeneity in the predicted correlation between the existence of entry barriers and competitive concerns over time. While the predicted correlation with concerns was essentially zero up to 1997, it becomes positive, significant, and of increasing importance since 1998. This development is also in line with the shift of DG Comp's merger policy toward a more economics based approach.

Figure 3.8: Effect of Barriers to Entry on Concerns over Time

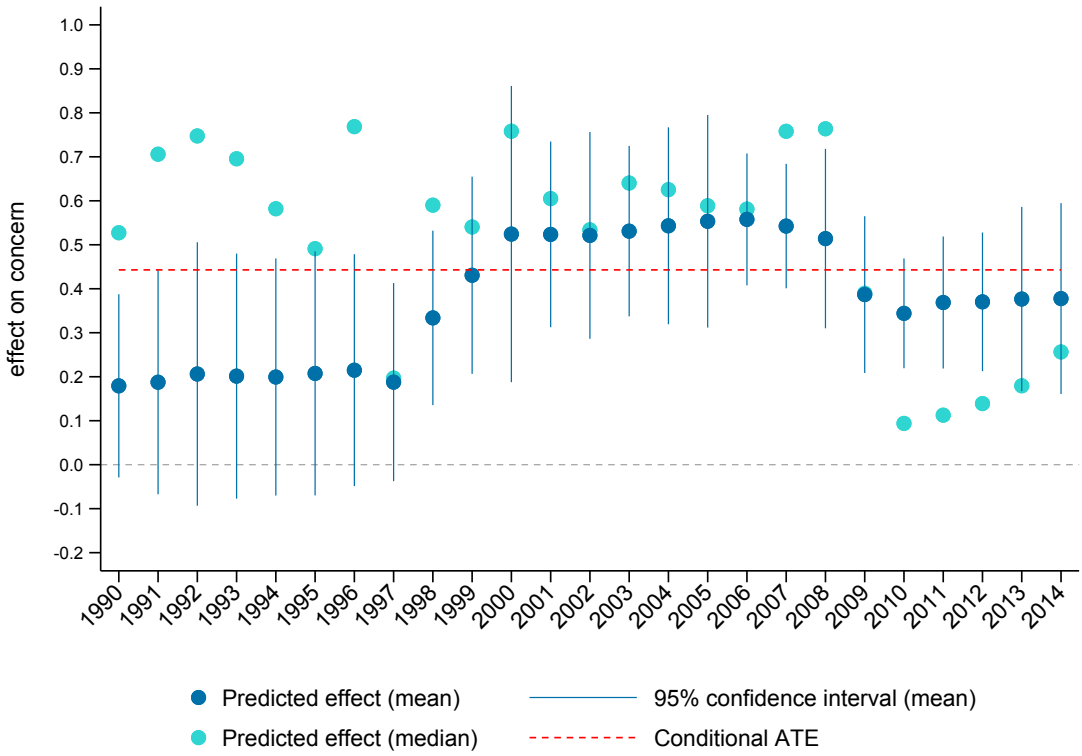

Predicted effect of barriers to entry on concerns over time, setting all other included explanatory variables equal to the sample mean/median.

3.5.2.4 Treatment - Risk of Foreclosure

Lastly, Figure 3.9 shows the predicted correlation between the indicator variable for risk of foreclosure in the concerned market and competitive concerns of DG Comp over time, setting all other covariates to their mean (dark blue), respectively median (light blue), value. The conditional average treatment effect predicted by the causal forest is 0.51, which is more than the double of the coefficient on the foreclosure indicator in the specifications in Table 3.6.

However, as shown in Table 3.2, DG Comp considered risk of foreclosure to exist in only about 3% of the concerned markets. Consequently, the confidence intervals for the predicted correlation are very wide, especially in the early years with fewer merger cases, and no clear pattern for the relationship between risk of foreclosure and competitive concerns emerges. However, there is a positive and mostly significant correlation that, if anything, seems to become more important over time.

Figure 3.9: Effect of Risk of Foreclosure on Concerns over Time

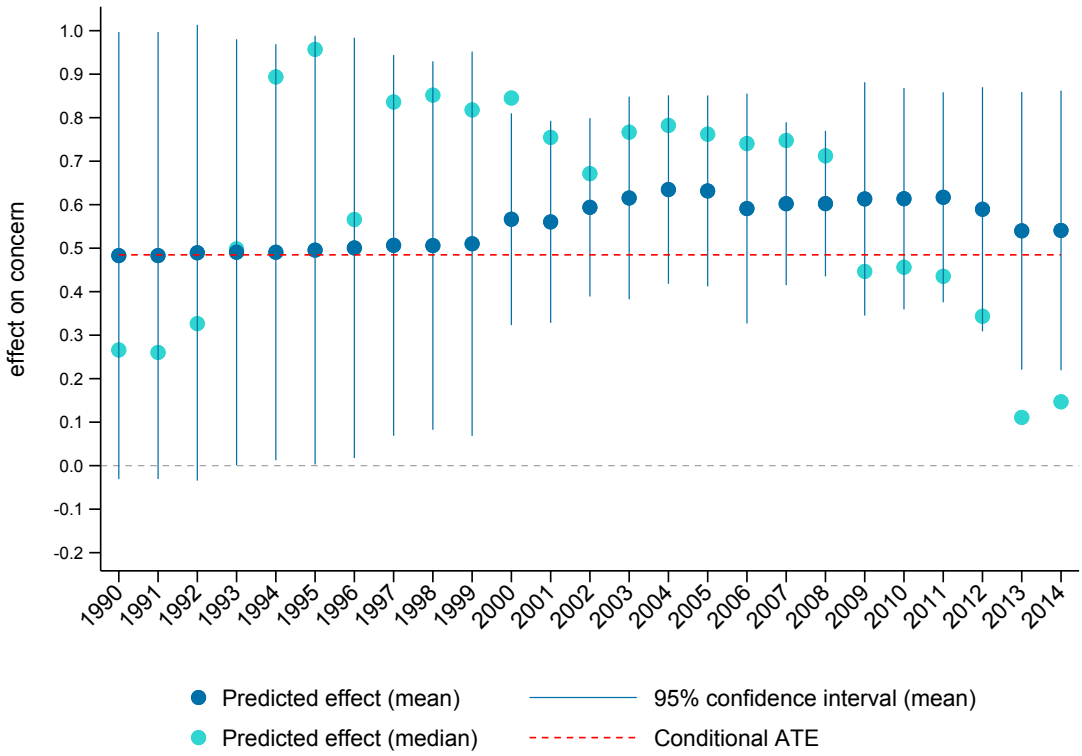

Predicted effect of risk of foreclosure on concerns over time, setting all other included explanatory variables equal to the sample mean/median.

3.6 Conclusion

In this paper, we study the time-dynamics of the EC's merger decision procedure over the first 25 years of European merger control using a new dataset containing all merger cases with an official decision documented by DG Comp (more than 5000 individual decisions). Specifically, we evaluate how consistently different arguments related to the structural market parameters – market shares, concentration, likelihood of entry, and foreclosure – are put forward to motivate a particular decision over time.

In a first step, and in line with the existing literature, we start by estimating the probability of intervention as a function of merger characteristics at the merger level. We find that the existence of barriers to entry, the increase of concentration measures and, in particular, the share of product markets with competitive concerns increase the likelihood of an intervention.

In order to obtain a more fine-grained picture of the decision determinants, we extend our analysis to the specific product and geographic markets concerned by

a merger. Instead of estimating the overall probability of an intervention, we estimate the likelihood that competitive concerns are found in that specific product/geographical market (our data contain more than 30,000 affected markets). This step is particularly important because larger mergers typically affect many different product markets in many different geographic regions. Therefore, by analyzing individual markets we not only get more statistical power but we are also able to conduct a more disaggregate analysis. We find that more determinants significantly affect the Commission's competitive concerns at the market level than seen at the merger level. Thus, the aggregation to – and the analysis at – the merger level hides some of the EC's more fine-grained considerations concerning specific markets. We find that, again, barriers to entry, but also the risk of foreclosure play an important role for the competitive analysis. Moreover, while tightly defined (national) markets increase the probability of concerns, the number of active competitors decreases it. Finally, structural indicators of market shares and concentration have the expected effects, which are however more relevant than in the merger-level analysis.

After this static analysis, we assess how the impact of these key determinants changes over time. We generally find that the importance of market shares and concentration seems to have declined over time. However, the parametric estimations are quite volatile and do not allow for uncovering clear patterns over time.

In the final step, we use non-parametric prediction methods, in particular the causal forest algorithm proposed by Athey and Imbens (2016), to more precisely explore how the correlation between the structural market parameters and competitive concerns varies with all other merger and market characteristics. Predicting the relationship between one structural market parameter and competitive concerns over time using the trained causal forests and holding all other merger and market characteristics constant, allows us to uncover clearer patterns over time. In particular, we find that concentration as well as the merging parties' market shares have become less important decision determinants over time and are even insignificant in most recent years. On the other hand, the importance of barriers to entry as well as the risk of foreclosure have increased in DG Comp's merger assessment since the early 2000s. This is in line with the goals of the 2004 merger policy reform, which aimed at adopting a more economics based approach of merger assessment and, consequently, putting less weight on simple structural indicators, such as HHI and market share.

3.7 Appendix

3.7.1 Regression Results OLS Concern over Time

Table 3.7: Linear Probability Model for Concern by Notification Year

	1990-1994	1995	1996	1997	1998	1999	2000
Barriers to	0.253**	0.730***	0.788***	-0.211***	0.499***	0.365***	0.395***
entry in submarket	(0.107)	(0.063)	(0.212)	(0.051)	(0.112)	(0.078)	(0.111)
Risk of			-0.017	0.693***		0.300***	0.613***
foreclosure in submarket			(0.111)	(0.091)		(0.083)	(0.098)
Joint market	0.015	0.137	0.383***	0.262**	0.155**	0.341***	0.411***
share above 50%	(0.075)	(0.091)	(0.099)	(0.093)	(0.072)	(0.051)	(0.077)
HHI $\geq$ 2000	0.076	0.079	-0.196**	0.081*	0.208	0.183***	0.149**
& Delta HHI $\geq$ 150	(0.066)	(0.048)	(0.068)	(0.039)	(0.155)	(0.038)	(0.066)
Fullmerger	-0.062	0.070	0.261	-0.176**	0.004	-0.067	-0.062
	(0.122)	(0.074)	(0.185)	(0.066)	(0.147)	(0.129)	(0.111)
Joint Venture	-0.201***	0.046	0.096	-0.268***	0.042	-0.088	-0.152*
	(0.067)	(0.067)	(0.119)	(0.055)	(0.160)	(0.130)	(0.088)
Conglomerate	0.074	0.066	1.098	0.057	-0.310*	-0.027	0.093***
merger in submarket	(0.116)	(0.038)	(0.810)	(0.045)	(0.157)	(0.050)	(0.024)
Vertical merger	-0.196**	0.012	-0.376*	0.237	0.067	0.010	-0.027
in submarket	(0.082)	(0.020)	(0.208)	(0.165)	(0.083)	(0.047)	(0.045)
Market	0.100*	0.516*	0.160	0.019	0.261*	0.065	0.050
definition national	(0.049)	(0.270)	(0.196)	(0.065)	(0.139)	(0.040)	(0.188)
Market	0.026	0.501*	0.233	0.188**	0.217	0.074**	-0.015
definition EU wide	(0.067)	(0.272)	(0.190)	(0.063)	(0.153)	(0.030)	(0.195)
Market	0.391	0.367*	0.160	0.138	0.430**	0.060	0.075
definition worldwide	(0.250)	(0.201)	(0.196)	(0.126)	(0.171)	(0.068)	(0.191)
Number of	-0.012**	-0.004	-0.009	0.002	-0.001	0.001	-0.001
concerned markets	(0.005)	(0.003)	(0.010)	(0.004)	(0.004)	(0.001)	(0.001)
Number of	-0.003	-0.002	0.020**	-0.019	-0.004	-0.005	0.022
competitors	(0.010)	(0.018)	(0.007)	(0.015)	(0.016)	(0.011)	(0.017)
Indicator no	-0.040	-0.069	0.141***	0.014	0.070	-0.045	0.076*
info on competitors	(0.047)	(0.073)	(0.026)	(0.069)	(0.132)	(0.046)	(0.044)
Constant	0.495***	-0.482	-0.017	-0.080	-0.354	0.239	0.126
	(0.097)	(0.292)	(0.094)	(0.083)	(0.312)	(0.161)	(0.157)
Industry Group FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R2	0.515	0.687	0.591	0.632	0.636	0.592	0.612
Observations	205	137	155	242	204	520	887

We report heteroskedasticity robust standard errors clustered at the industry group level. Significance at the 1%, 5%, and 10% levels is represented by ***,** and * respectively.

Table 3.8: Linear Probability Model for Concern by Notification Year (Continued)

	2001	2002	2003	2004	2005	2006	2007
Barriers to	0.241***	0.299**	0.328***	0.226**	0.326**	0.392***	0.366*
entry in submarket	(0.085)	(0.134)	(0.086)	(0.103)	(0.126)	(0.072)	(0.197)
Risk of	-0.043	0.060	-0.037	0.234	0.406***	0.131	0.241
foreclosure in submarket	(0.085)	(0.147)	(0.062)	(0.264)	(0.116)	(0.224)	(0.301)
Joint market	0.176***	0.181***	0.210**	0.246***	0.191***	0.143	0.356***
share above 50%	(0.038)	(0.058)	(0.084)	(0.049)	(0.058)	(0.086)	(0.084)
HHI $\geq$ 2000	0.111**	-0.015	0.205***	0.125*	0.108***	0.162*	0.093**
& Delta HHI $\geq$ 150	(0.044)	(0.042)	(0.069)	(0.070)	(0.036)	(0.090)	(0.039)
Fullmerger	0.118*	-0.006	-0.181	0.190**	-0.173**	-0.141**	-0.105
	(0.063)	(0.044)	(0.115)	(0.089)	(0.069)	(0.054)	(0.064)
Joint Venture	0.083	0.027	-0.151	0.445*	-0.208**	-0.231**	-0.127**
	(0.055)	(0.046)	(0.156)	(0.219)	(0.075)	(0.104)	(0.050)
Conglomerate	-0.085*	-0.195	-0.001	-0.393***		-0.001	-0.119
merger in submarket	(0.048)	(0.131)	(0.060)	(0.072)		(0.098)	(0.079)
Vertical merger	0.078	-0.015	-0.009	-0.226***	-0.075*	0.227**	-0.020
in submarket	(0.055)	(0.058)	(0.055)	(0.074)	(0.039)	(0.086)	(0.053)
Market	0.208**	-0.188*	0.270	0.032	-0.043	0.024	-0.007
definition national	(0.082)	(0.092)	(0.246)	(0.069)	(0.091)	(0.112)	(0.104)
Market	0.129**	-0.280***	0.226	-0.090	0.049	-0.066	0.011
definition EU wide	(0.049)	(0.094)	(0.241)	(0.065)	(0.078)	(0.118)	(0.100)
Market	0.299**	-0.201*	0.321		0.093	-0.003	-0.051
definition worldwide	(0.133)	(0.116)	(0.220)		(0.089)	(0.115)	(0.088)
Number of	0.001	-0.001	0.000	-0.004	0.002	0.000	-0.000
concerned markets	(0.001)	(0.000)	(0.001)	(0.002)	(0.002)	(0.000)	(0.000)
Number of	0.001	0.006	-0.002	-0.052**	-0.012	-0.009	-0.009
competitors	(0.017)	(0.011)	(0.021)	(0.021)	(0.010)	(0.014)	(0.006)
Indicator no	-0.049	-0.036	0.000	-0.363***	0.020	-0.131*	0.013
info on competitors	(0.061)	(0.113)	(0.085)	(0.093)	(0.047)	(0.064)	(0.045)
Constant	-0.316***	0.260	-0.058	0.308*	0.039	0.051	0.040
	(0.108)	(0.170)	(0.353)	(0.152)	(0.121)	(0.150)	(0.120)
Industry Group FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R2	0.698	0.403	0.508	0.483	0.446	0.547	0.445
Observations	774	569	494	546	1,209	1,408	1,423

We report heteroskedasticity robust standard errors clustered at the industry group level. Significance at the 1%, 5%, and 10% levels is represented by ***,** and * respectively.

Table 3.9: Linear Probability Model for Concern by Notification Year (Continued)

	2008	2009	2010	2011	2012	2013	2014
Barriers to	0.397***	0.435***	-0.083*	0.000	0.058***	0.113***	1.000***
entry in submarket	(0.110)	(0.081)	(0.042)	(.)	(0.016)	(0.007)	(0.000)
Risk of	0.046	0.419*	0.930***			0.065	
foreclosure in submarket	(0.335)	(0.239)	(0.108)			(0.048)	
Joint market	0.281***	0.142***	0.049*	0.000	0.109*	0.080***	0.000
share above 50%	(0.063)	(0.041)	(0.026)	(.)	(0.059)	(0.021)	(0.000)
HHI $\geq$ 2000	0.041	0.131	0.072	0.000	-0.079*	-0.004	0.000
& Delta HHI $\geq$ 150	(0.032)	(0.076)	(0.043)	(.)	(0.045)	(0.009)	(0.000)
Fullmerger	0.041	0.014	0.050***	0.000	0.044	-0.039	0.000
	(0.101)	(0.031)	(0.014)	(.)	(0.038)	(0.036)	(0.000)
Joint Venture	-0.038	0.024	-0.025	0.000	0.088*	0.004	
	(0.110)	(0.051)	(0.034)	(.)	(0.048)	(0.005)	
Conglomerate	0.052	-0.453*					
merger in submarket	(0.130)	(0.225)					
Vertical merger	-0.009	-0.026	-0.115	0.000	0.060	-0.008	-0.000
in submarket	(0.031)	(0.096)	(0.071)	(.)	(0.060)	(0.007)	(0.000)
Market	0.154***	0.042	0.331***	0.000	0.001	-0.010	0.000
definition national	(0.046)	(0.049)	(0.038)	(.)	(0.006)	(0.009)	(0.000)
Market	0.014	0.115**	0.250***	0.000	-0.201	0.003	0.000
definition EU wide	(0.046)	(0.041)	(0.084)	(.)	(0.117)	(0.013)	(0.000)
Market	-0.045	0.092*	0.196**	0.000	-0.088		
definition worldwide	(0.032)	(0.050)	(0.072)	(.)	(0.064)		
Number of	-0.001	0.001	0.001	0.000	0.002***	0.000	-0.000
concerned markets	(0.001)	(0.000)	(0.001)	(.)	(0.000)	(0.000)	(0.000)
Number of	-0.008	-0.004	0.003	0.000	-0.013	0.003	0.000
competitors	(0.007)	(0.010)	(0.003)	(.)	(0.012)	(0.002)	(0.000)
Indicator no	-0.003	-0.091*	0.027	0.000	-0.099	0.002	-0.000
info on competitors	(0.038)	(0.047)	(0.026)	(.)	(0.083)	(0.006)	(0.000)
Constant	0.274**	0.014	0.044	0.000	0.011	-0.010	-0.000
	(0.103)	(0.099)	(0.079)	(.)	(0.063)	(0.014)	(0.000)
Industry Group FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R2	0.496	0.415	0.542	.	0.468	0.122	1.000
Observations	1,534	761	411	179	519	595	38

We report heteroskedasticity robust standard errors clustered at the industry group level. Significance at the 1%, 5%, and 10% levels is represented by ***,** and * respectively.

3.7.2 Determinants of Concern - Market Level - Split Sample over Industries

Table 3.10: Linear Probability Model for Concern by Industry

	Group 1	Group 2	Group 3	Group 4	Group 5	Group 6	Group 7
Barriers to	0.412***	0.071	1.000***	0.637***	0.241***	0.487***	0.403***
entry in submarket	(0.070)	(0.067)	(0.000)	(0.054)	(0.032)	(0.038)	(0.095)
Risk of	0.326***	0.659***			0.469***		-0.364
foreclosure in submarket	(0.113)	(0.147)			(0.055)		(0.260)
Joint market	0.415***	0.329***		0.217***	0.265***	0.301***	0.302***
share above 50%	(0.047)	(0.029)		(0.046)	(0.022)	(0.028)	(0.061)
HHI $\geq$ 2000	0.135***	0.079***	-0.000	0.066*	0.076***	0.177***	0.072**
& Delta HHI $\geq$ 150	(0.029)	(0.020)	(0.000)	(0.034)	(0.017)	(0.029)	(0.031)
Fullmerger	0.068	0.153***	0.000	-0.223***	-0.067***	0.121***	-0.228***
	(0.053)	(0.026)	(0.000)	(0.051)	(0.025)	(0.043)	(0.073)
Joint Venture	-0.006	0.060**		0.089	-0.150***	-0.093*	-0.280***
	(0.054)	(0.030)		(0.101)	(0.034)	(0.056)	(0.079)
Conglomerate		-0.087*			-0.185***	0.355***	
merger in submarket		(0.048)			(0.069)	(0.075)	
Vertical merger	0.021	-0.042	-0.000	-0.009	-0.010	0.022	0.042
in submarket	(0.040)	(0.026)	(0.000)	(0.055)	(0.021)	(0.046)	(0.040)
Market	0.201**	0.043	0.000	0.148**	0.011	-0.244***	0.444**
definition national	(0.091)	(0.062)	(0.000)	(0.073)	(0.059)	(0.057)	(0.178)
Market	0.157*	0.045		0.106	-0.047	-0.171**	0.431**
definition EU wide	(0.089)	(0.066)		(0.068)	(0.057)	(0.069)	(0.173)
Market	0.157*	0.033		0.219	-0.002	-0.198***	0.348*
definition worldwide	(0.081)	(0.100)		(0.207)	(0.060)	(0.072)	(0.196)
Number of	-0.000	0.000	-0.000	0.002*	-0.001***	-0.001	-0.000
concerned markets	(0.001)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)
Number of	-0.004	0.007	-0.000	0.003	-0.006	-0.021***	0.024*
competitors	(0.006)	(0.008)	(0.000)	(0.011)	(0.006)	(0.005)	(0.013)
Indicator no	-0.061	-0.026		-0.123*	-0.089***	-0.061**	0.114*
info on competitors	(0.037)	(0.033)		(0.066)	(0.025)	(0.027)	(0.058)
Post reform	0.093	0.052	0.000	-0.715***	-0.865***	-0.103	-0.067**
indicator	(0.085)	(0.052)	(0.000)	(0.179)	(0.037)	(0.108)	(0.033)
Constant	-0.213*	-0.227***	0.000	0.485**	1.010***	0.218*	-0.294
	(0.123)	(0.087)	(0.000)	(0.198)	(0.070)	(0.129)	(0.205)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R2	0.671	0.409	1.000	0.586	0.507	0.483	0.577
Observations	455	1,022	39	435	1,919	1,035	339

We report heteroskedasticity robust standard errors.
Significance at the 1%, 5%, and 10% levels is represented by ***,** and * respectively.

Table 3.11: Linear Probability Model for Concern by Industry (Continued)

	Group 8	Group 9	Group 10	Group 11	Group 12	Group 13
Barriers to	0.066	0.467***	0.681***	0.268***	0.407***	0.328***
entry in submarket	(0.157)	(0.057)	(0.072)	(0.078)	(0.055)	(0.077)
Risk of	0.213*	0.502***	-0.322**	0.510***	-0.047	0.408***
foreclosure in submarket	(0.118)	(0.103)	(0.125)	(0.088)	(0.044)	(0.117)
Joint market	0.215***	0.155***	0.146**	0.132***	0.171***	0.187***
share above 50%	(0.061)	(0.036)	(0.057)	(0.031)	(0.036)	(0.050)
HHI $\geq$ 2000	0.057**	0.081***	-0.016	0.106***	-0.037	0.028
& Delta HHI $\geq$ 150	(0.027)	(0.019)	(0.020)	(0.020)	(0.035)	(0.018)
Fullmerger	0.058	-0.200***	-0.158***	-0.219***	-0.114**	0.061*
	(0.055)	(0.044)	(0.052)	(0.036)	(0.045)	(0.032)
Joint Venture	0.002	-0.218***	-0.126**	-0.213***		0.019
	(0.060)	(0.056)	(0.057)	(0.035)		(0.037)
Conglomerate	0.265*	-0.156***	0.022	-0.131	-0.016	-0.059*
merger in submarket	(0.143)	(0.057)	(0.032)	(0.096)	(0.040)	(0.036)
Vertical merger	-0.080***	0.005	0.031	-0.039**	-0.030	-0.050
in submarket	(0.028)	(0.019)	(0.029)	(0.016)	(0.033)	(0.031)
Market	0.178*	0.025	0.294***	0.078	0.182**	0.075*
definition national	(0.094)	(0.105)	(0.095)	(0.075)	(0.074)	(0.043)
Market	0.201**	0.087	0.132*	0.072	0.091	0.039
definition EU wide	(0.096)	(0.104)	(0.074)	(0.073)	(0.066)	(0.028)
Market	0.242**	0.062	0.079	0.149*		0.068
definition worldwide	(0.095)	(0.103)	(0.081)	(0.076)		(0.051)
Number of	-0.003***	0.005***	-0.003***	0.001	-0.001***	0.001
concerned markets	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)	(0.001)
Number of	0.002	-0.019**	-0.005	0.003	-0.009	-0.004
competitors	(0.006)	(0.008)	(0.005)	(0.005)	(0.006)	(0.008)
Indicator no	0.042	-0.145***	-0.109***	0.052*	-0.007	-0.046
info on competitors	(0.038)	(0.037)	(0.040)	(0.028)	(0.055)	(0.039)
Post reform	0.101	-0.109**	-0.351***	-0.021	0.632***	-0.028
indicator	(0.091)	(0.055)	(0.110)	(0.026)	(0.087)	(0.023)
Constant	-0.331***	0.079	0.240*	-0.109	0.053	-0.141*
	(0.124)	(0.119)	(0.129)	(0.082)	(0.042)	(0.079)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
R2	0.392	0.644	0.793	0.522	0.385	0.453
Observations	369	621	339	632	443	435

We report heteroskedasticity robust standard errors.
Significance at the 1%, 5%, and 10% levels is represented by ***,** and * respectively.

Table 3.12: Linear Probability Model for Concern by Industry (Continued)

	Group 14	Group 15	Group 16	Group 17	Group 18	Group 19
Barriers to	0.406***	0.000	0.346***	0.199***		0.581***
entry in submarket	(0.069)	(.)	(0.054)	(0.028)		(0.119)
Risk of	0.046		0.269***	-0.027		0.131
foreclosure in submarket	(0.066)		(0.104)	(0.040)		(0.174)
Joint market	0.253***	0.000	0.071	0.113***	0.000	0.221***
share above 50%	(0.048)	(.)	(0.045)	(0.020)	(.)	(0.052)
HHI $\geq$ 2000	0.205***	0.000	0.134***	0.197***	0.000	0.083***
& Delta HHI $\geq$ 150	(0.036)	(.)	(0.020)	(0.028)	(.)	(0.027)
Fullmerger	-0.297***	0.000	-0.120***	-0.029	0.000	0.171
	(0.064)	(.)	(0.036)	(0.087)	(.)	(0.115)
Joint Venture	-0.372***	0.000	-0.084**	0.003	0.000	0.155**
	(0.064)	(.)	(0.036)	(0.093)	(.)	(0.066)
Conglomerate		0.000	-0.025	0.130**		0.018
merger in submarket		(.)	(0.037)	(0.063)		(0.086)
Vertical merger	0.047	0.000	0.006	0.037	0.000	0.003
in submarket	(0.038)	(.)	(0.015)	(0.028)	(.)	(0.032)
Market	0.004	0.000	-0.026	0.092*		-0.004
definition national	(0.061)	(.)	(0.023)	(0.048)		(0.177)
Market	-0.166**	0.000	0.014	0.062		0.003
definition EU wide	(0.078)	(.)	(0.024)	(0.059)		(0.175)
Market		0.000	0.070*	0.052		-0.045
definition worldwide		(.)	(0.036)	(0.055)		(0.166)
Number of	-0.000	0.000	-0.001	0.000	0.000	-0.001
concerned markets	(0.001)	(.)	(0.001)	(0.000)	(.)	(0.001)
Number of	0.002	0.000	0.006*	-0.028***	0.000	0.025***
competitors	(0.006)	(.)	(0.004)	(0.006)	(.)	(0.009)
Indicator no	0.009	0.000	0.088***	-0.108***	0.000	0.076*
info on competitors	(0.035)	(.)	(0.022)	(0.034)	(.)	(0.039)
Post reform	0.106*	0.000	0.038***	-0.121	0.000	-0.185
indicator	(0.057)	(.)	(0.012)	(0.078)	(.)	(0.166)
Constant	0.212**	0.000	-0.034	0.128	0.000	-0.319
	(0.097)	(.)	(0.048)	(0.127)	(.)	(0.207)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
R2	0.657	.	0.548	0.326	.	0.640
Observations	547	85	680	1,398	60	420

We report heteroskedasticity robust standard errors.
Significance at the 1%, 5%, and 10% levels is represented by ***,** and * respectively.

Table 3.13: Linear Probability Model for Concern by Industry (Continued)

	Group 20	Group 21	Group 22	Group 23	Group 24	Group 25
Barriers to	0.362***	0.974***	0.215	0.178**		0.751***
entry in submarket	(0.062)	(0.042)	(0.147)	(0.082)		(0.194)
Risk of	-0.283***	0.957***		-0.274**		0.980***
foreclosure in submarket	(0.085)	(0.044)		(0.123)		(0.044)
Joint market	0.025	0.191	0.233***	0.268***	0.000	-0.021
share above 50%	(0.022)	(0.124)	(0.078)	(0.078)	(0.000)	(0.038)
HHI $\geq$ 2000	0.076***	-0.008	0.026	0.204***	0.000	0.079*
& Delta HHI $\geq$ 150	(0.024)	(0.012)	(0.052)	(0.043)	(0.000)	(0.041)
Fullmerger	0.082***	-0.002	0.057	0.267	-1.000***	0.124
	(0.027)	(0.014)	(0.052)	(0.168)	(0.000)	(0.140)
Joint Venture	-0.083	-0.031	-0.022	0.302*	-1.000***	
	(0.063)	(0.025)	(0.067)	(0.178)	(0.000)	
Conglomerate	0.145		-0.001	-0.141		
merger in submarket	(0.134)		(0.067)	(0.132)		
Vertical merger	0.062	0.015	-0.097	0.103*	0.000	0.039
in submarket	(0.038)	(0.022)	(0.114)	(0.062)	(0.000)	(0.047)
Market	-0.033	-0.042	-0.214***	-0.227***	-0.000	-0.158
definition national	(0.079)	(0.032)	(0.072)	(0.047)	(0.000)	(0.124)
Market	-0.022	-0.033	-0.075	-0.281***	-0.000	-0.054
definition EU wide	(0.088)	(0.027)	(0.112)	(0.075)	(0.000)	(0.073)
Market	-0.032	-0.027	-0.224***	-0.187	-0.000	-0.169
definition worldwide	(0.088)	(0.023)	(0.083)	(0.121)	(0.000)	(0.134)
Number of	-0.003*	0.000	0.013***	-0.001	0.000	0.001
concerned markets	(0.002)	(0.000)	(0.004)	(0.001)	(0.000)	(0.001)
Number of	-0.004	-0.006	-0.026*	-0.011	0.000**	-0.089
competitors	(0.003)	(0.009)	(0.014)	(0.011)	(0.000)	(0.057)
Indicator no	-0.002	-0.021	-0.275***	0.093	0.000**	-0.356*
info on competitors	(0.024)	(0.045)	(0.082)	(0.073)	(0.000)	(0.203)
Post reform	-0.044	-0.027	-0.135	0.137	-0.000	-0.099
indicator	(0.090)	(0.024)	(0.143)	(0.181)	(0.000)	(0.094)
Constant	0.055	0.091	0.389**	0.020	1.000***	0.355*
	(0.181)	(0.083)	(0.171)	(0.184)	(0.000)	(0.203)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
R2	0.479	0.889	0.427	0.282	1.000	0.724
Observations	442	251	244	434	50	116

We report heteroskedasticity robust standard errors.
Significance at the 1%, 5%, and 10% levels is represented by ***,** and * respectively.

Figure 3.10: OLS Regression Coefficient on High Concentration over Industry

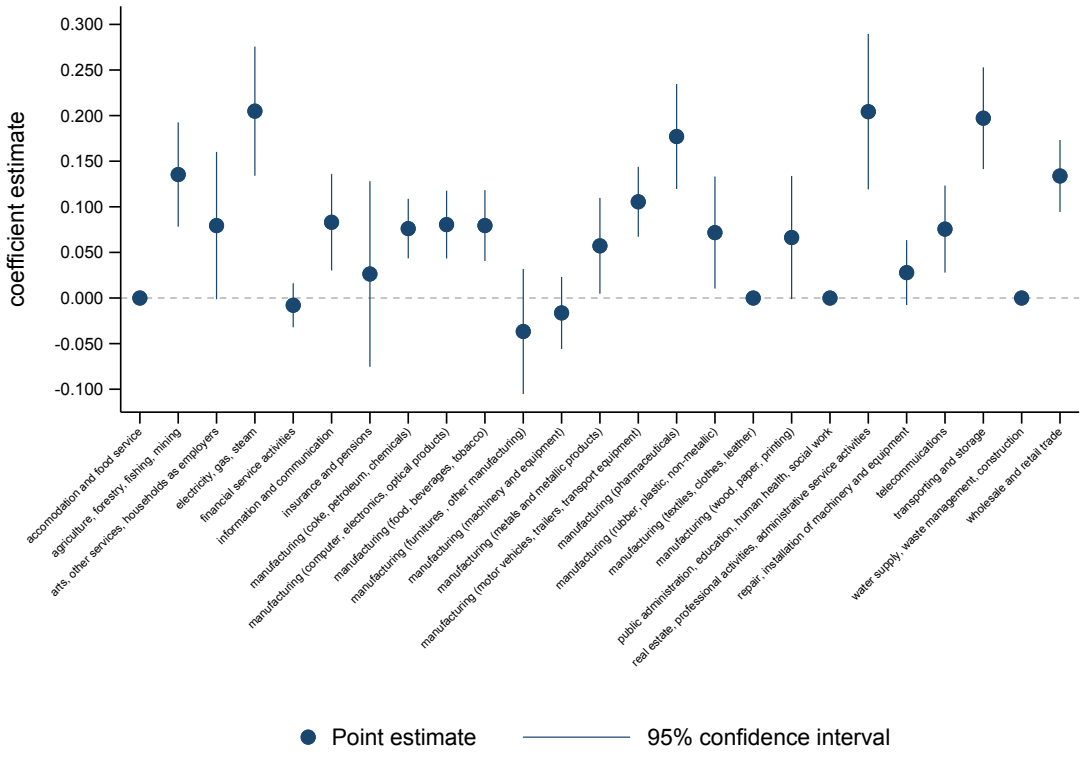

Regression coefficient on indicator variable for post-merger HHI above 2000 and change in HHI due to the merger larger than 150 in OLS regression on concerns. Each reported coefficient stems from a separate regression for the respective industry. Confidence intervals are based on heteroskedasticity robust standard errors.

Figure 3.11: OLS Regression Coefficient on Joint Market Share over Industry

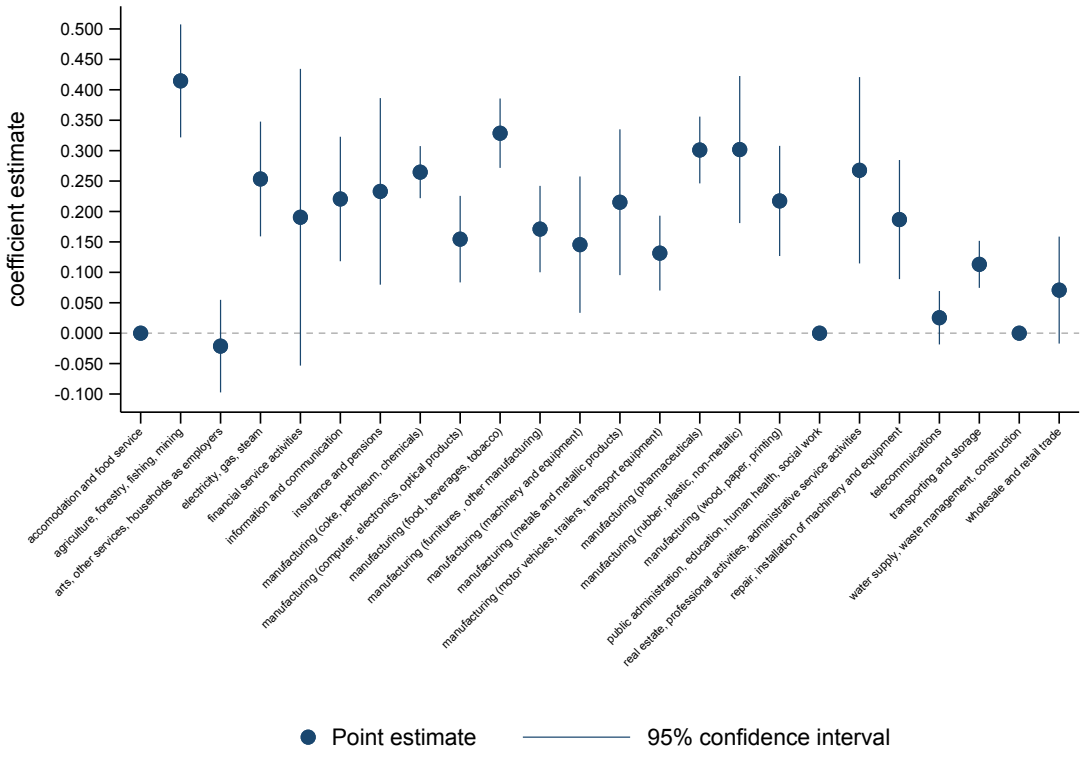

Regression coefficient on indicator variable for joint market share above 50% in OLS regression on concerns. Each reported coefficient stems from a separate regression for the respective industry. Confidence intervals are based on heteroskedasticity robust standard errors.

Figure 3.12: OLS Regression Coefficient on Barriers to Entry over Industry

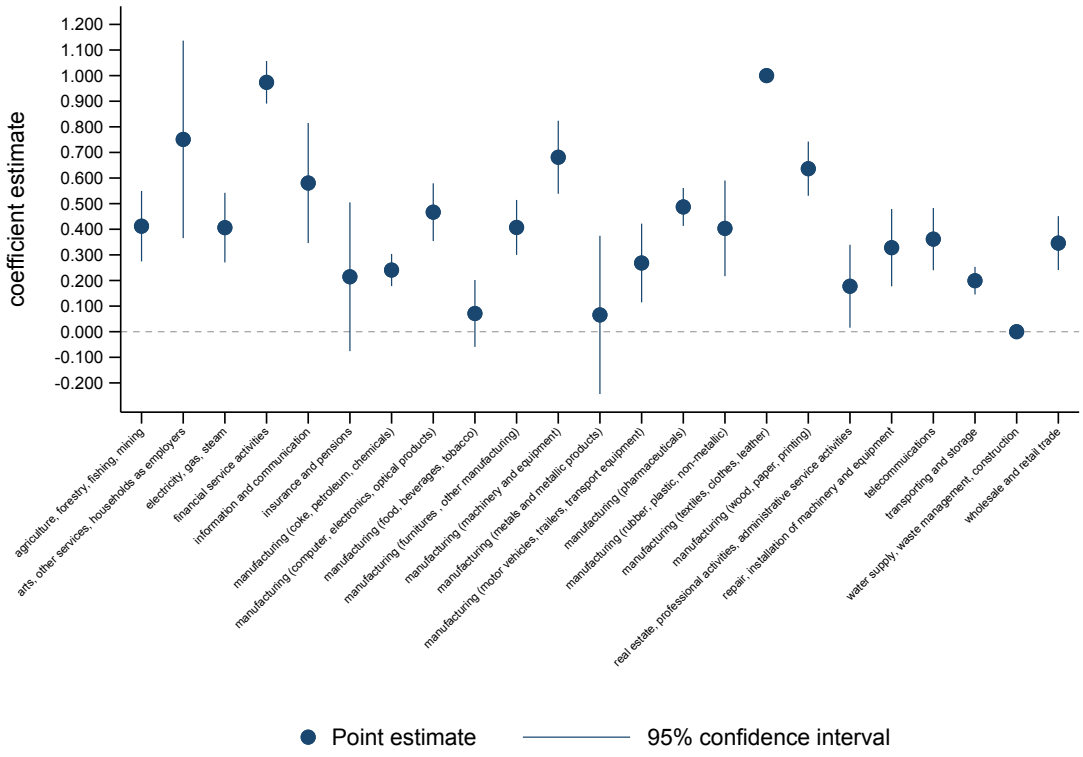

Regression coefficient on barriers to entry in OLS regression on concerns. Each reported coefficient stems from a separate regression for the respective industry. Confidence intervals are based on heteroskedasticity robust standard errors.

Figure 3.13: OLS Regression Coefficient on Risk of Foreclosure over Industry

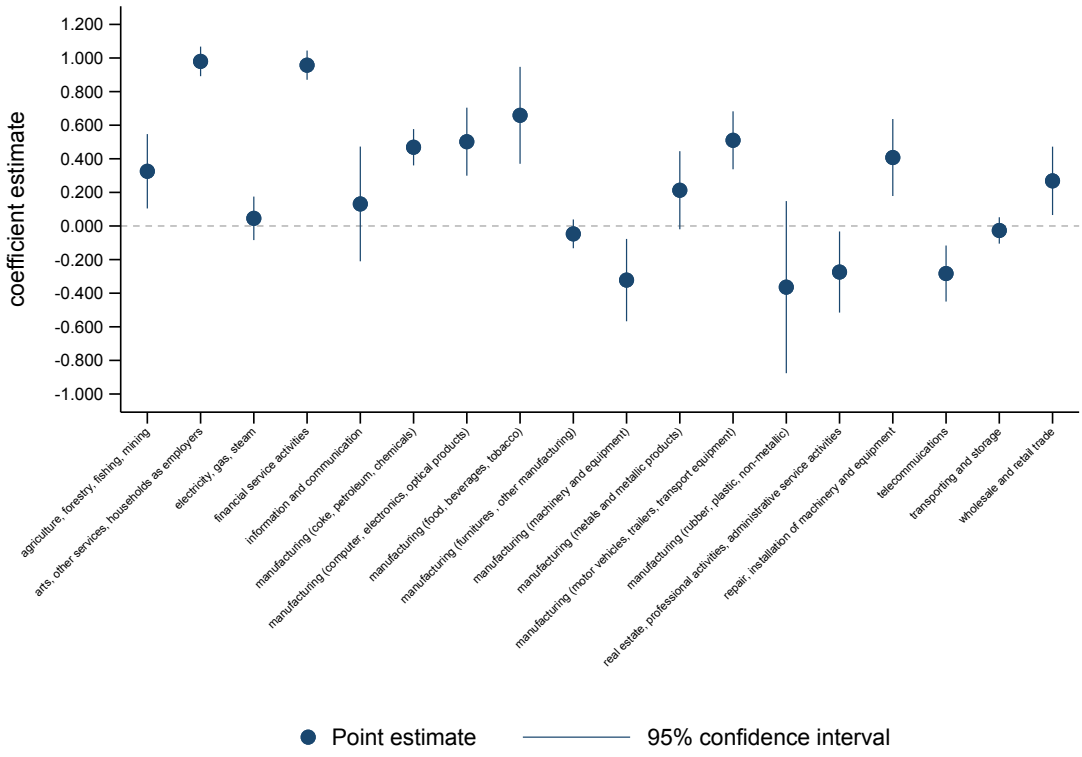

Regression coefficient on risk of foreclosure in OLS regression on concerns. Each reported coefficient stems from a separate regression for the respective industry. Confidence intervals are based on heteroskedasticity robust standard errors.

3.7.3 Technical Background on Causal Forests

3.7.3.1 Background on Causal Forests

Causal forests are based on the random forest methodology by Breiman (2001). They have been developed by Athey and co-authors in a series of papers (see Athey and Imbens (2016), Wager and Athey (2017) and Athey, Tibshirani, and Wager (2017)), extending the regression tree and random forest algorithms so as to estimate average treatment effects for different subgroups, rather than predicting outcomes as is the case for regression trees and random forests.

In a standard CART tree (Classification and Regression Tree), the goal is to predict individual outcomes Y_i using the mean outcome Y of observations that are "close" in X-space. To determine which observations are "close", the algorithm starts to recursively split the covariate space (binary splits) until it is partitioned into a set of so-called leaves L that contain only a few training samples. The outcome Y_i for observation i is then predicted by identifying the leaf containing observation i based on its characteristics X_i and setting the prediction to the mean outcome within that leaf:

$$\hat{\mu}(x) = \frac{1}{\mid \{i : X_i \in L(x)\} \mid} \sum_{\{i:X_i \in L(x)\}} Y_i \tag{3.5}$$

The algorithm automatically decides on the splitting variables and split points. This is done based on an in sample goodness-of-fit criterion (so essentially how close the predicted outcomes are to the actual outcomes). For regression trees (continuous outcome variable Y) the goodness-of-fit criterion used is the mean squared error, for classification trees (categorical outcome variable Y) the goodness-of-fit criterion is a measure of classification error based on the empirical classification probabilities in the leaves. The algorithm then splits on the covariate at the cut-off value that leads to the greatest improvement in the goodness-of-fit criterion. Once the best split at a given point in the tree is found, the splitting process is repeated in each of the resulting two regions. For CART trees, the splitting process is usually stopped when a specified minimum node size is reached - by default this is a node size of 5 for regression and 1 for classification trees. The tree is then pruned based on some cost-complexity trade-off measure in order to avoid over-fitting (See Hastie, Tibshirani, and Friedman (2008, chapter 9) for further details).

A random forest is then an ensemble of regression or classification trees, where the predictions are averaged across trees (for classification problems, the random forest obtains a class vote from each tree and then classifies based on majority vote). Each individual tree in the forest is grown using a random sample with replacement from the training set. One third of the data is not used for training and can be used for

testing (out-of-bag error). Differently from growing a single tree, splitting for each node in a tree in the forest is done based on only a subset of the covariates X and each tree is grown to the largest extent possible without pruning. The idea behind random forests is to reduce variance and produce more robust predictions compared to a single tree. The splitting on only a subset of variables at each node reduces the correlation between the trees in the forest and the variance of the predictions further (See Breiman (2001) and Hastie, Tibshirani, and Friedman (2008, chapter 15) for further details).

In case of a causal forest, we are not interested in predicting individual outcomes Y_i but individual treatment effects $Y_i^1 - Y_i^0$ to study how treatment effects vary by subgroup. This implies that standard fit measures used in regression trees and random forests, such as the mean squared error, are not available since one of the potential outcomes and hence the actual treatment effect is never observed. However, the causal forest methodology builds on regression tree methods in that it also applies a "goodness-of-fit" criterion in treatment effects to decide on splits. Athey and Imbens (2016) show that the mean squared error function of a causal tree can be estimated and is a function of the variance of the estimated treatment effect. Basically, the goodness-of-fit measure to be minimized rewards a partition of the data for finding strong heterogeneity in treatment effects and penalizes a partition for high variance in leaf estimates. Minimizing the expected mean squared error of predicted treatment effects (rather than the infeasible mean squared error), is shown to be equivalent to maximizing the variance of the predicted treatment effects across leaves with a penalty for within-leaf variance (variance of means of treatment and control group outcomes within leaves).

Causal trees are similar to nearest-neighbour methods as they also rely on the unconfoundedness assumption and use "close" observations to predict treatment effects. However, rather than defining closeness based on some pre-specified distance measure (such as Euclidean distance in k-nearest-neighbour matching), closeness is defined with respect to a decision tree and the closest control observations to i are those that fall in the same leaf. Analogously to CART regression trees, the leaves in causal trees should be small enough so that the (Y_i, W_i) pairs in a given leaf act as though they had come from a randomized experiment (Wager and Athey, 2017). The treatment effect for observation i with covariates $X_i = x$ falling into leaf L is then simply estimated as the difference of mean outcomes between treated and control observations within that leaf:

$$\hat{\tau}(x) = \frac{1}{|\{i : W_i = 1, X_i \in L\}|} \sum_{\{i:W_i=1,X_i\in L\}} Y_i - \frac{1}{|\{i : W_i = 0, X_i \in L\}|} \sum_{\{i:W_i=0,X_i\in L\}} Y_i$$

Given the procedure for generating a single causal tree, a causal forest then generates B such trees, each of which delivers an estimate $\hat{\tau}_b(x)$. The causal forest as developed by Wager and Athey (2017) then aggregates the predictions of the single trees by averaging:

$$\hat{\tau}(x) = \frac{1}{B} \sum_{b=1}^{B} \hat{\tau}_b(x) \tag{3.6}$$

The causal forest algorithm by Athey, Tibshirani, and Wager (2017) (the one we use here), predicts treatment effects slightly differently. For each observation i, the algorithm weights the nearby control observations according to the fraction of trees in which a control observation appears in the same leaf as the treated observation i. The treatment effect is then calculated as the difference between observation i's actual outcome and the weighted average outcome of its control observations. This implies that for each observation an individual treatment effect τ_i can be estimated.

As for CART trees and random forests, the advantage of a causal forest over a causal tree is that it is not always clear what the "best" causal tree is. The aggregation across trees helps to reduce variance, the estimates of the causal effects change more smoothly with covariates and individual treatment effects τ_i can be estimated while in a causal tree all individuals assigned to a given terminal leaf have the same estimated treatment effect (Wager and Athey, 2017).

Athey and Imbens (2016) further introduce so-called "honesty" in causal trees to ensure correct inference: the data is divided in half, where one half of the data is used to build the tree (so determine the splits in covariate space) and the other half is used to predict treatment effects. Wager and Athey (2017) extend this idea to causal forests and develop asymptotic theory for inference in causal forests. Thus, the causal forest algorithm by Athey, Tibshirani, and Wager (2017) does not only allow to predict heterogeneous treatment effects in a very flexible way but also provides confidence intervals for these estimates.

3.7.3.2 Background on grf package

We use the generalized random forest (grf) R package of Athey, Tibshirani, and Wager (2017). The package allows, among others, to train a causal forest, obtain the

conditional average treatment effect and predict treatment effects, either in-sample using out-of-bag training samples or out-of-sample using prediction datasets as we do in our application. As the package also predicts the variance of treatment effects, it is possible to compute point-wise confidence intervals for predicted treatment effects.

To build the trees in the forest, the package uses by default 50% of the data to grow each tree. When honesty is used, these sub-samples are further cut in half, where one half is used to place the splits within the tree and the other half is used to estimate treatment effects within the leaves.

While the causal forest algorithm is based on the regression tree methodology, to our understanding, it can still be applied to a binary outcome variable Y as is the case in our application. Athey, Tibshirani, and Wager (2017) apply the causal forest methodology themselves in the example of the effect of child rearing on female labor-force participation where the outcome variable is an indicator variable for whether the mother did not work in the year preceding the census.

In case of a binary outcome variable, the causal forest function gives estimates of $\tau(x) = \mathbf{E}\left[Y(1) - Y(0)|X = x\right]$ and according to a forum discussion on the grf package by the authors, the provided confidence intervals are also formally justified for binary Y as long as $Y(w)$ is not a deterministic function of X (i.e. there is still some randomness in the outcome Y given X and W). For binary outcome Y, the prediction function for causal forests then returns the estimated change in the probability of Y associated with the treatment W, which should be between -1 and 1.

3.7.4 Variable Importance Plots of Causal Forests

The variable importance measures the frequency with which the causal forest splits over a given covariate. It is based on the split frequencies function provided in the grf R package by Athey, Tibshirani, and Wager (2017) that shows how often the forest chose to split on each covariate at different split depths. For the plots shown here, we take into account splits within trees up to a split depth of 4. The variable importance function first counts the fraction of times the forest splits on each covariate at split levels 1, 2, 3 and 4. To calculate the overall variable importance measure, splits on a given covariate are weighted differently depending on the split depth. In the variable importance plots below, we use a decay exponent of 2, implying weights for splits at depth 1,2,3 and 4 of 1, 0.25, 0.1111 and 0.0625 respectively.

3.7.4.1 Treatment - High Concentration

Figure 3.14: Variable Importance Plot for Correlation between High Concentration and Concerns

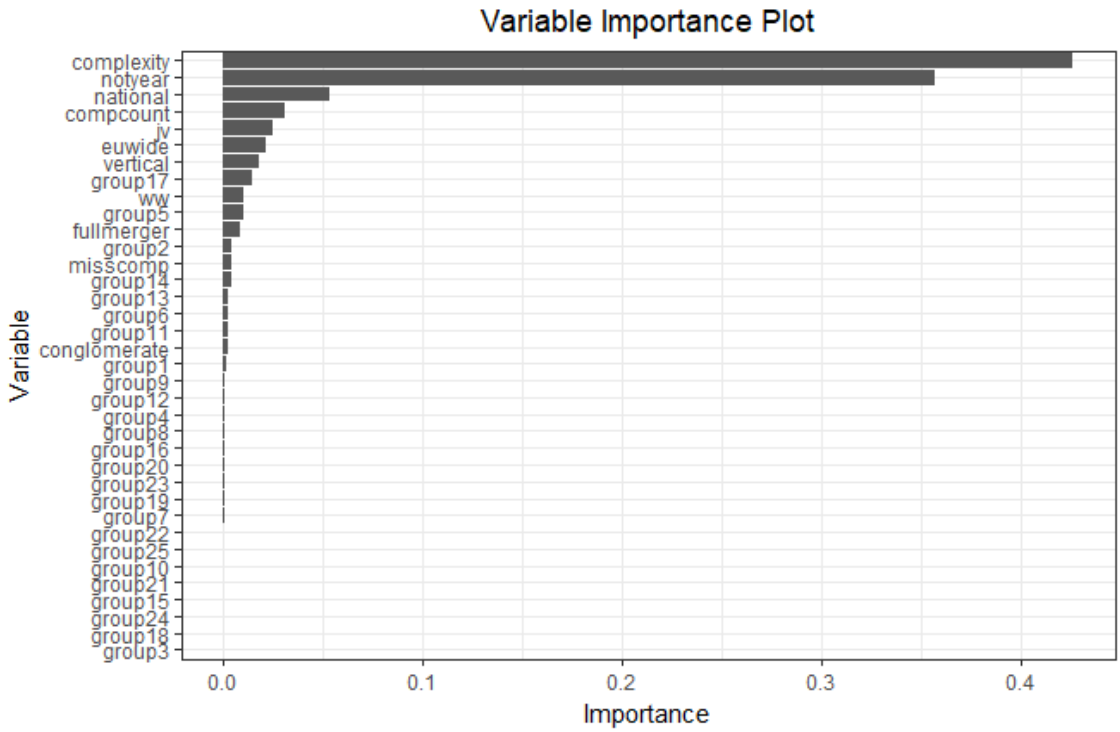

3.7.4.2 Treatment - Joint Market Share above 50%

Figure 3.15: Variable Importance Plot for Correlation between Joint Market Share and Concerns

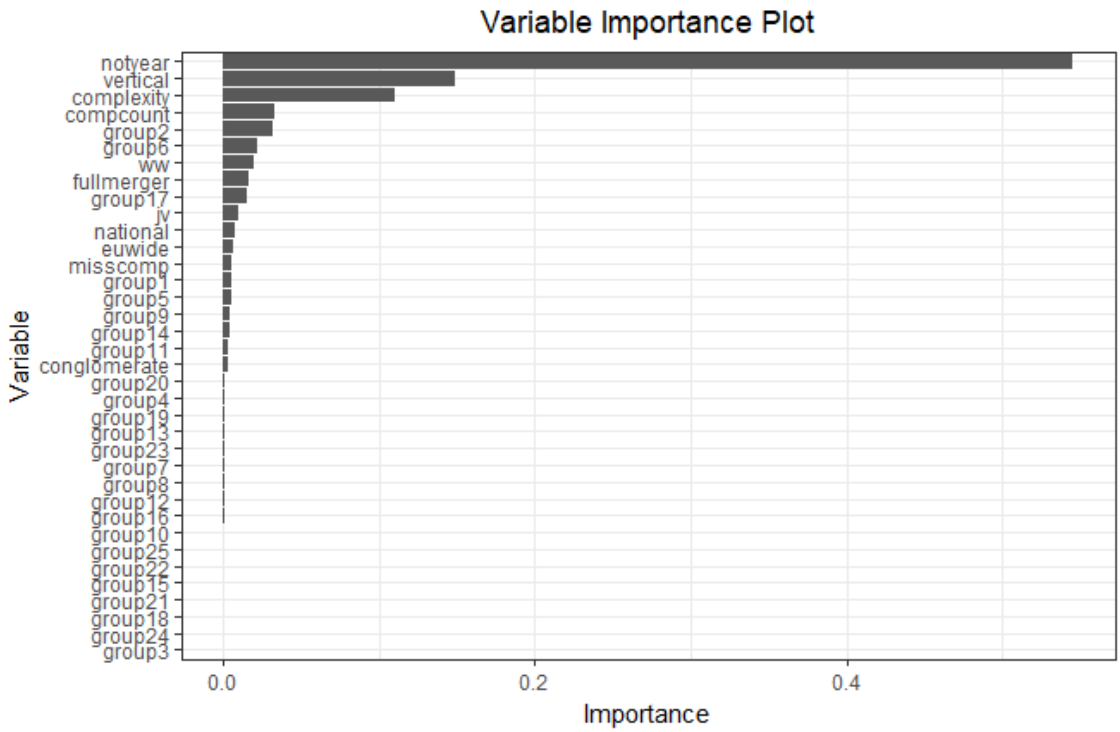

3.7.4.3 Treatment - Barriers to Entry

Figure 3.16: Variable Importance Plot for Correlation between Barriers to Entry and Concerns

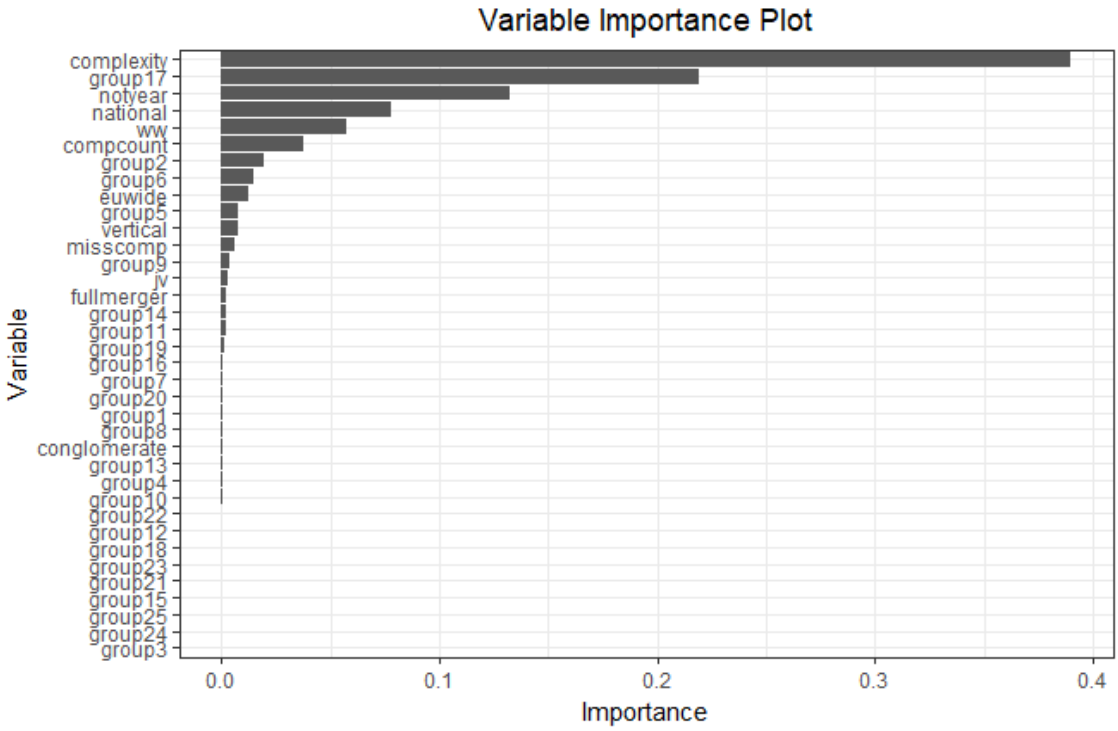

3.7.4.4 Treatment - Risk of Foreclosure

Figure 3.17: Variable Importance Plot for Correlation between Risk of Foreclosure and Concerns

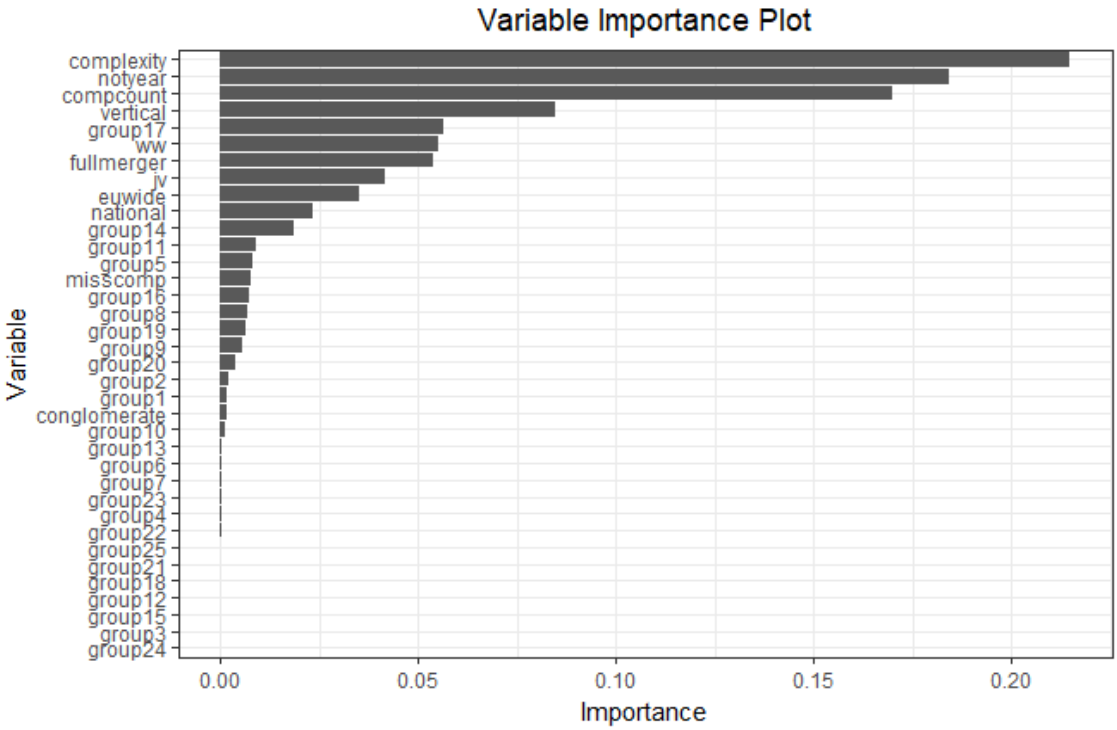

3.7.5 Determinants of Concern - Causal Forest Predictions over Industries

Figure 3.18: Effect of High Concentration on Concerns over Industries

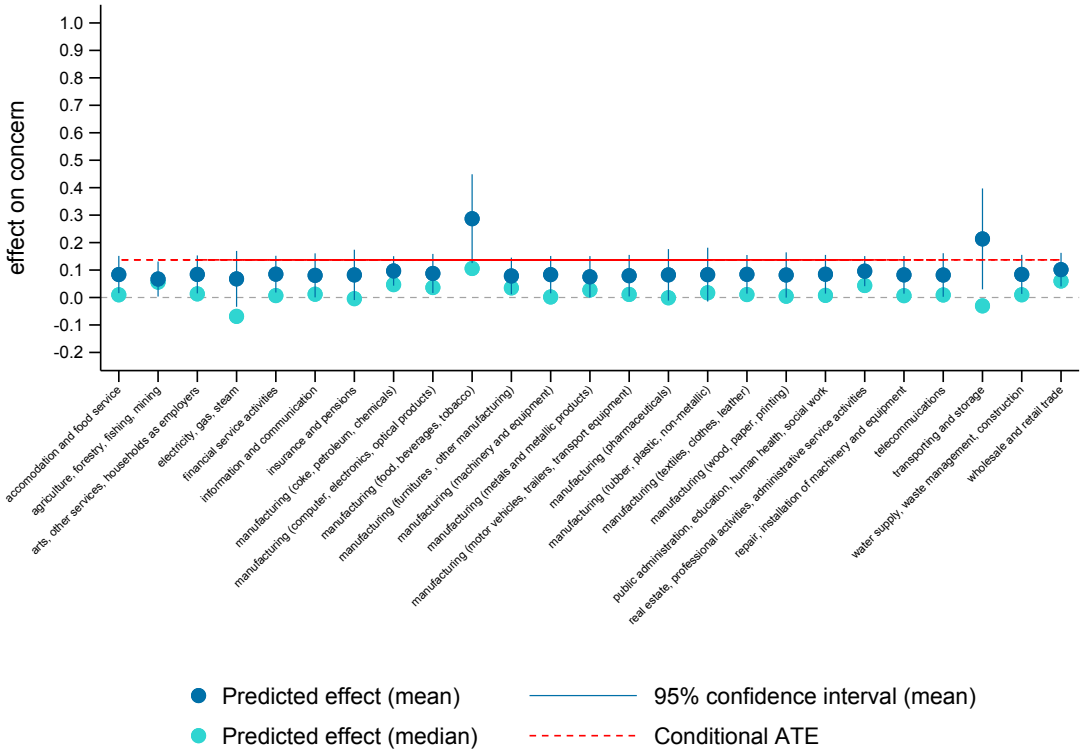

Predicted effect of indicator variable for post-merger HHI above 2000 and change in HHI larger than 150 on concerns over industries, setting all other included explanatory variables equal to the sample mean/median.

Figure 3.19: Effect of Joint Market Share on Concerns over Industries

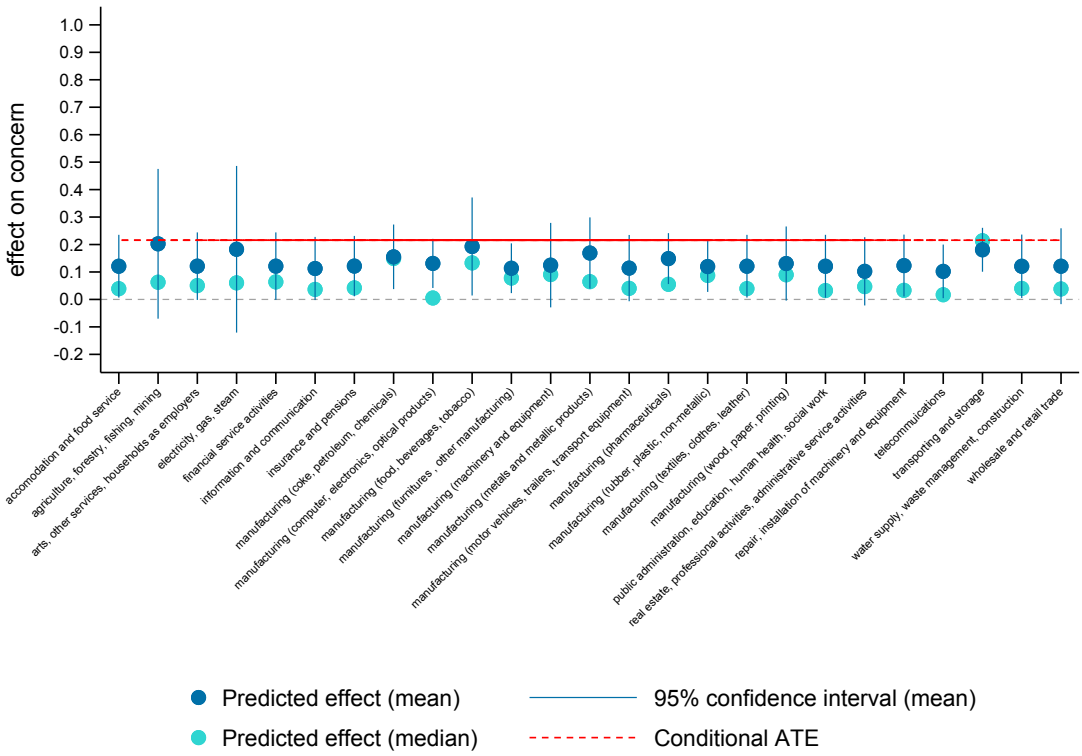

Predicted effect of indicator variable for joint market share above 50% on concerns over industries, setting all other included explanatory variables equal to the sample mean/median.

Figure 3.20: Effect of Barriers to Entry on Concerns over Industries

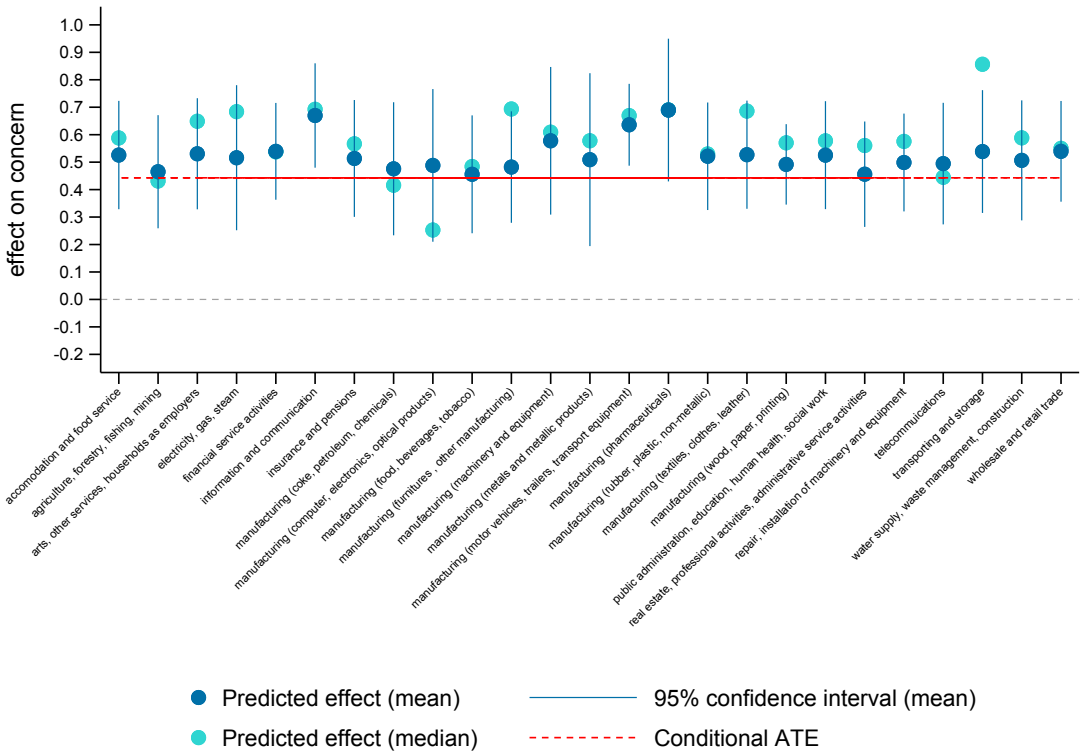

Predicted effect of barriers to entry on concerns over industries, setting all other included explanatory variables equal to the sample mean/median.

Figure 3.21: Effect of Risk of Foreclosure on Concerns over Industries

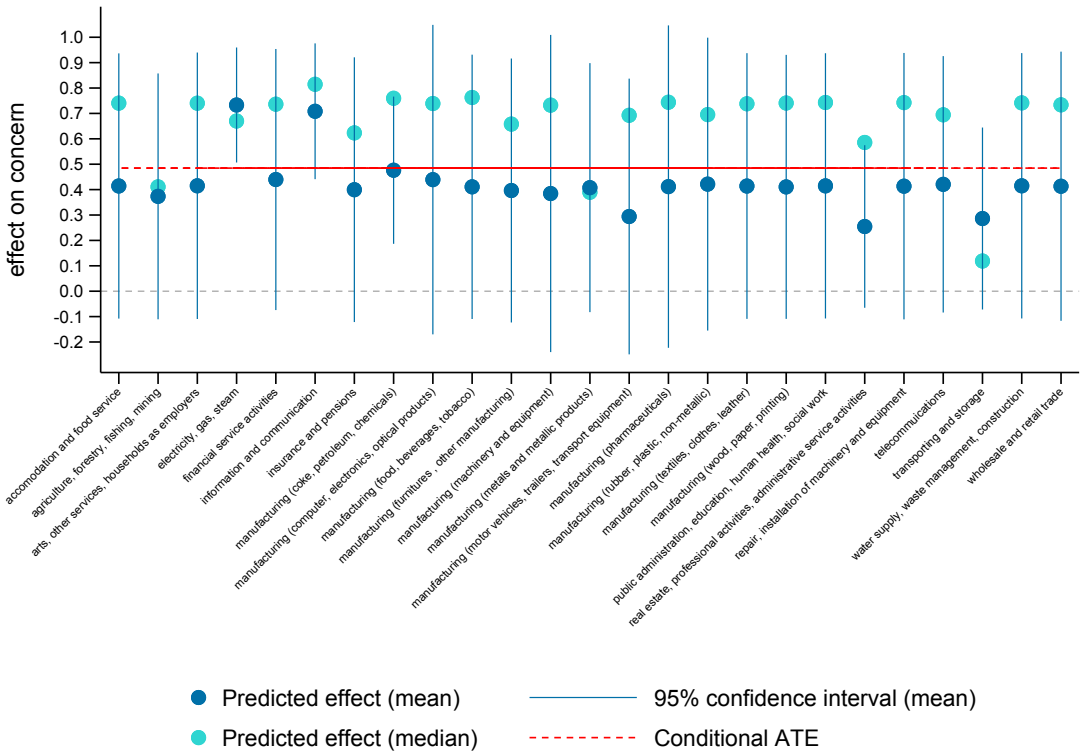

Predicted effect of risk of foreclosure on concerns over industries, setting all other included explanatory variables equal to the sample mean/median.

Chapter 4

EU Merger Policy Predictability Using Random Forests [1]

4.1 Introduction

European competition policy, i.e. the design and enforcement of competition rules, is designed to ensure fair and equal conditions for businesses. Competition policy covers the monitoring and, where necessary, blocking of anticompetitive agreements (in particular hardcore cartels), abuses by dominant firms, mergers and acquisitions as well as state aid. Among the different areas of the European Commission's (EC) competition policy, this paper focuses on merger policy.

The European Communities Merger Regulation (ECMR), the legal basis for common European merger policy, came into force in 1990. Over the course of the next 25 years, European merger control has seen significant changes, most prominently with the 2004 amendment to the ECMR. The goal of this reform was to adopt a "more economic approach" to European merger control, i.e. an approach closer to economic principles. With the reform, a new substantive test, the so-called "significant impediment of effective competition" (SIEC) test, as well as the concept of an efficiency defense were introduced, a chief economist was appointed, significant procedural changes were made, and horizontal merger guidelines were issued.

Merger policy can be approached and assessed from various angles. Duso, Gugler, and Szücs (2013) identify three dimensions along which merger policy effectiveness can be evaluated: the predictability, the correctness, and the deterrence effects of a decision. This paper considers only the first part of merger policy effectiveness:

[1] This chapter is the accepted manuscript published in the DIW Discussion Paper Series as: Affeldt, P. (2019). EU Merger Policy Predictability Using Random Forests. *DIW Discussion Paper No. 1800.* I thank Ivan Mitkov, Fabian Braesemann, David Heine, Juri Simons, and Isabel Stockton for their help with data collection.

its predictability. While ultimately the correctness of a decision is one of the main aspects of effective merger control, the predictability of decisions based on *ex ante* observable merger characteristics is of interest in its own respect.

Prior to the notification of a merger, legal certainty and the predictability of the merger control procedure are important for judges, competition lawyers, and, in particular, for firms' choices of which kind of mergers to propose. In particular, transparent antitrust policy fosters accountability of the antitrust authority and reduces personal bias in decisions taken, as decisions and the reasoning behind them must be made public. In turn, this will lead to a more fair and consistent process, and, consequently, decisions. The fact that the entire assessment process is transparent, with decisions and reasoning made public after a decision has been taken, also increases confidence in the assessment process, thus enhancing the authority's credibility. Lastly, a transparent and predictable process allows firms to better understand the authority's merger review process and, ultimately, predict the outcome of a merger review to a certain extent. Therefore, it should encourage self-compliance: firms should be encouraged to propose pro-competitive mergers and discouraged from proposing anti-competitive mergers (McAfee, 2010). While McAfee (2010) also discusses the costs of transparency in antitrust policy, the desirability of merger control policy with clear, transparent, and traceable rules and proceedings that decrease uncertainty for firms as well as the risk of political influence, has long been stressed (see for example Smith (1958) and Elman (1965)).[2]

One specific goal of the 2004 merger reform was precisely to increase legal certainty and transparency of the merger review process as evidenced by the publication of merger guidelines and the institutional changes made (See for example Gerber (2014)). However, the effect of the reform on the predictability of Directorate General Competition (DG Comp)'s decisions is ambiguous, as the use of a "more economic approach" in the merger review implies a shift from simple general rules, such as concentration thresholds, toward a more in depth case-by-case economic analysis. The question is hence whether the merger reform increased the *ex ante* predictability of decisions based on market and merger characteristics and also how the merger reform changed the decision criteria on which DG Comp bases its merger assessment.

To evaluate the predictability of DG Comp's merger policy, in this paper, I use a dataset containing almost all mergers notified to the European Commission between 1990 and 2014. This amounts to 25 years of data on European merger control. Un-

[2]In particular, McAfee (2010) argues that transparency encourages the use of simple rules that can lead to repeated and consistent mistakes in case assessment, increases the cost of investigations due to additional disclosure costs, might delay innovation in assessment techniques and increases the importance of precedents.

like most of the existing literature, rather than assessing mergers at the aggregate level, the data is collected at a more fine-grained level, defining an observation as a particular product and geographic market combination concerned by a merger.[3] Importantly, this allows for studying the factors that cause competitive concerns in specific sub-markets, as mergers typically concern several product and geographic markets that can be affected differently by the concentration. The final dataset used in the empirical analysis contains 22,812 product/geographic market level observations belonging to 2,417 DG Comp merger decisions.

The goal of this paper is, firstly, to study the predictability of DG Comp's merger policy and, secondly, to assess how it changed following the 2004 merger reform. Unlike the existing literature studying the determinants of DG Comp's merger intervention decisions and their predictability, I use the non-parametric random forests algorithm by Breiman (2001) to predict DG Comp's assessment of competitive concerns arising in affected markets due to the merger. This machine learning algorithm is designed to maximize predictive performance rather than estimating causal effects and allows for highly flexible, non-linear interactions between covariates. First, I train two random forests, one pre-reform and one post-reform, and compare their predictive performance to the predictive performance of a Linear Probability Model (LPM). Second, I study how the predictions of the pre-reform and post-reform random forests differ and how the merger assessment of DG Comp changed following the 2004 merger reform.

I find that the predictive performance of the random forests is much better than the performance of the LPM models. While all models are able to predict the majority outcome of *no competitive concerns* very well (between 80% and 90% correct predictions), the LPM models do very poorly in predicting the minority outcome of *competitive concerns*, with only between 16% and 44% correct predictions. In particular, the pre-reform LPM model only correctly predicts about 30% of the competitive concern markets in the pre-reform prediction set, the post-reform LPM model only predicts about 31% of the markets with concerns in the post-reform prediction set correctly. Furthermore, based on these predictions as well as the R^2, the LPM models would wrongly suggest that the predictability of DG Comp's assessment decreased after the 2004 reform. The random forests however, are able to correctly classify the minority class cases in about 60% of the cases, both pre- and post-reform. Thus, it is not true that the predictability of DG Comp's merger policy decreases post-reform.

Secondly, the random forest trained on pre-reform data only predicts well for the pre-reform period but does significantly worse than the random forest trained on

[3] One exception is Mini (2018), who also uses affected markets as the level of observation.

the post-reform data in predicting the minority outcome of *competitive concerns* for the post-reform period. This suggests, in line with the results of Affeldt, Duso, and Szücs (2019), that DG Comp's assessment criteria changed post-reform. Based on the two random forests and correcting for the case mix effect, the policy effect shows a decrease in the concern rate of about 1.5 percentage points post-reform. However, this decomposition only considers the average concern rate rather than investigating for which observations the predictions of the pre-reform and the post-reform random forests differ. Studying post-reform cases for which the predictions of the two random forests differ, I find that pre-reform DG Comp seems to have relied more on structural indicators in its merger assessment, while post-reform DG Comp seems to base its assessment of competitive concerns more on a case-by-case analysis and less on simple structural indicators such as market shares or concentration measures. Nevertheless, the highly flexible random forest algorithm still allows for high prediction precision. These results are in line with the findings from Affeldt, Duso, and Szücs (2019), who find the same time-dynamics of DG Comp's decision procedures using causal forests without imposing a different pre- and post-reform model.

The paper is structured as follows. In Section 4.2, I discuss the institutional background of European merger control and the 2004 reform of the EU merger guidelines. In Section 4.3, I review studies that empirically investigate the determinants and predictability of merger intervention and present some recent papers that employ machine learning techniques for prediction tasks. I describe the dataset used in the empirical analysis in Section 4.4 and the random forest algorithm employed for prediction in Section 4.5. I present the prediction results and discuss the change in the decision rules after the merger reform in Section 4.6 and conclude in Section 4.7.

4.2 Institutional Background

Mergers and merger control are important for firms as well as society and, ultimately, consumers. While mergers can reduce competition and lead to increases in market power, and, consequently, price increases to the detriment of consumers, they can also lead to important efficiency gains due to economies of scale and scope that are partly passed on to consumers, increasing not only producer surplus but also consumer surplus.

In the European Union (EU), regulation of competition has been undertaken by the European Commission since 1975. Specifically, DG Comp is responsible for enforcing EU rules regarding antitrust, mergers, state aid, and liberalization with the goal of protecting consumer surplus.

The basis for DG Comp's merger policy is the European Communities Merger Regulation (ECMR), which was passed in 1989 and came into force in September 1990.[4] It specifies the scope of intervention and legal competence of the European Commission in merger cases with a "community dimension". The definition of this "community dimension" was broadened by the passing of regulation 1310/97[5] in 1997. In particular, mergers must be notified to the EC if the combined worldwide turnover of the merging parties is sufficiently high, if their combined intra-community turnover is sufficiently high and not too concentrated in one Member State only.[6] This implies that, from 1990 onwards, all major combinations had to be notified and have been scrutinized by DG Comp. Notice that these definitions also include companies that are located, produce, and sell outside of the European Union, as long as their sales to European markets are sufficiently high. Thus, a merger can be subject to the jurisdiction of more than one competition authority. For example, the merger of the two U.S. companies *General Electric* and *Honeywell* was ratified by American authorities, but prohibited by the European Commission in 2001.[7]

Once it is established that a combination is subject to EC jurisdiction, the merging parties must notify the concentration to the EC prior to its implementation. After the reception of the official notification, the EC publishes a note in the Official Journal of the European Communities and third parties can comment on the proposed merger.

After notification to the EC (and the receipt of all necessary information), so-called phase-1 proceedings are initiated. DG Comp then has 25 working days (which can be extended to a maximum of 35 working days) for an initial assessment of the merger. DG Comp can then clear the proposed merger (phase-1 clearance), clear it

[4]Council Regulation (EEC) No 4064/89 of 21 December 1989 on the control of concentrations between undertakings [Official Journal L 395 of 30 December 1989].

[5]Council Regulation (EC) No 1310/97 of 30 June 1997 [Official Journal L 180 of 9 July 1997].

[6]In particular, in article 1.2 of regulation 4064/89, a combination is defined to have community dimension by meeting the following conditions: a) the aggregate worldwide turnover of all the undertakings concerned is more than ECU 5 000 million, and b) the aggregate Community-wide turnover of each of at least two of the undertakings concerned is more than ECU 250 million, unless each of the undertakings concerned achieves more than two-thirds of its aggregate Community-wide turnover within one and the same Member State. Regulation 1310/97 assesses a community dimension even if a merger does not meet the original two conditions, provided it satisfies the following four conditions: (a) the combined aggregate worldwide turnover of all the undertakings concerned is more than EUR 2 500 million; (b) in each of at least three Member States, the combined aggregate turnover of all the undertakings concerned is more than EUR 100 million; (c) in each of at least three Member States included for the purpose of point (b), the aggregate turnover of each of at least two of the undertakings concerned is more than EUR 25 million; and (d) the aggregate Community-wide turnover of each of at least two of the undertakings concerned is more than EUR 100 million, unless each of the undertakings concerned achieves more than two-thirds of its aggregate Community-wide turnover within one and the same Member State.

[7]See for example Patterson and Shapiro (2001) for a discussion of the case.

subject to remedies proposed by the merging parties (phase-1 remedy), or initiate a more in-depth investigation (phase-2 investigation) depending on whether the proposed transaction raises competitive concerns and whether these can be addressed by initial remedies or not. Furthermore, the merging parties can also withdraw the proposed merger during phase-1 (phase-1 withdrawal).

If DG Comp initiates an in-depth phase-2 investigation, it may take up to 90 working days. Based on the conclusions from this in-depth investigation of the effects of the merger, DG Comp can again unconditionally clear the merger (phase-2 clearance), clear the merger subject to commitments by the merging parties (phase-2 remedy), or prohibit the merger (phase-2 prohibition). Again, the merging parties can also withdraw the proposed merger in phase-2 (phase-2 withdrawal). Bergman, Jakobsson, and Razo (2005) argue that withdrawing a merger in phase-2 of the investigation process is virtually equivalent to a prohibition as parties often withdraw a merger before an actual prohibition takes place. Hence, both a prohibition as well as a phase-2 withdrawal suggest that the EC and the notifying parties were unable to find suitable remedies to address the anti-competitive concerns of the proposed transaction.

The ECMR was amended in 2004. This amendment brought significant changes to European merger control with the aim of bringing merger control closer to economic principles: the concept of an efficiency defense was introduced, a chief economist was appointed, the timetable for remedies was improved, and horizontal merger guidelines were issued. The reception of the new merger regulation was generally favorable (Lyons, 2004).

One of the most significant changes in the horizontal merger guidelines was the move from the old "Dominance Test" (DT) for market power to a "significant impediment of effective competition test" (SIEC). In the old DT test, a merger was declared as incompatible with the common market if it *"creates or strengthens a dominant position as a result of which effective competition would be significantly impeded"*. This implies that pre-reform, the creation or strengthening of a dominant position was a necessary condition for the prohibition of a merger. Thus, mergers that reduced effective competition without creating a dominant position could not be challenged under the old legislation. It is argued that the dominance test was deficient in cases of collective dominance and tacit collusion and that the "substantial lessening of competition" test employed by the United States' Federal Trade Commission (FTC) would be preferable.

In the revised 2004 Merger Regulation the wording of article 8.3 (prohibition) hence reads:

> *"A concentration which would significantly impede effective competition, in the common market or in a substantial part of it, in particular as a result of the creation or strengthening of a dominant position, shall be declared incompatible with the common market".*[8]

This revision implies that the creation of a dominant position is no longer a necessary condition for intervention and, therefore, aligns the test used by DG Comp more closely with U.S. practice (Bergman, Coate, Jakobsson, and Ulrick, 2010; Szücs, 2012).

4.3 Previous Literature

Mergers are an important research topic in the field of industrial organization. There are large bodies of theoretical and empirical literatures on questions such as firms' incentives to merge and merger policy effectiveness. Duso, Gugler, and Szücs (2013) identify three dimensions along which merger policy effectiveness can be evaluated: the predictability, correctness, and deterrence effects of a decision. A large part of the literature studying the effectiveness of merger control looks at whether the competition authority made the correct decision in a particular case (*ex-post* evaluations of merger policy) (Duso, Neven, and Röller, 2007; Duso, Gugler, and Yurtoglu, 2011; Kwoka, 2013).[9] A correct decision in this context is a decision that achieves the goals set in the legal framework - in the European Union as well as in most other jurisdictions the goal of competition policy is the protection of consumer surplus. A merger that decreases consumer surplus is considered to be anti-competitive. Thus, in order to judge whether a particular decision was correct, one must determine whether a given merger harmed consumer surplus. For example, Duso, Neven, and Röller (2007) use the reaction of competitors' stock market prices to evaluate the degree of pro- or anti-competitiveness for a sample of mergers.[10] They then employ a probit model to estimate the frequency of type I (prohibition of a pro-competitive merger) and type II (clearance of an anti-competitive merger) errors in the decisions.

[8] "Dominance" has been defined as *"... a position of economic strength enjoyed by an undertaking which enables it to prevent effective competition being maintained on the relevant market by giving it the power to behave to an appreciable extent independent of its competitors, customers and ultimately of consumers"* by the European Court of Justice in *United Brands* (27/67, E.C.R. 207, para. 65).

[9] Duso (2012) provides a literature review of *ex-post* merger evaluation studies.

[10] If the merger is anti-competitive, it will likely lead to rising share prices of competitors, as in an oligopolistic setting an increase in concentration leads to an increase in mark-ups and profits also of competitors.

Instead, the present paper focuses only on the first part of merger policy effectiveness: its predictability. The goal of the paper is to understand how DG Comp decides on interventions in merger cases and whether it is possible to predict DG Comp's decision based on *ex ante* merger and market characteristics. However, these predictions do not allow for judging whether DG Comp's decision was correct in the sense that it protected consumer surplus.

Therefore, the present paper is specifically related to two strands of the literature. Firstly, by predicting the intervention decision by DG Comp based on *ex ante* merger characteristics, this paper relates to the literature on the determinants of merger policy intervention by competition authorities and the predictability of decisions, in particular to those papers that empirically review the EU merger policy reform. Secondly, in this paper, I employ machine learning techniques to predict DG Comp's decisions. Hence, on the methodological side, this paper relates to the literature on prediction using machine learning. In the following, I give a brief overview of these two strands of literature.

4.3.1 Policy Predictability

First, to the best of my knowledge, all of the related literature, except the two studies by Bradford, Jackson, and Zytnick (2018) and Mini (2018), investigate the determinants of merger intervention decisions *at the merger level* and for a *sample of merger cases* only. As discussed in detail in Affeldt, Duso, and Szücs (2019), the scope and depth of the present data allow for going beyond the existing literature by allowing for heterogeneity within merger cases by examining the individual product and geographic markets concerned.

Secondly, the existing literature focuses on studying and discussing the relevant determinants of merger intervention decisions, the difference in policy between the EU and the U.S., or between the EU pre- and post-merger reform. None of the existing studies focus on the predictive performance of the employed empirical models.

Thirdly, all of the existing literature uses parametric models to empirically study the determinants of merger intervention decisions. Instead, I employ flexible, non-parametric machine learning techniques designed to maximize predictive performance to predict DG Comp's decisions. In particular, as I argue below, these machine learning techniques allow for predicting outcomes well, even for the rare outcome. However, as the focus of the existing literature is not on predictive performance, most papers report the percentage of correct predictions but do not distinguish between correct predictions of the majority and minority class. However, at the extreme, if e.g. 90% of the observations are class 1 and the model classifies

100% of the observations as class 1, then 90% of the cases are correctly predicted but the model is essentially useless as it is unable to correctly predict any of the rare class 2 cases.

Bergman and co-authors study European merger control in a series of papers. In the first study, Bergman, Jakobsson, and Razo (2005) employ a logit model for a sample of 96 EU merger cases to estimate the likelihood of going to phase-2 or prohibition decisions as a function of market-relevant and political variables. They report an overall rate of correctly classified cases of going to phase-2 of 79% - 91% and of the decision to prohibit a merger of 84% - 86%. The authors also discuss that while this overall rate of correctly classified cases is relatively high, only between 20% and 27% of the prohibitions are correctly classified while all non-prohibitions are correctly classified. When grouping prohibitions and phase-2 approvals with commitments together, the overall level of predictability lies between 69% and 93% depending on the specification. For one of the models, the authors report that they are now able to classify 100% of the approvals and 75% of the prohibitions/commitments correctly. However, note that the regressions are based only on a sample of 96 merger cases, out of which 17 are prohibitions, 30 are other phase-2 cases, and 49 are other cases. In terms of unbalancedness of the data, only the specification not regrouping prohibitions and approvals with commitments is comparable to the rate of competitive concerns in the dataset that I use. In this specification, the percentage of correctly predicted prohibitions is below 30%, which is much lower than the percentage of correct predictions for the rare class I achieve with my model (see Section 4.6).

Bergman, Coate, Jakobsson, and Ulrick (2010) examine instead similarities between EU and U.S. merger decisions using a sample of horizontal phase-2 mergers between 1990-2004 for both the EU (109 cases) and the U.S. (166 cases). They estimate a probit model for each regime to evaluate enforcement policy, where the dependent variable is an indicator for intervention (one for prohibition, approval subject to substantial remedies, or withdrawal by the parties at least one month into the phase-2 investigation). The authors report 83% - 84% and 87% - 91% correct predictions for the EU model and the U.S. model, respectively, depending on the specification. While they state that in both regimes "challenges are predicted more accurately than are closed investigations" (Bergman, Coate, Jakobsson, and Ulrick, 2010, p.321), they do not report the percentage of correct predictions for the different outcomes.

In the most recent study of the series, Bergman, Coate, Mai, and Ulrick (2016) update the dataset of Bergman, Coate, Jakobsson, and Ulrick (2010) by adding observations both to the EU as well as the U.S. dataset for the time period following

the 2004 EU merger policy reform. Their final dataset contains a sample of 151 EU phase-2 cases and 260 U.S. cases, covering the 1993-2013 period. Separate logit models on an intervention indicator are estimated for EU cases (distinguishing pre- and post-reform) and U.S. cases. While the authors report about 82% correct predictions for the EU models and above 90% correct predictions for the U.S. models, they do not discuss these results in the paper nor do they report the percentage of correct predictions for the intervention and no-intervention cases separately.

Szücs (2012) investigates the convergence between U.S. and EU merger policy following the 2004 EU merger policy reform using a sample of 309 EU and 286 U.S. merger cases decided by DG Comp and the FTC, respectively, between 1991 and 2008. For each of the pre-reform EU, post-reform EU and U.S. merger samples, he estimates a logit model on the decision to intervene and then uses the estimated models to predict the probability of intervention for each merger case from the point of view of both competition authorities (similar to Bergman, Coate, Jakobsson, and Ulrick (2010)). While the U.S. model classifies 86% of the cases correctly, the percentage of correctly classified cases is above 90% for the EU model both pre- and post-reform. The author does not report the percentage of correct predictions for the intervention and no-intervention outcomes separately.

Duso, Gugler, and Szücs (2013) evaluate European merger policy effectiveness along three dimensions: the predictability, correctness, and deterrence effects of a decision. Regarding predictability of European merger policy, Duso, Gugler, and Szücs (2013) estimate two probit models (one pre-reform, one post-reform) for a sample of 368 EU merger cases where the intervention decision of DG Comp (remedies or prohibition) is a function of *ex ante* observable merger characteristics. Model fit is discussed in terms of pseudo R^2 and the percentage of correctly classified observations (71% pre-reform and 76% post-reform). Once again, the percentage of correctly classified cases is not reported for the intervention and no-intervention outcomes separately.

Mai (2016) studies the effect of the EU merger policy reform on the probability of a merger being challenged by DG Comp based on a sample of 341 phase-1 and phase-2 horizontal mergers between 1990 and 2012. The probability of a challenge in a probit model pooling pre- and post-reform cases is driven by the market shares of the merging parties, entry barriers, and some other factors. Mai (2016) also estimates separate pre- and post-reform models and applies the methodology used by Bergman, Coate, Jakobsson, and Ulrick (2010), Szücs (2012) and Bergman, Coate, Mai, and Ulrick (2016) by predicting the probability of a challenge for pre-reform mergers using the post-reform model and vice versa. The author reports an overall rate of correct predictions above 80% and up to 90% depending on the specification,

where the post-reform models generally perform better than the pre-reform models. As in the papers discussed previously, correct predictions for the different classes are not reported.

Bradford, Jackson, and Zytnick (2018) empirically investigate whether European merger control is used for protectionism and find no evidence that DG Comp intervened more frequently or extensively in transactions involving non-EU or U.S.-based firms. Differently from the previous literature and similar to the present dataset, they collect information on all merger cases decided by DG Comp between 1990 and 2014. However, their analysis is still conducted at the level of the merger rather than the concerned product and geographic market. Furthermore, they do not collect information on the structural parameters of market shares, concentration, likelihood of entry, and foreclosure from the case documents. For the linear probability models estimating the probability of challenge, the authors only report R^2 as a measure of model fit and do not discuss the predictive performance of the model at all.

The paper with the most closely related dataset to the one used here is the one by Mini (2018). Like this paper, and unlike all other existing studies, Mini (2018) also collects information on the universe of EU merger decisions from the publicly available case documents between 1990 and 2013, recording each market concerned by the transaction as a separate observation. Thus, for each merger, he records potentially many observations and collects similar merger and market level characteristics from the case documents, like those included in the dataset used here. He then estimates probit models at this concerned market level for horizontal overlap markets, interacting all explanatory variables with a post-reform indicator variable. In the first model, the main variables of interest are the merging parties' market shares and the change in market shares, in the second model, the main variables of interest are post-merger Herfindahl-Hirschman-Index (HHI)[11] as well as the change in HHI due to the merger. For these models, the author reports about 92% correctly predicted observations; however also here predictions are not distinguished by class.

Thus, while most of the existing literature reports the percentage of correctly predicted observations together with the R^2 as an indicator of model fit, Bergman, Jakobsson, and Razo (2005) is the only paper that discusses the lower predictive performance of the model for the rare class. Most papers report percentages of correctly predicted observations above 80% and up to 90%. However, once again, if the data contains for example 90% class 1 observations and only 10% class 2 observations and the model classifies 100% of the observations as class 1, then 90% of the cases are correctly predicted but the model is not able to correctly predict any of the rare class 2 cases.

[11]The HHI is defined as the sum of squared market shares of all firms active in the market.

4.3.2 Prediction using Machine Learning

Unlike the existing literature on the effects of the EU merger policy reform, in this paper I employ non-parametric machine learning techniques to predict the intervention decision by DG Comp and to evaluate how the decisions changed post EU merger reform. On the methodological side, this paper therefore relates to the economics literature employing machine learning techniques for prediction tasks. While this list is by no means exhaustive, I mention a few papers here. The topics these papers study are very different from my application, but all studies try to predict an outcome based on observables using machine learning techniques. Not all of them use the random forest algorithm, but they all focus on prediction rather than causal inference employing machine learning techniques that allow for more complex interactions between covariates than do parametric models.

According to Kleinberg, Ludwig, Mullainathan, and Obermeyer (2015) an increasing number of empirical studies consider prediction policy problems. In particular, the authors mention prediction problems in education (e.g. value added of teachers), labor market policy (e.g. unemployment spell length), social policy (e.g. predicting highest risk youth for targeting interventions), and finance (e.g. identifying creditworthiness of borrowers). Kleinberg, Ludwig, Mullainathan, and Obermeyer (2015) further include an illustrative healthcare application in which they use a sample of Medicare beneficiaries to predict the pay-off of hip or knee replacement surgery. The question the authors ask is whether one can predict which surgeries are futile based on patient characteristics. In order to do so they predict mortality in the 1-12 months after hip or knee replacement using Lasso based on patient demographics, co-morbidities, symptoms, injuries, acute conditions, and health care utilization. The predictive performance of the model is however not discussed.

Chalfin, Danieli, Hillis, Jelveh, Luca, Ludwig, and Mullainathan (2016) study the selection of the most productive labor input, which is also a prediction policy problem. In particular, they illustrate the use of machine learning techniques in this respect with two applications: they predict worker productivity to improve police hiring practices (lowering police use of force or misbehavior) and teacher tenure decisions (improving teacher value added) using stochastic gradient boosting in the first and regression with lasso penalty in the second application. Explanatory variables include socio-demographic attributes of workers (i.e. police officers or teachers), students, schools and surveys capturing e.g. prior behavior. However, the authors do not discuss the predictive performance of the models in detail.

Björkegren and Grissen (2018) use machine learning techniques to predict loan repayment for post-paid mobile subscriptions in a developing country context. In

particular, they use random forests and logistic regression with a model selection procedure to predict the probability of loan default using behavioural patterns derived from raw mobile operator transaction records. Model performance is mainly evaluated based on the area under the receiver operating characteristic curve (AUC). For the random forests, the AUC lies between 0.62 and 0.71 depending on the model, which is much lower than the AUCs I achieve with the random forests (see Section 4.5).

Kleinberg, Lakkaraju, Leskovec, Ludwig, and Mullainathan (2018) and Ribers and Ullrich (2018) go one step further by not only predicting outcomes but also studying whether the machine learning algorithm makes better decisions than humans. Kleinberg, Lakkaraju, Leskovec, Ludwig, and Mullainathan (2018) analyze the problem of predicting risk of defendants' committing a crime in the context of judges' bail decisions using gradient boosted decision trees and judging whether machine prediction can improve judges' bail decisions. The authors highlight the importance of unobservable characteristics, selected labels, and omitted pay-offs to properly compare machine prediction and human decisions. In particular, the selective labels problem arises because crime outcomes can be observed only for those defendants who were released by the judges. Hence, crime rates of those defendants who were kept in bail have to be predicted. However, this is problematic as judges might have based their decisions on unobservables, so the crime rates of released defendants might not be a good proxy for crime rates of detained defendants with similar observable characteristics. To deal with this problem the authors use different econometric strategies including quasi-random assignment of cases to judges.

Ribers and Ullrich (2018) train a random forest on a high-dimensional administrative dataset from Denmark to predict bacterial causes of urinary tract infections. Unlike many machine learning papers tackling prediction centred policy problems, the authors have the advantage of observing labels, i.e. they have patient test outcomes independent of physician prescription choices. This implies that they can actually evaluate whether physicians or the machine made the right prediction. They show that machine predicted bacterial risk is highly correlated with the actual presence of bacterial urinary tract infection. They model the prescription decision as a trade-off between the social cost of prescribing (as it increases resistance) and the curative effect on the patient in case she truly has a bacterial infection. They find that antibiotic prescribing can be lowered by up to 10% with no reduction in the number of treated patients suffering from a bacterial infection based on a machine learning assisted prescription rule that allows to redistribute prescriptions from low risk to high risk patients.

I cannot make this additional step of evaluating whether DG Comp or the machine learning algorithm took the right merger intervention decision. In order to judge the correctness of a decision, I would need a measure of whether a proposed merger is anti- or pro-competitive, i.e. whether it decreases or increases consumer surplus. For example Duso, Neven, and Röller (2007) and Duso, Gugler, and Szücs (2013) use the reaction of competitors' stock market prices for a sample of mergers to determine whether a proposed merger is pro- or anti-competitive in order to evaluate whether DG Comp took the right decision.

4.4 Data

The initial merger database contains almost the entire population of DG Comp's merger decisions, both in the dimension of time and with regard to the scope of the decisions encompassed. The data were obtained from the publicly accessible cases published by DG Comp on the EC's webpage.[12] We start data collection with the very first year of common European merger control, 1990, and include all years up to 2014. This amounts to data on the first 25 years of European merger control. Rather than taking a particular merger case as the level of observation, we collected data at a more fine-grained level and defined an observation as a particular product and geographic market combination affected by a merger. For further details on the merger database as well as the data collection procedure, see the data documentation Affeldt, Duso, and Szücs (2018).

For the analysis in this study, I do not use all observations contained in the merger database. First, I drop cases that were referred back to member states as well as phase-1 withdrawals.[13] This leads to a dataset containing 5,109 DG Comp merger decisions, where each decision occupies a number of rows equal to the number of product/geographic markets affected in the specific transaction. This dataset contains a total of 30,995 market level observations and is used in the analysis in Affeldt, Duso, and Szücs (2019). Appendix 4.8.1 contains summary statistics for this comprehensive dataset.

Secondly, I further reduce the dataset, as I only keep observations of product/geographic markets for which information on the merging parties' joint market share is available and the calculation of the post-merger HHI is possible. The reason is that both the merging parties' joint market share as well as HHI are important criteria taken into account by DG Comp when assessing a merger proposal and, thus,

[12]The types of notified mergers, decisions taken and reports for each of the EC's decisions can be downloaded from: `http://ec.europa.eu/competition/mergers/cases/` and `http://ec.europa.eu/competition/mergers/legislation/simplified_procedure.html`.

[13]There is only two phase-1 withdrawals contained in the dataset.

I want to include them in the set of explanatory variables. This leads to a final dataset used in the empirical analysis containing 22,812 product/geographic market level observations belonging to 2,417 DG Comp merger decisions. All summary statistics and analyses presented in the following are based on this dataset.

As market share information is most frequently missing for phase-1 clearance cases, keeping only observations for which market share information is available leads to a selection issue. If mergers for which market share information is missing are systematically different from mergers for which market share information is available, then the estimated correlation between market share and concentration and any intervention decision also captures this unobserved difference and results would be biased. However, the selection issue is studied in Affeldt, Duso, and Szücs (2019), where separate OLS regressions are estimated based on the entire sample, based on the sample without available market share information, based on the sample with available market share information, and based on the sample with available market share information including the merging parties' joint market share as well as post-merger HHI as additional explanatory variables. The estimation results show that DG Comp's decision determinants are rather similar across all sub samples.[14] Furthermore, in this paper, I am not seeking to estimate a causal relationship between for example market shares and an intervention decision. I am just interested in a correct prediction of the intervention decision, which does not require unbiased or causal coefficients on the explanatory variables (see Section 4.5).

The dataset contains, first, information on the name of the acquirer and the target firm as well as the countries of the merging parties, the dates of the notification to the EC and the final decision[15] as well as the type of decision eventually taken by DG Comp (clearance, remedy, and prohibition) or whether the proposing parties withdrew the notification. The data also allows for distinguishing between a phase-1 and a phase-2 decision.

Table 4.1 reports the number of phase-1 clearances, phase-1 remedies, phase-2 clearances, phase-2 remedies, as well as prohibitions distinguishing the pre-reform and post-reform periods. Compared to the full merger database, keeping only observations that contain market share information leads to a slight over-representation of phase-2 cases.[16] Note also that, while the full merger database contains 19 prohi-

[14]Mini (2018) claims in his paper that sample selection is not an issue because he uses the universe of horizontal merger cases in his estimation. However, I do not agree with this statement, as he also uses the merging parties' market share as explanatory variable in his estimation and therefore only records cases "provided that the EC disclosed data on merging parties' shares" (Mini, 2018, p.5).

[15]Note that the notification of a merger and the decision do not necessarily take place in the same year.

[16]See Affeldt, Duso, and Szücs (2018) for comparison: In the full merger database about 90% of

bitions and 5 phase-2 withdrawals, for one of the prohibitions and all of the phase-2 withdrawals market share information is not available. Thus, my estimation dataset does not contain any phase-2 withdrawals.

Table 4.1: Type of Decisions, 1990-2014

	Pre-reform		Post-reform	
Type of decision	frequency	percent	frequency	percent
Phase-1 clearance	1,108	84.90	916	82.37
Phase-1 remedy	99	7.59	128	11.51
Phase-2 clearance	21	1.61	26	2.34
Phase-2 remedy	63	4.83	38	3.42
Prohibition	14	1.07	4	0.36
Total	1,305	100.00	1,112	100.00

The dataset contains a number of *ex ante* characteristics of the merger, some of them at the merger level, some of them at the level of the product/geographic markets affected. While Table 4.2 contains summary statistics of the variables contained in the dataset that vary at the market level, Table 4.3 contains summary statistics for the merger level variables. As I predict DG Comp's assessment separately for the pre- and post-reform periods in Section 4.6, I also report summary statistics distinguishing pre- and post-reform cases.

The first variable in Table 4.3, *Intervention*, is a dummy variable that indicates whether DG Comp intervened in a particular merger case. I define this variable to be equal to one if DG Comp prohibited the merger, cleared the merger subject to remedies in phase-1, or cleared the merger subject to remedies in phase-2.[17] The corresponding variable at the product/geographic market level is the first variable in Table 4.2, *Concern*, which is an indicator variable equal to one if the merger raised competitive concerns in a specific product/geographic market according to DG Comp. This is the case in about 14% of the markets pre-reform and 13% of the markets post-reform. As this variable indicates in which particular markets the merger is likely to be problematic, this is the dependent variable of the empirical analysis presented in Section 4.6. Thus, instead of estimating the overall probability of an intervention, I estimate the probability that competitive concerns are found in a market affected by the merger. The higher the fraction of concerned markets

the cases are phase-1 clearances.

[17]In principal, I would also treat a phase-2 withdrawal as an intervention by DG Comp as in Affeldt, Duso, and Szücs (2019). Given that the phase-2 withdrawals fall out of my estimation dataset due to lack of market share information, the treatment of withdrawals by the merging parties in phase-2 is not an issue here.

in which competitive concerns are found, the higher the likelihood that DG Comp will intervene in a merger case.[18]

Table 4.2: Summary Statistics Variables at Market Level

	Pre-reform			Post-reform		
	mean	sd	obs	mean	sd	obs
Concern	0.14	0.34	8,531	0.13	0.33	14,281
Vertical merger	0.18	0.39	8,531	0.35	0.48	14,281
Conglomerate merger	0.05	0.21	8,531	0.00	0.07	14,281
National market	0.66	0.47	8,531	0.64	0.48	14,281
EU wide market	0.24	0.43	8,531	0.22	0.41	14,281
Worldwide market	0.07	0.25	8,531	0.13	0.33	14,281
Left open market	0.03	0.17	8,531	0.01	0.10	14,281
Entry barriers	0.12	0.33	8,531	0.16	0.37	14,281
Risk of foreclosure	0.05	0.21	8,531	0.02	0.14	14,281
Number of competitors	1.85	2.23	8,531	2.04	2.58	14,281
No competitor information	0.45	0.50	8,531	0.46	0.50	14,281
Joint market share	30.49	22.31	8,531	33.63	24.27	14,281
Post-merger HHI (low)	1,890.51	2,112.91	8,531	2,301.39	2,495.94	14,281
Post-merger HHI (high)	5,658.73	2,252.89	8,531	5,627.19	2,249.99	14,281

The other variables contained in both tables are used as covariates in the empirical analysis and describe the merger as well as how the merger affects the concerned markets according to DG Comp's *ex ante* assessment. However, note that of course all of these variables are based on what the official decision documents state, so to some extent they might reflect the assessment or subjective view (or mistakes) of DG Comp. This issue is present in most papers in this literature. One exception is Duso, Gugler, and Szücs (2013), who include only truly *ex ante* observable merger characteristics (such as for example the nationality of the merging parties or whether the concentration is a full merger) as explanatory variables in the probit estimation for intervention. In the working paper version (Duso, Gugler, and Szücs, 2012), a second model, the so-called "investigation model", is estimated; it contains results from the merger investigation as additional explanatory variables. The predictive power of this second model, measured by pseudo R^2 as well as the percentage of correctly classified observations, increases significantly compared to the first.[19] Therefore, the predictability of DG Comp's intervention decisions would

[18]See also Affeldt, Duso, and Szücs (2019). In the regressions explaining an intervention decisions at the merger level, the fraction of affected markets in which the merger leads to competitive concerns according to DG Comp positively affects the probability of intervention.

[19]In particular, the pseudo R^2 increases from 0.19 (pre-reform) and 0.25 (post-reform) to 0.68 and 0.59 pre- and post-reform, respectively. Also the percentage of correctly classified observations increased from 71% (pre-reform) and 76% (post-reform) to 90% for both the pre- and post-reform models.

likely decrease if I based my estimations only on unambiguously objective merger and market characteristics contained in the decision documents (such as, for example, whether the merger is a full merger or not). However, there is a trade-off between basing the estimation only on *ex ante* observable merger characteristics and omitting important determinants of whether a merger raises competitive concerns, such as market shares or entry barriers. Furthermore, to the extent that the legal system, in which DG Comp is operating, provides a consistency check (e.g. on how market shares should be calculated or in which types of markets entry barriers are high), I believe that DG Comp's merger assessment should be consistent.

The nature of the merger is described by a number of indicator variables. The dataset contains an indicator variable distinguishing full mergers from acquisition of shares as well as an indicator variable for joint ventures; see Table 4.3. At the market level, the dataset contains information on whether a product/geographic market is vertically affected by the merger. Vertically affected markets are markets where one or more of the merging parties operate in a market that is upstream or downstream of a market in which another merging party is active and any of their individual or combined market shares at either level is 25% or more.[20] The dataset further includes an indicator variable that is one if the merger is conglomerate in nature in the particular concerned market[21], see Table 4.2.

Table 4.3: Summary Statistics Variables at Merger Level

	Pre-reform			Post-reform		
	mean	sd	obs	mean	sd	obs
Intervention	0.13	0.34	1,305	0.15	0.36	1,112
Full merger	0.60	0.49	1,305	0.69	0.46	1,112
Joint Venture	0.35	0.48	1,305	0.19	0.39	1,112
Number of concerned markets	7.47	11.53	1,305	14.50	23.42	1,112
EU acquirer	0.69	0.46	1,305	0.64	0.48	1,112
EU target	0.74	0.44	1,305	0.70	0.46	1,112
Indicator for July/August	0.18	0.38	1,305	0.17	0.37	1,112
Indicator for December	0.06	0.23	1,305	0.05	0.21	1,112

[20] Commission Regulation (EC) No 802/2004 of 7 April 2004 implementing Council Regulation (EC) No 139/2004 on the control of concentration between undertakings [Official Journal L 133 of 30 April 2004]. The market share threshold has been raised to 30% at the end of 2013. See Commission Implementing Regulation (EU) No 1269/2013 of December 2013 [Official Journal L 336 of 14 December 2013].

[21] Conglomerate mergers are *"mergers between companies that are active in closely related markets (e.g. mergers involving suppliers of complementary products or products that belong to the same product range)."* Guidelines on the assessment of non-horizontal mergers under the Council Regulation on the control of concentrations between undertakings (2008/C 265/07), paragraph 5 [Official Journal of 18 October 2008].

Furthermore, the dataset contains information on the geographic market definition adopted in each market by DG Comp, distinguishing geographic markets that are defined as being national, EU wide, or worldwide. Lastly, the geographic market definition can also be left open.

Further indicator variables record whether DG Comp considered barriers to entry to exist and whether DG Comp raised concerns that the merger would foreclose other firms in a particular market. The database also contains a count of the number of competitors in the concerned market and an indicator variable equal to one if no information on competitors is available. Merging parties face about two competitors on average; however, information on competitors is missing in about 50% of the markets - these are mainly mergers that were cleared in phase-1.

As said before, I only use observations for which information on the market shares of the merging parties could be collected from DG Comp's competitive assessment in the decision document. Thus, data availability is constrained by the extent of DG Comp's analysis. Market share information is collected at the level of the relevant product/geographic market combination. Since the publicly available case documents generally report only market share ranges, the dataset contains the midpoint of the reported market share interval.[22] Therefore, I cannot avoid that market shares contain measurement error leading to endogeneity bias. This is an issue that this study shares with the existing literature. To my knowledge, Mini (2018) is the only one who constructs expected market shares and expected concentration measures rather than using the midpoints of the market share ranges reported by DG Comp. He highlights the issue of measurement error in market shares and HHI and explicitly accounts for it in estimation. In the dataset we only recorded the midpoints of the market shares ranges. Thus, I cannot follow Mini (2018)'s approach. To my understanding it is also unclear in which direction the bias goes.[23] However, as we always recorded the lower bound of the market share range whenever the market

[22]Since DG Comp generally reports only a range of market shares in the publicly available documents, the market shares are defined to be equal to the central value of the interval. If, for example, the market share range indicated is [0-10] percent, a market share of 5 percent is recorded. If however the interval given in the decision is only 5 percentage points wide, the conservative lower market share bound is reported. If for example the market share interval is [15-20] percent, 15 percent market share is reported.

[23]According to Mini (2018), the midpoint approach yields the correct expected value if the random variable has a symmetric marginal distribution which is centered on the midpoint. However, if for a merger case the sum of the upper bounds of all market share intervals (including competitors) is larger than 100%, the midpoint approach is no longer correct: according to Mini (2018) the domains of the marginal distributions of market shares are no longer necessarily the whole range from lower to upper market share bound and the distributions are even no longer necessarily symmetric. And even in cases where the midpoint of the market share range is the expected market share, the change in HHI as well as the contribution of the merging parties to the post-merger HHI would be underestimated.

share intervals given in the decision were only 5 percentage points wide, I expect that the market shares are underestimated. This measurement error leads to an attenuation bias in the estimated relation between market shares and intervention decision even if the size of the measurement error is not correlated with market share itself. Therefore, I expect that the relation between market shares and intervention decision is underestimated. As stated previously in relation to the selection issue though, I am interested in prediction and not in estimating causal effects of market shares or HHI on intervention decisions. Therefore, I do not require unbiased coefficients (see Section 4.5).

The market share information allows the calculation of the merging parties' combined market shares and the construction of a post-merger HHI. Table 4.2 also includes summary statistics for the market share related variables. The average merging parties' joint market share is slightly above 30%, with average post-merger HHI between 1,891 and 5,659 depending on the period and the calculation method. In the database, we include two different HHI measures. The variable *Post-merger HHI (low)* is a lower bound of the post-merger HHI: it is calculated as the square of the merging parties joint markets share plus the sum of squared market shares of competitors whenever information on competitors' market shares is available. This assumes that competitors are very small, whenever market share information of competitors is not available but market shares do not add up to 100%. The variable *Post-merger HHI (high)*, on the other hand, is an upper bound for the post-merger HHI: it adds the square of all missing market shares (100% minus all available market share information) to *Post-merger HHI (low)*. This hence treats all missing market share information as one missing competitor. In the empirical analysis, I use the *Post-merger HHI (high)* variable in order to be conservative as this measure should be an upper bound for market concentration.[24] If anything, I will overestimate market concentration and how problematic a given merger will be in a particular market. If a merger is unlikely to raise competitive concerns in a market for which concentration is measured by the upper bound of HHI, it will also not raise competitive concerns if the lower bound of HHI is used instead.

Lastly, the data include information on the main industry in which a merger took place. The industry is identified by NACE codes, which is the industry classification system used by the European Union to classify different economic activities. For the empirical analysis, I grouped the industries into 25 groups as shown in Table 4.4, where some NACE codes are grouped together while primarily the manufacturing

[24]To the extent that I underestimate the merging parties' market shares, any HHI measure will also be underestimated. However, this is especially the case for the *Post-merger HHI (low)* measure as it assumes that any remaining competitors, for which no information is available, are very small. This is therefore an additional reason to use the *Post-merger HHI (high)* variable.

Table 4.4: Industry Groups, 1990-2014

Industry group	Pre-reform obs	Pre-reform cases	Post-reform obs	Post-reform cases
Accomodation and food service	56	11	66	10
Agriculture, forestry, fishing, mining	256	41	550	40
Arts, other services, households as employers	53	7	188	12
Electricity, gas, steam	181	47	617	56
Financial service activities	255	54	341	38
Information and communication	312	56	570	55
Insurance and pensions	299	61	303	42
Manufacturing (coke, petroleum, chemicals)	1,234	128	1,998	118
Manufacturing (computer, electronics, optical products)	394	56	881	82
Manufacturing (food, beverages, tobacco)	737	75	824	75
Manufacturing (furnitures , other manufacturing)	127	6	463	19
Manufacturing (machinery and equipment)	248	43	393	41
Manufacturing (metals and metallic products)	284	62	484	52
Manufacturing (motor vehicles, trailers, transport equipment)	582	96	496	65
Manufacturing (pharmaceuticals)	681	28	1,090	42
Manufacturing (rubber, plastic, non-metallic)	219	45	517	31
Manufacturing (textiles, clothes, leather)	26	4	87	5
Manufacturing (wood, paper, printing)	303	50	554	39
Public administration, education, human health, social work	28	8	62	6
Real estate, professional activities, administrative service activities	234	49	562	30
Repair, installation of machinery and equipment	649	91	101	15
Telecommuications	244	64	414	49
Transporting and storage	407	76	1,702	73
Water supply, waste management, construction	134	28	166	24
Wholesale and retail trade	588	119	852	93
Total	8,531	1,305	14,281	1,112

industry has been further divided into smaller subgroups. For 280 pre-reform and 205 post-reform observations, the industry classification was missing. To avoid losing these observations in the analysis, I returned to the decision documents and manually assigned these mergers to the 25 industry groups.

I additionally define a few further variables. The variables *EU acquirer* and *EU target* are indicator variables that are equal to one if at least one of the acquiring or targeted firms is located within the EU, respectively. I include these variables in order to test for differential treatment of mergers by DG Comp depending on the nationality.[25] In order to test whether the holiday season matters for the likelihood of intervention by DG Comp due to resource constraints during these periods, I define two further indicator variables for July/August and December respectively.

Comparing the summary statistics pre- and post-reform, the post-reform mergers seem to be slightly more often full mergers and less often joint ventures than

[25]See, for example, Bradford, Jackson, and Zytnick (2018), who empirically investigate whether EU merger control is used for protectionism.

pre-reform. Additionally, they concern more markets post-reform and more product/geographic markets are affected vertically post-reform. Lastly, post-reform, DG Comp defined more geographic markets to be worldwide in nature and identified entry barriers in more markets than pre-reform.

4.5 Prediction using Random Forests

In this paper, I use the random forest algorithm by Breiman (2001) implemented in the randomForest package in R (Liaw and Wiener, 2002) to predict whether DG Comp found competitive concerns in the markets affected by a merger and to evaluate how DG Comp's assessment and its predictability changed post EU merger reform. I chose the random forest algorithm over alternatives, such as logistic regression or lasso, because it is able to uncover highly flexible, non-linear functions in a high-dimensional feature space without over-fitting.

Random forests are one example of supervised machine learning techniques that typically use a set of features X to predict an outcome Y. Thus, the goal is to construct $\hat{Y}(x)$, which is an estimator of $\mathbf{E}[Y|X = x]$, rather than estimating the causal effect of X on Y (Athey, 2018). Therefore, the aim is to reach a low error in the prediction $\hat{Y}$, which does not require the coefficients to be unbiased or causal. Our usual econometric tools, such as regression techniques, are made for causal inference, i.e. for obtaining unbiased estimates of the causal effect of covariates X on outcome Y. As prediction error is a function of not only bias but also variance, these tools do not yield the most accurate prediction $\hat{Y}$ (Chalfin, Danieli, Hillis, Jelveh, Luca, Ludwig, and Mullainathan, 2016). The tools from machine learning, on the contrary, are designed to do exactly this: they adaptively use the data to decide on how to trade off bias and variance to maximize prediction performance while allowing for a rich set of covariates X and functional forms (so higher order interaction terms or trees that allow for a high degree of interactivity between covariates). While the analyst has to provide the list of covariates X, the functional form is at least in part determined as a function of the data.

The random forest algorithm by Breiman (2001) uses regression or classification trees to predict an outcome Y. In a standard CART tree (Classification and Regression Tree), the goal is to predict individual outcomes Y_i using the mean outcome Y of observations that are "close" in X-space. To determine which observations are "close", the algorithm starts to recursively split the covariate space (binary splits) until it is partitioned into a set of so-called leaves L that contain only a few observations. The outcome Y_i for observation i is then predicted by identifying the leaf containing observation i based on its characteristics X_i and setting the prediction

to the mean outcome within that leaf:

$$\hat{Y}(x) = \frac{1}{|\{i : X_i \in L(x)\}|} \sum_{\{i:X_i \in L(x)\}} Y_i \tag{4.1}$$

The algorithm automatically decides on the splitting variables and split points. This is done based on an in sample goodness-of-fit criterion (so essentially how close the predicted outcomes are to the actual outcomes). For regression trees (continuous outcome variable Y), the goodness-of-fit criterion used is the mean squared error; for classification trees (categorical outcome variable Y), the goodness-of-fit criterion is a measure of classification error based on the empirical classification probabilities in the leaves.[26] The algorithm then splits on the covariate at the cut-off value that leads to the greatest improvement in the goodness-of-fit criterion. Once the best split at a given node in the tree is found, the splitting process is repeated in each of the resulting two regions. For CART trees, the splitting process is usually stopped when a specified minimum node size is reached - by default this is a node size of 5 for regression and 1 for classification trees. The tree is then pruned based on some cost-complexity trade-off measure in order to avoid over-fitting (See Hastie, Tibshirani, and Friedman (2008) chapter 9 for further details).

A random forest is essentially an ensemble of regression or classification trees, where the predictions are averaged across trees.[27] A random forest introduces two layers of randomness compared to a single classification tree. First, each individual tree in the forest is grown using a random sample with replacement from the training set. In each tree, one-third of the data is not used for training and can be used for testing (out-of-bag error). Secondly, differently from CART trees, splitting a node in a tree is done based on only a random subset of the covariates X rather than the full set of covariates and each tree is grown to the largest extent possible without pruning. The idea behind random forests is to reduce variance and produce more robust predictions compared to a single tree. Using a different bootstrap sample of the data to grow each tree in the forest as well as splitting each node based on only a subset of the covariates reduces the correlation between the trees in the forest and, therefore, the variance of the predictions (See Breiman (2001) and Hastie, Tibshirani, and Friedman (2008) chapter 15 for further details). Random forests are robust to overfitting and while one must choose the number of trees, the number of

[26] The randomForest R package uses the Gini index as measure for node impurity. Specifically, if $\hat{p}_{mk}$ is the proportion of class k observations in node m, the Gini index for node m is given by $\sum_{k=1}^{K} \hat{p}_{mk}(1 - \hat{p}_{mk})$.

[27] For classification problems, the random forest obtains a class vote from each tree and then classifies based on majority vote.

variables to be considered for splitting at each node, as well as the node size[28], it is argued that results do not usually seem to be very sensitive to the choice of these parameters (Liaw and Wiener, 2002). However, I use 5-fold cross-validation to tune these parameters (see below).

In order to predict the competitive concerns found by DG Comp and to evaluate how the assessment and its predictability changed post EU merger reform, I train two random forests: one based on pre-reform data and one based on post-reform data. For each of the two time periods, I randomly split the data into one training set (80% of the observations) used for building the random forest and one hold-out set (20% of the observations) used for out-of-sample prediction. After training the two random forests on the two training datasets, I first use the pre-reform and post-reform prediction sets to evaluate the predictive performance of the random forests and then, secondly, to study how DG Comp's assessment differs pre- and post-reform.

The random forests are trained to predict competitive concerns in a particular market affected by the merger. Thus, instead of estimating the overall probability of an intervention *at the merger level*, I estimate the likelihood that competitive concerns are found in a particular market affected by the merger. However, this directly relates to the intervention decision of DG Comp, as the probability of intervention increases with the fraction of affected markets in which competitive concerns are found.

However, training the model at the market level, rather than the merger level, has an important implication. It is unlikely that the observations of the different markets affected by a merger are independent. On the contrary, it is likely that there are unobservables at the affected market level that are correlated across the different markets concerned by a given merger and that influence whether DGComp identified competitive concerns. Furthermore, also the observable characteristics are correlated across affected markets: Firstly, because some observable characteristics only vary at the merger level (and, hence, take exactly the same value across markets). Secondly, because concerned markets are often very similar markets where, for example, if entry barriers are found to be high in one market, they are likely to be high in another as well. If the data is then simply randomly split into training and hold-out set, for a given merger some concerned markets might end up in the training set and some might end up in the hold-out set used for prediction. Given that the outcome of competitive concerns is highly correlated across affected markets and that the observable characteristics of these observations are very similar or identical in some cases, the predictive performance of this random forest would be overstated if it was

[28]The node size is the minimum size of final nodes within a tree.

assessed based on this hold-out set. The predictive performance should instead be assessed based on a completely independent prediction set.

Consequently, I split the data into training and hold-out sets keeping all markets affected by a given merger together in one set to avoid leakage. I split the data randomly using the R package caTools, which allows for balancing the class distribution of the dependent variable in the training and hold-out sets (Tuszynski, 2014). This is particularly relevant in this application, as competitive concerns arise in less than 15% of the concerned markets and both the training as well as the hold-out set need to contain observations of both classes. An overview of the means both of the dependent variable *Concern* as well as of the variables used as predictors in the random forests is given in Table 4.5. As the table shows the concern rate of DG Comp is similar for the respective training and prediction sets. However, the table also shows that the training and prediction sets differ in many of the observables, given that the size of the dataset is not extremely large and that I require all markets affected by a merger to be allocated to the same subset of the data. Compared to a balanced split, where the training and prediction sets are similar in all dimensions, I impose a stricter performance test on my prediction models as they also need to be able to classify a quite different dataset well.

Table 4.5: Mean Observables Training and Prediction Sets

	Pre-reform			Post-reform		
	Training mean	Prediction mean	Diff. p-value	Training mean	Prediction mean	Diff. p-value
Concern	0.14	0.14	0.99	0.13	0.13	0.99
Entry barriers	0.13	0.09	0.00	0.16	0.14	0.00
Risk of foreclosure	0.05	0.01	0.00	0.02	0.04	0.00
Full merger	0.73	0.67	0.00	0.77	0.77	0.58
Joint Venture	0.23	0.26	0.00	0.11	0.09	0.02
Conglomerate merger	0.05	0.03	0.01	0.00	0.01	0.01
Vertical merger	0.19	0.18	0.44	0.34	0.35	0.58
National market	0.68	0.59	0.00	0.66	0.59	0.00
EU wide market	0.23	0.29	0.00	0.20	0.28	0.00
Worldwide market	0.06	0.07	0.15	0.13	0.12	0.34
Number of competitors	1.81	2.04	0.00	1.93	2.47	0.00
No competitor information	0.45	0.42	0.02	0.47	0.40	0.00
EU acquirer	0.69	0.69	0.93	0.62	0.60	0.02
EU target	0.71	0.77	0.00	0.66	0.69	0.02
Indicator for July/August	0.16	0.17	0.28	0.15	0.18	0.00
Indicator for December	0.15	0.13	0.06	0.07	0.13	0.00
Joint market share	31.25	27.44	0.00	34.62	29.70	0.00
Post-merger HHI (high)	5,637.12	5,745.20	0.08	5,693.24	5,363.11	0.00
Observations	6,825.00	1,706.00	.	11,424.00	2,857.00	.

When growing a random forest, one needs to set some tuning parameters. First, I chose to grow forests containing 1,000 trees each, even though preliminary random

forests showed that the overall out-of-bag error rate already stabilizes at about 200 to 300 trees. However, since the stability of the variable importance measures (see Section 4.6) depends on the number of trees in the random forest (Liaw and Wiener, 2002), I use 1,000 trees in each of the random forests.

In order to choose both the number of covariates that are considered when splitting a node as well as the node size, I perform 5-fold cross-validation over a tuning grid both for the pre-reform as well as the post-reform random forest. I first split the respective training set (pre-reform and post-reform) into 5 folds, once again keeping all markets concerned by a particular merger together in a fold. 5-fold cross-validation then implies that a random forest is grown using four of the folds as training set and the fifth fold for evaluating model performance, permuting the folds used in training and evaluation. Hence, five different random forests are grown for each set of parameters on the tuning grid, where I tune over the node size and the number of variables considered at each split.[29] For the parameter tuning stage, I use the train function implemented in the caret package in R (Kuhn, 2008).[30]

Based on the results of this tuning stage, I choose a node size of 15 and to consider 12 variables at each split for the pre-reform random forest. The post-reform random forest has a node size of 20 but considers only 7 variables at each split. The results of the parameter tuning using 5-fold cross-validation show however that the model performance is relatively robust over the tuning range of node sizes and covariates considered at each split (see Appendix 4.8.2).[31] The final models of the pre- and

[29]For the tuning grid, I use node sizes of 1, 3, 5, 10, 15, 20, 25, 30, 35 and 40, as well as number of covariates considered at each split of 2, 4, 5, 6, 7, 8, 9, 10, 12, 14, 16, 18, 20, 22, 24, 26, 28 and 30. For classification, the default node size is one, while the default number of covariates that are considered for splitting a node is the square root of the number of covariates included in the random forest. As the default in this application would be to consider six covariates at each split, I chose a grid that is finer around the default value.

[30]Beside the tuning grid, one also needs to decide based on which summary metric the optimal model should be selected. For classification forests, such as in the present application, "accuracy" and "kappa" can be used. The default is to use accuracy, which is simply the fraction of all observations that are correctly classified. However, this metric does not distinguish the classes. In a highly imbalanced dataset with, for example, 80% class 1 observations and 20% class 2 observations, a classifier that would simply predict class 1 for all observations would not be very useful despite achieving accuracy of 80%. Kappa is similar to accuracy but is normalized at the baseline of random chance. Therefore, it compares the overall accuracy with expected random chance accuracy, where a Kappa of zero implies that the classifier performs no better than random classification. Given that the majority of mergers do not raise competitive concerns, I use Kappa to select the optimal model in the tuning stage.

[31]In particular, for the default values of a node size of one and six covariates considered at each split, the Kappa for the pre-reform random forest is about 0.5 compared to 0.52 for the optimal values chosen based on the cross-validation results. For the post-reform random forest, the Kappa for the default values is about 0.34 compared to 0.35 for the values chosen based on cross-validation. Furthermore, the two random forests using the default values for the node size and the number of covariates considered at each split also have very similar out-of-bag error rates as the random forests using the values chosen based on the cross-validation results.

post-reform random forests are estimated using the full respective training set.

Lastly, I use the "classweight" option available in the randomForest package. The reason is that the random forest algorithm is constructed so as to minimize the overall error rate. Hence, it will tend to concentrate more on prediction accuracy of the majority class, which can result in poor prediction accuracy of the minority class (Chen, Liaw, and Breiman, 2004). In my dataset, competitive concerns arise in only about 14% of the affected markets pre-reform and about 13% of the affected markets post-reform. When I trained random forests without class weights, the out-of-bag error rate for the minority class (markets affected by the merger in which DG Comp found competitive concerns) was above 30%, while the out-of-bag error rate for the majority class was close to zero. However, a prediction model predicting competitive concerns of mergers should be able to filter out problematic cases well, rather than predicting no competitive concerns more or less by default just because in more than 80% of the observations in the training data, there are no competitive concerns. The classweight option of the randomForest package is based on the idea of cost sensitive learning: as the randomForest algorithm is biased towards the majority class, we give the minority class larger weight, i.e. we put a heavier penalty on misclassifying the minority class. According to Chen, Liaw, and Breiman (2004), these weights are used in two places when training the random forests: firstly, the class weights are used to weight the goodness-of-fit criterion when deciding on splitting variables and split points within the trees; secondly, the class prediction within each terminal node is done by weighted majority vote. I use class weights of $1 : 6$ in training the random forests, as this corresponds roughly to the proportion of concern to no concern markets. Using the optimal parameters from the tuning stage, I however grew also pre- and post-reform random forests with class weights of $1 : 4$, $1 : 5$, and $1 : 7$ as a robustness check. These random forests have very similar out-of-bag error rates as the ones using class weights of $1 : 6$.

Figures 4.4 and 4.5 in Appendix 4.8.3 plot the out-of-bag error rate of the final random forests, distinguishing overall, for class 1 (Concern), and for class 2 (No concern) over the number of trees in the forest. For the pre-reform random forest, the overall out-of-bag error stabilizes at about 9% already at about 200-300 trees. The overall out-of-bag error rate of the post-reform random forest stabilizes at around 15% even earlier. The plots also show that using class weights in the random forests allows for achieving a very low out-of-bag error rate for the minority class (about 8% in the pre-reform random forest, about 5% in the post-reform random forest).

The two random forests produce a prediction of whether competitive concerns arise in a particular market affected by the merger. This prediction is based on the class vote from each tree and then classifying based on majority vote: if more than

50% of the trees in the forest predict competitive concerns for a particular observation, the random forest will predict competitive concerns. However, it could also be that DG Comp's classification threshold is different, maybe DG Comp would even predict competitive concerns if the probability of concerns arising was lower than 50%. Any classification threshold between 0 and 1 leads to a predicted outcome that, when compared to the true outcome, can be classified as correctly predicted positives (true positive), correctly predicted negatives (true negatives), falsely predicted positives (false positives), and falsely predicted negatives (false negatives). A common metric measuring prediction precision is the area under the receiver operating characteristic curve (AUC). The receiver operating characteristic curve (ROC curve) plots the model's true positive rate versus the false positive rate as the classification threshold varies over $[0, 1]$ (see Appendix 4.8.4). For the pre-reform random forest, the AUC is 0.9711, for the post-reform random forests, the AUC is 0.9512, which is very high compared to, for example, an AUC of 0.707 reported by Kleinberg, Lakkaraju, Leskovec, Ludwig, and Mullainathan (2018) or an AUC between 0.62-0.71 reported by Björkegren and Grissen (2018). An AUC of 0.5 represents random prediction, while an AUC of 1 means perfect prediction.

4.6 Estimation Results

In this section, I present and discuss the results based on the two random forests trained as explained in the previous section. Specifically, I first discuss the predictive performance of the random forests compared to the benchmark of a Linear Probability Model (LPM) estimated on the same covariates as the random forests. Secondly, I look into how the assessment of competitive concerns by DG Comp changed post-merger reform and investigate for which type of mergers and affected markets the predictions between the pre-reform and post-reform random forests differ.

4.6.1 Predictive Performance

In this section, I discuss the predictive performance of the random forests trained on the pre-reform and post-reform data. In order to assess the ability of the random forests to correctly predict DG Comp's competitive concerns, I compare the results to predictions of DG Comp's competitive concerns based on pre- and post-reform LPM models.[32]

[32]Instead of an LPM model, one could as well compare the predictive performance to a logit or probit model for competitive concerns raised by DG Comp.

Figure 4.1 shows the variable importance for the random forest trained on the pre-reform and the random forest trained on the post-reform training datasets, respectively. The plot shows the weighted total decrease in node impurity from splitting on the respective covariate, where node impurity is measured by the Gini index (see Section 4.5).

Figure 4.1: Variable Importance Plot for Pre- and Post-Reform Random Forests

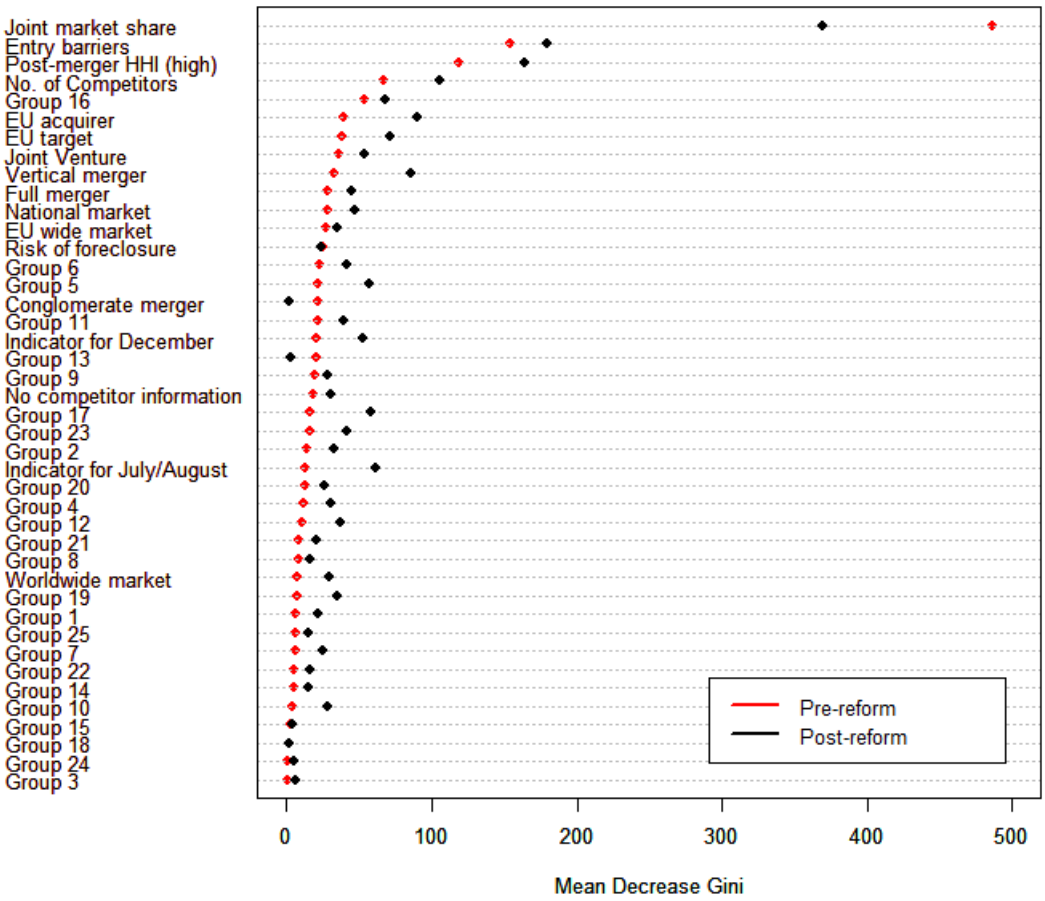

Although I am not primarily interested in the relative importance of the different determinants of a decision but in correct predictions, the variable importance plot essentially confirms the results of Affeldt, Duso, and Szücs (2019). While, in the pre-reform random forest, the joint market share of the merging parties is by far the most important covariate, its importance measured by the mean decrease in the Gini index decreases in the post-reform random forest. Furthermore, the variable importance plot shows that, while most of the same covariates show up in the upper part of the plot pre- and post-reform, the importance of covariates other than joint market share increases post-reform. In particular, the importance of entry

barriers, the post-merger HHI, the number of competitors, vertical aspects, as well as whether the merging parties are from the EU increase in the post-reform random forest compared to the pre-reform random forest. Note also that most of the industry indicator variables seem not to be very important predictors. These results are in line with Affeldt, Duso, and Szücs (2019), who, without imposing different models pre- and post-reform, find that the importance of structural indicators of market power has declined over time.

I also run simple LPM models on both the pre- and post-reform training datasets containing the same merger and market characteristics as used in training the random forests (see Appendix 4.8.5 for the regression results). I did this in order to compare the predictive performance of a parametric linear model with the predictions resulting from the highly non-linear and non-parametric random forests.

In line with the LPM results of Affeldt, Duso, and Szücs (2019), in these simple linear models, entry barriers as well as the merging parties' joint market share positively affect the probability that DG Comp will find competitive concerns in a given market affected by the merger both pre- and post-reform.[33] Furthermore, the magnitude of the estimated coefficients on entry barriers and joint market share decreases post-reform. Lastly, the R^2 is equal to 0.44 for the pre-reform LPM model but decreases to 0.28 in the post-reform LPM model, suggesting that DG Comp's merger assessment became less predictable post-reform. The generalized Hausman specification test clearly rejects the null hypothesis of no significant difference in model coefficients, indicating that separate models pre- and post-reform are in order.[34]

Both in the random forest and in the LPM models, the merging parties' joint market share, as well as entry barriers, seem to be the two most important determinants of competitive concerns by DG Comp. While in the pre-reform LPM model the coefficient on the risk of foreclosure is also positive and statistically significant, this variable is less prominent in the variable importance plots of the random forests. The results are also mostly in line with the existing literature, where the presence of entry barriers and the merging parties' combined market share are often found to be the most important predictors of an intervention decision by DG Comp (Bergman,

[33]In the pre-reform model, the risk of foreclosure as well as a merger being notified in the summer season also positively affect the probability of DG Comp finding competitive concerns, while on the other hand vertical aspects as well as an EU wide geographic market definition negatively impact this probability. Post-reform, joint ventures, conglomerate aspects as well as the indicator variable for missing competitor information negatively affect the probability of DG Comp finding competitive concerns.

[34]The generalized Hausman specification test used to compare the estimated coefficients across models is based on the joint variance/covariance matrix of the models being tested. The null hypothesis of the Hausman test is that there is no significant difference in the model coefficients. This test is clearly rejected for the two LPM models (chi-squared statistic of 62,559).

Jakobsson, and Razo, 2005; Bergman, Coate, Jakobsson, and Ulrick, 2010; Duso, Gugler, and Szücs, 2013; Mai, 2016). Mai (2016) also finds that the joint market share is a less important predictor post-reform. Importantly, Affeldt, Duso, and Szücs (2019) find these patterns over time without imposing different models pre- and post-reform.

Note, however, as explained previously, that the goal of the random forest algorithm is to correctly predict out-of-sample. Thus, I compare the predictive performance of the random forests to the predictive performance of the LPM models in the pre-reform and post-reform prediction datasets. Tables 4.6 and 4.7 show the actual as well as predicted concern rates (i.e. the percentage of markets with competitive concerns) by the pre-reform and post-reform random forest and LPM models, respectively. The models trained on the pre-reform training data should predict well the pre-reform prediction data, while the models trained on the post-reform training data should do well in predicting the post-reform prediction data.

Table 4.6: Actual and Predicted Concern Rates - RF Model

Dataset	Actual rate	Predicted rate (Pre-reform model)	Predicted rate (Post-reform model)
Pre-reform prediction set	13.6	13.8	17.4
Post-reform prediction set	12.8	14.3	24.2

Table 4.7: Actual and Predicted Concern Rates - LPM Model

Dataset	Actual rate	Predicted rate (Pre-reform model)	Predicted rate (Post-reform model)
Pre-reform prediction set	13.6	5.9	2.8
Post-reform prediction set	12.8	8.3	5.1

The actual percentage of markets in which DG Comp identified competitive concerns was 13.6% in the pre-reform prediction set and 12.8% in the post-reform prediction set. The random forest trained on the pre-reform training set predicts the concern rate in the pre-reform prediction set very well and over-predicts the concern rate in the post-reform prediction set by only 1.5 percentage points. The random forest trained on the post-reform training set over-predicts the concern rate both in the pre-reform, but particularly in the post-reform prediction set, predicting a

post-reform concern rate of 24.2%. On the other hand, both LPM models, estimated based on the pre-reform and post-reform training sets, respectively, under-predict the concern rates. In particular, the LPM model estimated based on the post-reform training set most severely under-predicts the concern rates: it predicts an intervention rate of only 2.8% and 5.1% for the pre-reform and post-reform prediction sets, respectively.

Thus, judging based on the predicted concern rates, it seems that the random forest is better than the LPM model in predicting competitive concerns pre-reform. However, it is unclear whether the random forest or the LPM model does better at predicting competitive concerns post-reform, where the random forest heavily over-predicts while the LPM model heavily under-predicts competitive concerns. Thus, I explain in detail why the random forest clearly outperforms the LPM model in the following. In particular, judging simply based on the predicted concern rate might be misleading when drawing conclusions about the predictive performance of different models. Ultimately, we should not care about whether the average concern rate is correctly predicted but in how far the model is able to correctly classify particular observations; i.e. whether the model is able to identify both class 1 and class 2 cases correctly. Tables 4.8 and 4.9 show the percentage of correct predictions for the pre- and post-reform prediction sets based on the random forest as well as the LPM models, respectively.

As discussed in Section 4.3, most existing literature only presents the overall rate of correct predictions but does not distinguish between correct predictions of class 1 and class 2. Mai (2016) reports an overall rate of correct predictions above 80% and up to 90%, Bergman, Coate, Mai, and Ulrick (2016) report about 82% correct predictions for the EU models and Mini (2018) even reports about 90% correctly predicted observations. Table 4.8 shows between 88% and 90% correct predictions in the LPM - this is therefore in line with the existing literature. However, the random forests also predict between 79% and 88% of observations correctly, which is slightly lower than for the LPM (see Table 4.9). However, at the extreme, if about 90% of the observations are class 1 and the model classifies 100% of the observations as class 1, then 90% of the cases are correctly predicted, even though the model is essentially useless as it is not able to correctly predict any of the rare class 2 cases. Consequently, I distinguish between predictions for markets without competitive concerns (*No concern*) and markets for which DG Comp found competitive concerns (*Concern*).

Looking at the results presented in Table 4.8, it is noticeable that the *No concern* outcome is predicted well by both the pre-reform as well as the post-reform random forest both on the pre-reform as well as the post-reform prediction sets. This is

Table 4.8: Percentage of Correct Predictions - RF Model

	Pre-reform model	Post-reform model
Pre-reform prediction set		
Concern	58.2	67.2
No concern	93.1	90.4
Overall	88.4	87.3
Post-reform prediction set		
Concern	38.8	61.2
No concern	89.3	81.2
Overall	82.8	78.6

Table 4.9: Percentage of Correct Predictions - LPM Model

	Pre-reform model	Post-reform model
Pre-reform prediction set		
Concern	30.2	15.9
No concern	97.9	99.3
Overall	88.7	87.9
Post-reform prediction set		
Concern	44.3	31.4
No concern	97.0	98.7
Overall	90.3	90.1

also true for the prediction of the *No concern* outcome by the pre-reform and post-reform LPM models presented in Table 4.9. The percentage of correctly predicted *No concern* observations is above 80% in all models. However, as *No concern* is the majority class in the data, it is not surprising that all the models do well in predicting this outcome.

Hence, the predictive performance of the different models should be judged based on how well these models are able to correctly classify observations of the minority class, i.e. markets in which DG Comp identified competitive concerns. Both the pre-reform and post-reform LPM models do very poorly in this respect: the pre-reform LPM model only correctly predicts about 30% of the competitive concern markets in the pre-reform prediction set, the post-reform LPM model only predicts about 31% of the markets with concerns in the post-reform prediction set correctly. The random forest models on the other hand, do significantly better in correctly predicting the minority class. The random forest trained on pre-reform data correctly predicts 58% of the competitive concern markets in the pre-reform prediction set and the random forest trained on post-reform data correctly predicts 61% of the markets

raising competitive concerns in the post-reform prediction set.

This shows that the very flexible non-parametric random forests, which are designed to maximize predictive performance, do much better in correctly classifying observations into markets with and without competitive concerns than the LPM models presented for comparison. Even though the minority class is still harder to predict than the majority class, the percentage of correct predictions of competitive concerns achieved by the random forests is doubled compared to the percentage of correct predictions by the LPM models.[35]

One reason for the difference in predictive performance between the random forests and the LPM models is that the LPM imposes a linear and additively separable relationship between the covariates and the outcome variable. However, whether a merger raises competitive concerns in a particular market might be determined by a very complex combination of market and merger characteristics that is very unlikely to be linear and additively separable. A random forest on the other hand, is able to find these potentially complex and non-linear combinations of characteristics that determine whether competitive concerns arise or not and "cuts" the high dimensional characteristic space into regions separating problematic from non-problematic markets. For example, within a given tree, concerns might arise if there are entry barriers and a combined market share above 50% but also if there are no entry barriers, market share above 60% and the merger takes place within a certain industry.

Of course, one could add various interaction terms of covariates in the LPM. However, this can quickly become a very large number of interaction terms (even here with 42 covariates included in the random forests) and it is unclear *ex ante* which of these interaction terms are actually relevant for predicting the outcome. The random forest, on the other hand, performs model selection - i.e. deciding which combinations of covariates are actually relevant for predicting the outcome - and estimation of the model at the same time.

A second reason for why the random forest performs better than the LPM in correctly classifying the rare class is that by using class weights as one of the tuning parameters, the random forest is able to give the correct classification of minority

[35] As a robustness check I also computed the percentage of correct predictions based on a pre- and a post-reform random forest using the default values for the node size (one) and the number of covariates considered at each split (six in this case). These random forests do similar in predicting the *No concern* outcome with about 95% correct predictions for the pre-reform random forest on the pre-reform prediction set and about 86% for the post-reform forest on the post-reform prediction set. They do however worse in predicting the minority class: the pre-reform random forest correctly predicts about 55% of the *Concern* outcomes in the pre-reform prediction set, the post-reform random forest correctly predicts about 47% of the *Concern* outcomes in the post-reform prediction set. Nevertheless, the two random forests still do much better in correctly classifying markets with competitive concerns than the LPM models.

class observations higher weight than the correct classification of majority class cases.

One caveat of using random forests for prediction tasks must be noted though. While random forests perform very well in predicting outcomes out-of-sample for observations that are similar in their covariates to what the random forest has seen before in the training data, random forests perform poorly in predicting outcomes for out-of-sample observations that are very different in their covariates to the cases contained in the training data. Therefore, if one expects to see out-of-sample observations that look very different from the training data, random forests might not be the optimal predictor. However, in this particular application, this is not a big issue as most of the covariates included in the random forests are dummy variables. Furthermore, for the continuous covariates of joint market share and post-merger HHI, the respective training sets contain observations with market shares ranging from 0% to 100% and post-merger HHI ranging from as low as 650 to the maximum of 10,000. The training data therefore covers the entire characteristics space.

Lastly, Table 4.8 also shows that the random forest trained on pre-reform data does well in predicting the minority class for the pre-reform prediction set but not for the post-reform prediction set. The equivalent is however not true for the random forest trained on post-reform data. Even though the post-reform random forest is trained on post-reform data only, it actually does slightly better in classifying cases from the pre-reform prediction set into concern and no concern markets than from the post-reform prediction set. Nevertheless, the post-reform random forest is still able to filter out competitive concern markets in the post-reform prediction set much better than the pre-reform random forest. This indicates that the models - and hence also the assessment rule of DG Comp - differ pre- and post-reform. I look at this aspect next.

4.6.2 Pre-Reform versus Post-Reform Predictions

While the random forest trained on pre-reform data does well in predicting outcomes for the pre-reform prediction set, the random forest trained on post-reform data does well in predicting outcomes for the post-reform prediction set. However, the pre-reform random forests does very poorly in predicting the minority outcome in the post-reform period. Based on the actual and predicted concern rates in Table 4.6, the so-called case mix and policy effects can be distinguished. Following the Oaxaca decomposition of Bergman, Coate, Jakobsson, and Ulrick (2010):

(a) The difference between the actual post-reform concern rate and the predicted post-reform concern rate using the pre-reform random forest is the policy effect: $12.8\% - 14.3\% = -1.5$ percentage points

(b) The difference between the predicted post-reform concern rate using the pre-reform random forest and the actual pre-reform concern rate is the case-mix effect: $14.3\% - 13.6\% = 0.7$ percentage points

This decomposition shows that the merger reform led to an overall slightly less aggressive merger policy by DG Comp. Mini (2018) reports an overall policy effect of about 7 percentage points, which is higher than what I find based on the predictions of the random forests. However, the markets contained in his study are not identical to those I include in my dataset, as Mini only considers horizontal overlap markets, i.e. markets in which both parties are active and where their combined market share is 15% or more.

The Oaxaca decomposition is based only on the average concern rate rather than investigating for which observations the predictions of the pre-reform and the post-reform models differ. Mini (2018) goes a step further and investigates how the importance of market shares and concentration measured by HHI differs pre- and post-reform. He reports the concern rates as well as the policy and case-mix effects by combined market share and HHI buckets, concluding that DG Comp did not change its stance towards mergers with very low market shares nor to the ones with very high combined market share (such as mergers to monopoly). However, in the middle ranges of joint market share, DG Comp challenged fewer mergers post-reform. Table 4.10 reports the difference in concern rates pre- and post-reform by markets share buckets for my dataset.[36]

Unlike Mini (2018), I find that the challenge rates are statistically significantly different from each other in all market share buckets. As in the results reported by Mini (2018), DG Comp is less likely to find competitive concerns post-reform in affected markets where the merging parties' market share is above 40%. However, I find that post-reform DG Comp is more likely to find competitive concerns in markets with merging parties' combined market share below 40%, whereas Mini (2018) finds no significant difference for markets with merging parties' market shares below 30%. The difference to the results reported by Mini (2018) could be due to the fact that I include all markets in my dataset, whereas Mini only considers horizontal overlap markets. Furthermore, the dataset by Mini includes mergers up to and including 2013, while my dataset also contains merger cases notified in 2014. However, the EC raised the threshold for a horizontally affected market from 15% to 20% combined market share in December 2013. This could partly explain the

[36]The table is based on the entire pre- and post-reform dataset, i.e. including both the observations from the training as well as the prediction sets. Note that I use the market shares as they are coded in my dataset, which is the midpoint of the interval reported in the decision documents. Therefore, I cannot be sure that actual combined market shares fall within that range. However, the table still provides an indication of the difference in concern rates pre- and post-reform.

Table 4.10: Concern Rates Pre- and Post-Reform by Combined Market Share

Market share range (%)	Pre-reform			Post-reform			Difference	
	Obs.	Fraction of obs. (%)	Concern rate (%)	Obs.	Fraction of obs. (%)	Concern rate (%)	Rate diff. (post-pre)	Two-sided p-value
<30	4,734	55.5	1.80	7,559	52.9	3.69	1.90	0.000
30 - <40	1,257	14.7	5.65	2,182	15.3	8.02	2.37	0.009
40 - <50	854	10.0	23.19	1,418	9.9	15.09	-8.09	0.000
50 - <60	580	6.8	34.66	959	6.7	27.95	-6.71	0.006
60 - <70	395	4.6	45.06	584	4.1	37.16	-7.91	0.013
70 - <80	276	3.2	57.61	456	3.2	42.76	-14.85	0.000
80 - <90	201	2.4	64.68	349	2.4	47.85	-16.83	0.000
90 - 100	234	2.7	58.55	774	5.4	40.44	-18.11	0.000
Total	8,531	100.0	13.59	14,281	100.0	12.80	-0.79	0.089

The table is based on the entire dataset, combining training and prediction sets pre- and post-reform. I use a two-sample test of proportions to determine whether the concern rates differ between the pre- and post-reform period.

higher post-reform concern rate for mergers with relatively low combined market share. Secondly, the higher post-reform concern rate for mergers with relatively low market shares also highlights one major goal of the merger reform: moving from the DT to the SIEC test made it easier to challenge mergers in markets where the merging parties are not dominant, while pre-reform the creation or strengthening of a dominant position was a necessary condition for the prohibition of a merger.

Mini (2018) only investigates the differences in concern rates as well as policy and case-mix effects between pre- and post-reform period by market share and concentration buckets. However, it could be the case that DG Comp's merger policy changed in more dimensions than just the importance of market share and concentration levels. In particular, the lower R^2 of the LPM model post-reform as well as decrease in the percentage of correct predictions post-reform also of the random forests, seems to suggest that DG Comp's assessment might be more complex and case-by-case post-reform. The fact that the post-reform random forest is still able to filter out competitive concern markets in the post-reform prediction set well and much better than the pre-reform random forest also suggests that decision criteria changed post-reform.

Table 4.11 shows the predicted *Concern* and *No concern* markets based on the two random forests. In particular, the models differ in their predictions of markets with competitive concerns. In order to investigate how the assessment of DG Comp differs pre- and post-reform, I look at the predictions for the post-reform prediction

set only and how the predictions of the pre-reform and post-reform random forests differ for the post-reform period.

Table 4.11: Actual and Predicted Concerns - RF Model

	Pre-reform model		Post-reform model	
	Predicted Concern	Predicted No concern	Predicted Concern	Predicted No concern
Pre-reform prediction set				
Actual Concern	135	97	156	76
Actual No concern	101	1,373	141	1,333
Post-reform prediction set				
Actual Concern	142	224	224	142
Actual No concern	266	2,225	468	2,023

Table 4.12 shows the differences in the post-reform predictions of the two random forests distinguishing markets where DG Comp found and did not find competitive concerns. While in 106 markets, both random forests correctly predict competitive concerns (29% of actual markets with competitive concerns), they wrongly both predict no competitive concerns in 207 markets (57% of actual markets with competitive concerns). In 1,964 markets, both random forests correctly predict no concerns (79% of actual markets with no concerns) and in 106 markets they wrongly both predict competitive concerns (4% of actual markets with no concerns). This translates to an overall agreement rate in the predictions of the two random forests of about 83%. However, the interesting cases are in particular those affected markets for which the random forests trained on pre-reform and post-reform data differ in their predictions. In 118 markets where DG Comp found competitive concerns, the post-reform random forest correctly predicts those, while the pre-reform random forest predicts no concerns (32% of actual markets with competitive concerns). Secondly, in 59 markets for which DG Comp raised no concerns, the prediction of the post-reform random forest is correct while the pre-reform random forest incorrectly predicts competitive concerns in these markets (2% of actual markets with no concerns) and in 261 markets with no concerns, the prediction of the pre-reform random forest is correct while the post-reform random forest incorrectly predicts competitive concerns in these markets (10% of actual markets with no concerns).

Consequently, I study in which dimensions the affected markets differ depending on the prediction of the pre- and post-reform random forests. Table 4.13 compares the mean of the covariates of markets for which the pre- and the post-reform random forests predict competitive concerns, respectively. The markets for which the pre-reform random forest predicts competitive concerns are significantly more likely than

Table 4.12: Differences in Post-Reform Predictions by RF Models

Predictions	Actual Concern	Actual No concern
Both pre- and post-model: Concern	106	207
Both pre- and post-model: No concern	106	1,964
Pre-model: No concern/Post-model: Concern	118	261
Pre-model: Concern/Post-model: No concern	36	59
Total	366	2,491

the markets for which the post-reform random forest predicts competitive concerns to exhibit risk of foreclosure, to be vertically affected markets and defined as being EU wide geographic markets. On the other hand, they are significantly less likely to exhibit entry barriers and to be defined as national geographic markets. They are also markets in which the merging parties have lower joint market shares than the markets for which the post-reform model predicts competitive concerns.

Table 4.13: Equality of Means Test - Predicted Concern

	Pre-reform model Concern mean	Post-reform model Concern mean	t statistic	two-sided p-value
Entry barriers	0.48	0.59	-3.75	0.000
Risk of foreclosure	0.12	0.04	4.52	0.000
Full merger	0.95	0.92	1.71	0.087
Joint Venture	0.01	0.03	-3.17	0.002
Conglomerate merger	0.00	0.02	-3.95	0.000
Vertical merger	0.33	0.21	4.14	0.000
National market	0.62	0.70	-2.68	0.007
EU wide market	0.27	0.21	2.31	0.021
Worldwide market	0.11	0.09	1.06	0.289
Number of competitors	2.97	3.27	-1.40	0.161
No competitor information	0.30	0.27	0.97	0.333
EU acquirer	0.62	0.40	7.11	0.000
EU target	0.78	0.80	-0.94	0.348
Indicator for July/August	0.18	0.28	-3.77	0.000
Indicator for December	0.17	0.08	4.13	0.000
Joint market share	46.93	51.77	-3.30	0.001
Post-merger HHI (high)	5,029.36	4,865.25	1.25	0.211
Observations	1,100			

Table 4.14 compares the mean of the covariates of markets for which the pre- and the post-reform random forests predict no competitive concerns, respectively. The markets for which the pre-reform random forest predicts no competitive concerns are significantly less likely to exhibit entry barriers or risk of foreclosure than markets for which the post-reform random forests predicts no competitive concerns. Furthermore, the mergers are less likely to be full mergers and the merging parties

have lower combined market share in these markets than in those for which the post-reform random forest predicts no competitive concerns.

Table 4.14: Equality of Means Test - Predicted No Concern

	Pre-reform model No concern mean	Post-reform model No concern mean	t statistic	two-sided p-value
Entry barriers	0.03	0.06	-5.20	0.000
Risk of foreclosure	0.01	0.04	-5.81	0.000
Full merger	0.71	0.74	-2.42	0.015
Joint Venture	0.12	0.11	1.87	0.061
Conglomerate merger	0.01	0.01	1.90	0.058
Vertical merger	0.36	0.37	-1.11	0.269
National market	0.58	0.57	0.57	0.571
EU wide market	0.28	0.29	-0.67	0.501
Worldwide market	0.13	0.13	-0.07	0.945
Number of competitors	2.31	2.34	-0.34	0.737
No competitor information	0.43	0.42	0.70	0.486
EU acquirer	0.59	0.63	-2.74	0.006
EU target	0.66	0.67	-0.70	0.485
Indicator for July/August	0.17	0.16	1.39	0.166
Indicator for December	0.12	0.14	-2.06	0.039
Joint market share	24.19	26.02	-3.38	0.001
Post-merger HHI (high)	5,469.79	5,446.05	0.33	0.738
Observations	4,614			

These differences between the predictions based on the pre-reform and post-reform random forests confirm that pre-reform DG Comp seems to have relied more on market shares in its merger assessment. Thus, the pre-reform model already predicts competitive concerns at lower market shares than the post-reform model and, conversely, the markets for which the pre-reform model predicts no competitive concerns have lower average combined market share than those markets for which the post-reform model still predicts no competitive concerns. Furthermore, the results show that post-reform DG Comp seems to assess mergers more on a case-by-case basis, where cases are not automatically considered to raise competitive concerns just because markets are concentrated, exhibit entry barriers, or are geographically narrowly defined, while mergers in broad geographic markets with relatively low market shares or vertical aspects might still raise competitive concerns. Therefore, although the overall intervention rate decreased slightly post-reform, DG Comp seems to base its assessment of competitive concerns more on a case-by-case economic analysis where a complex interaction between all merger characteristics determines the intervention decision. Of course these potentially complex interactions cannot be easily detected by a simple comparison of means as in Tables 4.13 and 4.14. However, as the prediction results show, the random-forest trained on the

post-reform data is able to find these relevant interactions of characteristics and filter out problematic markets much better than the pre-reform random forest or any simple linear model.

4.7 Conclusion

One goal of the 2004 EU merger reform was to bring merger control closer to economic principles. Another was to increase legal certainty and transparency of the merger review process as evidenced by the publication of merger guidelines and the institutional changes made. However, the effect of the reform on the predictability of DG Comp's decisions is ambiguous, as the use of a "more economic approach" in the merger review implies a shift from simple general rules, such as concentration thresholds, toward a more in depth case-by-case economic analysis. Thus, the question is whether the merger reform increased the *ex ante* predictability of decisions based on market and merger characteristics and also how the merger reform changed the decision criteria on which DG Comp bases its merger assessment.

In this paper, I study the predictability of DG Comp's merger policy and assess how it changed following the 2004 merger reform based on a comprehensive dataset covering almost all mergers notified to the EC between 1990 and 2014. Unlike most of the existing literature, rather than assessing mergers at the aggregate, the data is collected at a more fine-grained level, defining an observation as a particular product and geographic market combination concerned by a merger. This allows studying the factors that cause competitive concerns in specific sub-markets, as mergers typically concern several product and geographic markets that can be affected differently by the merger.

In addition, and unlike the existing literature studying the determinants of DG Comp's merger intervention decisions and their predictability, I use non-parametric random forests to predict DG Comp's assessment of competitive concerns arising in affected markets due to the merger. This machine learning algorithm is designed to maximize predictive performance rather than estimating causal effects and allows for highly flexible, non-linear interactions between covariates.

First, I find that the predictive performance of the random forests is much better than the performance of the LPM models I estimate for comparison reasons. While all models are able to predict the majority outcome of *no competitive concerns* very well (between 80% and 90% correct predictions), the LPM models do very poorly in predicting the minority outcome of *competitive concerns* with only between 16% and 44% correct predictions. Furthermore, based on these predictions as well as the R^2, the LPM models would wrongly suggest that the predictability of DG Comp's

assessment decreased after the 2004 reform. The random forests however, are able to correctly classify the minority class cases in about 60% of the cases both pre- and post-reform. Thus, it is not true that the predictability of DG Comp's merger policy decreases post-reform.

Secondly, I study how the predictions of the pre-reform and post-reform random forests differ for the post-reform period. Based on the two random forests and correcting for the case mix effect, the policy effect shows a decrease in the concern rate of about 1.5 percentage points post-reform. However, this decomposition only considers the average concern rate rather than investigating for which type of mergers and affected markets the predictions of the pre-reform and the post-reform random forests differ. Studying post-reform cases for which the predictions of the two random forests differ, I find that pre-reform DG Comp seems to have relied more on structural indicators, such as market shares and HHI, in its merger assessment, while post-reform DG Comp seems to base its assessment of competitive concerns more on a case-by-case analysis and less on simple structural indicators such as market shares or concentration measures. The highly flexible random forest algorithm is able to detect these potentially complex interactions between merger and market characteristics on which DG Comp's decision is based, therefore still allowing for high prediction precision compared to the overly simplistic LPM models.

4.8 Appendix

4.8.1 Summary Statistics Entire Dataset

Table 4.15: Summary Statistics Variables at Market Level (Entire Dataset)

	mean	sd	min	max	observations
Concern	0.11	0.31	0	1	30,995
Vertical merger	0.26	0.44	0	1	30,995
Conglomerate merger	0.02	0.13	0	1	30,995
National market	0.58	0.49	0	1	30,995
EU wide market	0.20	0.40	0	1	30,995
Worldwide market	0.10	0.29	0	1	30,995
Left open market	0.12	0.33	0	1	30,995
Entry barriers	0.12	0.32	0	1	30,995
Risk of foreclosure	0.03	0.16	0	1	30,995
Number of competitors	1.58	2.33	0	34	30,995
No competitor information	0.56	0.50	0	1	30,995
Joint market share	32.46	23.60	0	100	22,812
Post-merger HHI (low)	2,147.73	2,368.30	0	10,000	22,812
Post-merger HHI (high)	5,638.99	2,251.08	650	10,000	22,812

Table 4.16: Summary Statistics Variables at Merger Level (Entire Dataset)

	mean	sd	min	max	observations
Intervention	0.07	0.26	0	1	5,109
Full merger	0.55	0.50	0	1	5,109
Joint Venture	0.37	0.48	0	1	5,109
Number of concerned markets	6.07	13.43	1	245	5,109
EU acquirer	0.66	0.47	0	1	5,109
EU target	0.72	0.45	0	1	5,109
Indicator for July/August	0.18	0.39	0	1	5,109
Indicator for December	0.06	0.23	0	1	5,109

4.8.2 Random Forest Tuning Stage

The two plots show the results of the tuning of the node size and the number of covariates considered at each split using 5-fold cross-validation. For the pre-reform random forest, the highest kappa is achieved with a node size of 15 and considering 12 randomly selected covariates at each split. For the post-reform random forest, the highest kappa is achieved for a node size of 20 and considering only 7 randomly selected predictors at each split.

Figure 4.2: Parameter Tuning Pre-Reform Random Forest

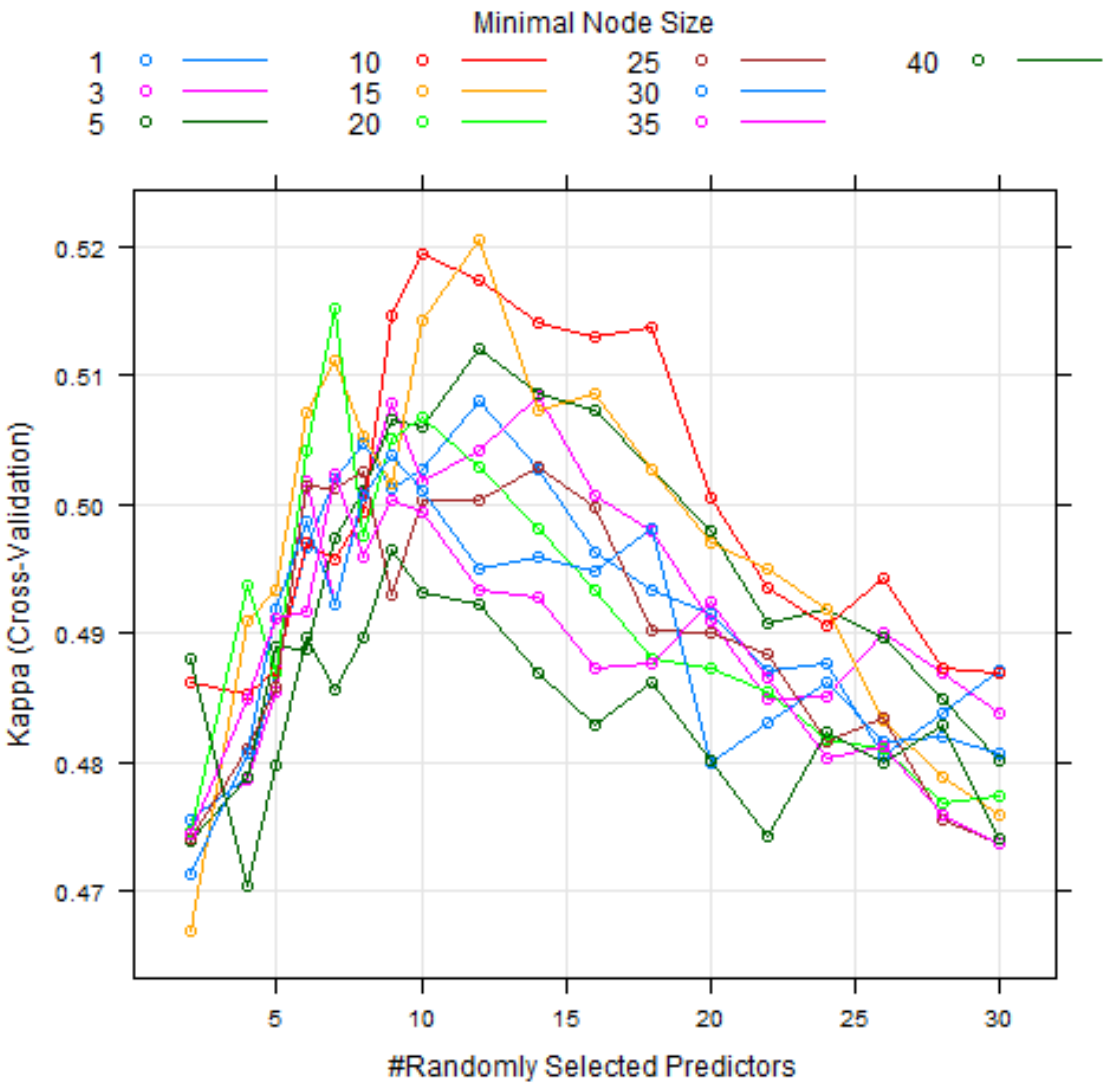

Figure 4.3: Parameter Tuning Post-Reform Random Forest

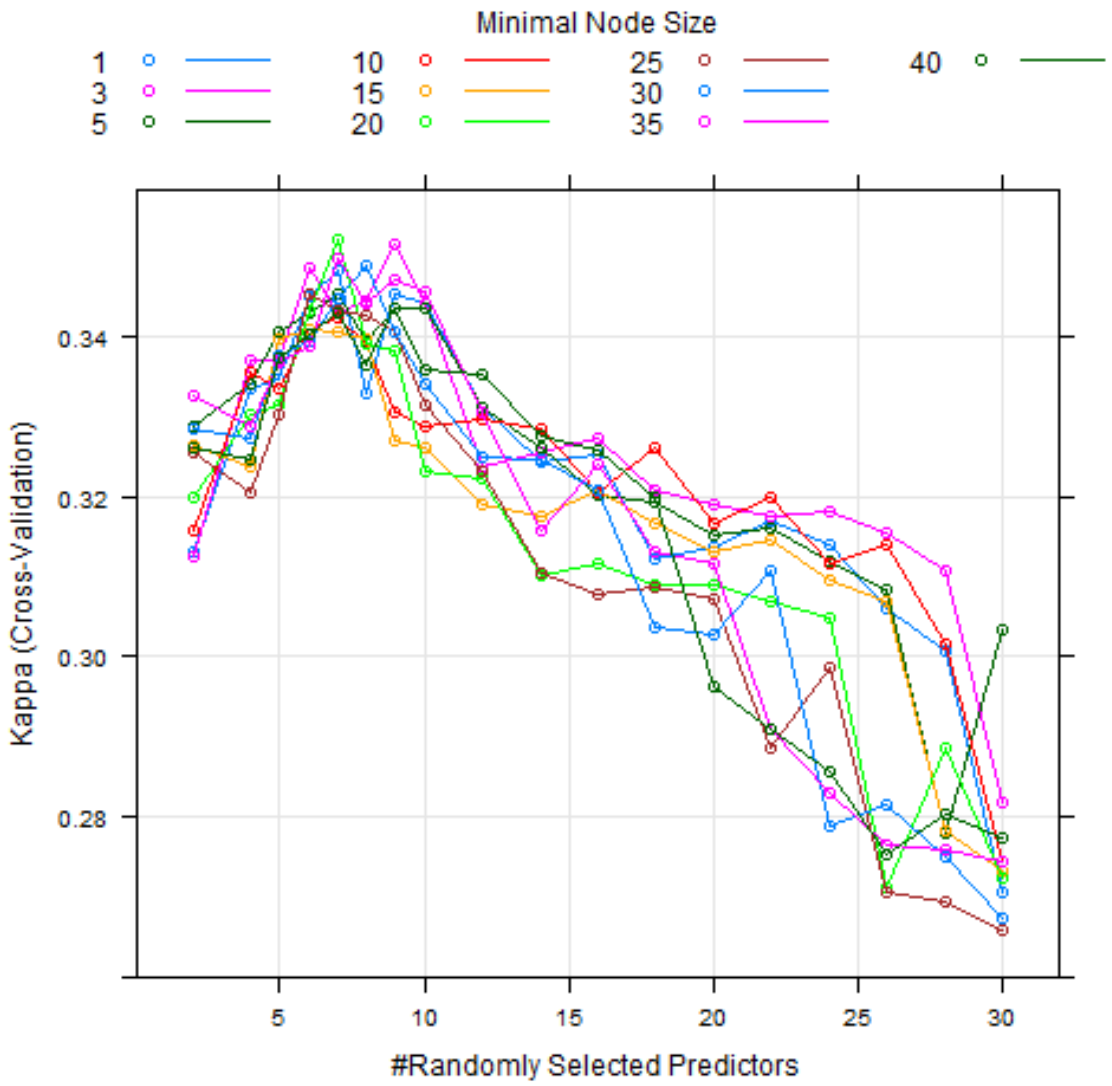

4.8.3 OOB Error Rate

The two plots show the development of the out-of-bag error rate as the number of trees included in each of the random forests increases. As random forests bootstrap the training data when constructing the individual trees, one can evaluate the prediction error for an observation by computing the mean prediction error using only the trees in the forest which do not include this particular observation - this is the out-of-bag error (OOB error).

Figure 4.4: OOB Error for Pre-Reform Random Forest

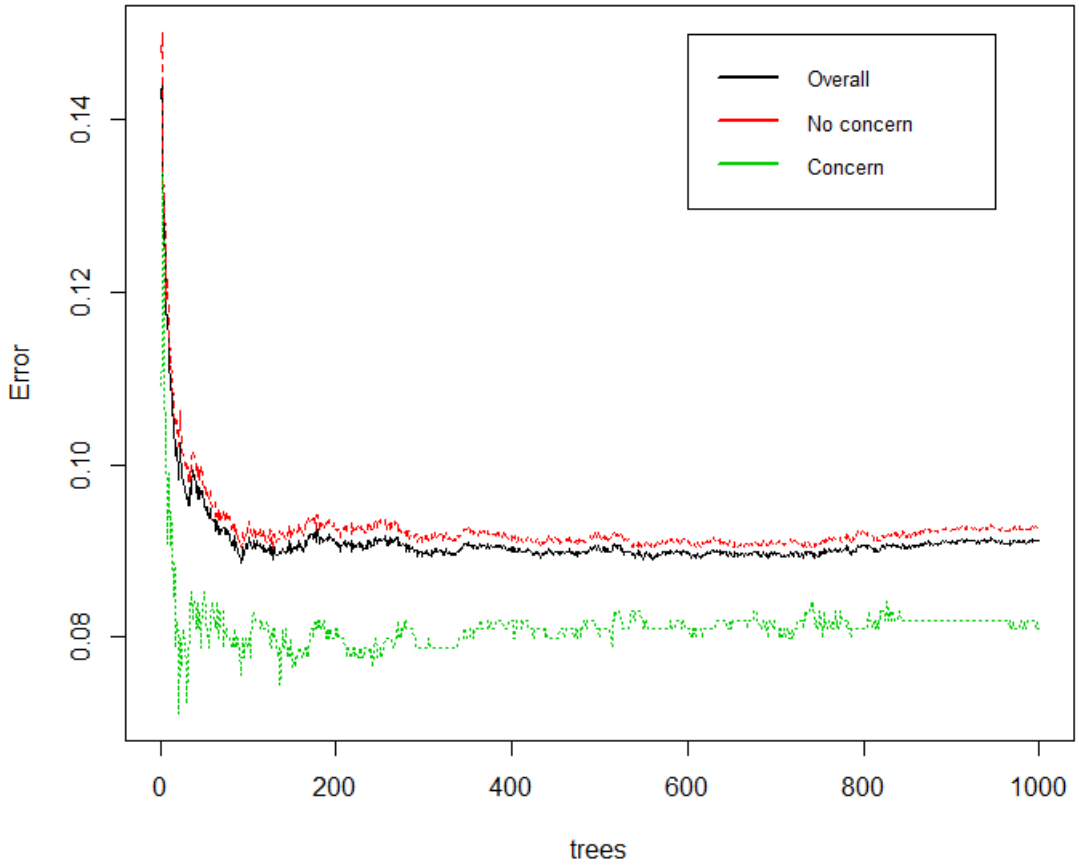

Figure 4.5: OOB Error for Post-Reform Random Forest

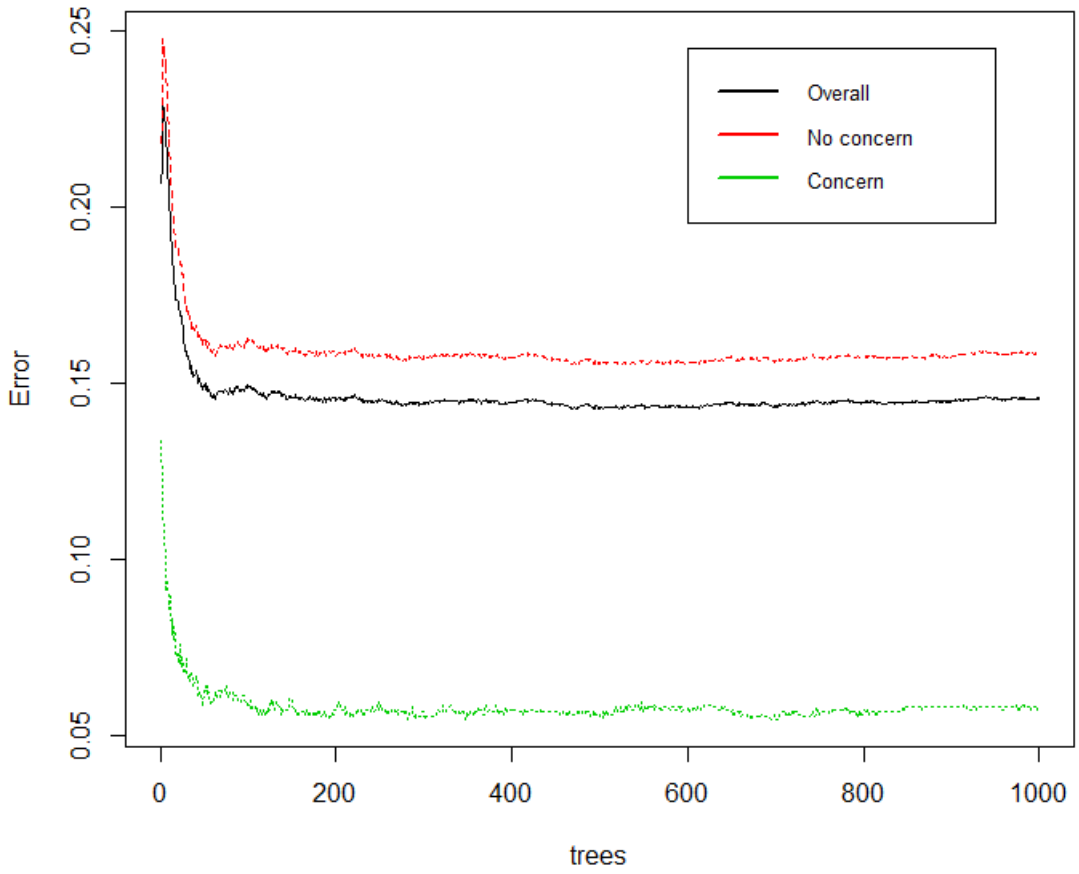

4.8.4 ROC Curve

The two plots show the receiver operating characteristic (ROC) curves for each of the two trained random forests. The ROC curve plots the model's true positive rate (sensitivity) versus the false positive rate (1-specificity) as the classification threshold varies over $[0, 1]$. While the 45 degree line represents random guessing, any point above the 45 degree line is better than random guessing. Points in the top left corner would represent perfect prediction with no false positives and no false negatives. The area under the curve (AUC) is a commonly used measure of precision. For the pre-reform random forest, the AUC is 0.9711, for the post-reform random forests, the AUC is 0.9512, which is very high compared to for example an AUC of 0.707 reported by Kleinberg, Lakkaraju, Leskovec, Ludwig, and Mullainathan (2018) or an AUC between 0.61-0.77 reported by Björkegren and Grissen (2018). Again, an AUC of 0.5 represents random prediction, while an AUC of 1 represents perfect prediction.

Figure 4.6: ROC Curve for Pre-Reform Random Forest

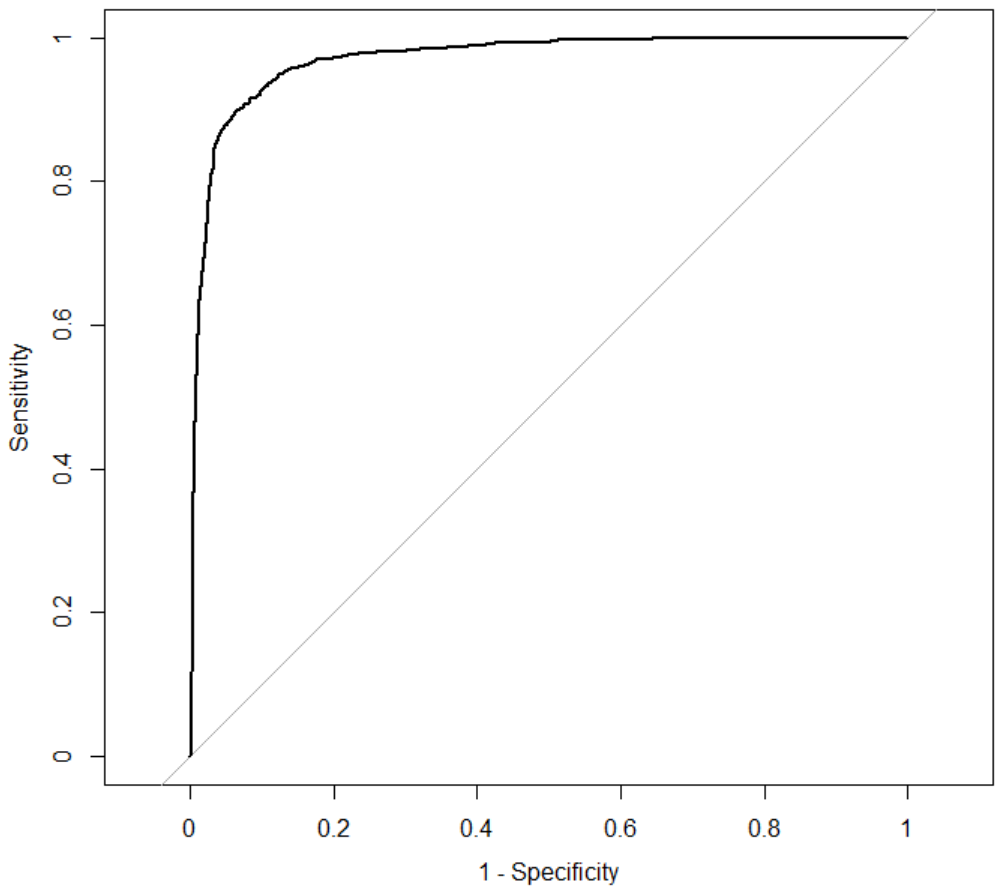

Figure 4.7: ROC Curve for Post-Reform Random Forest

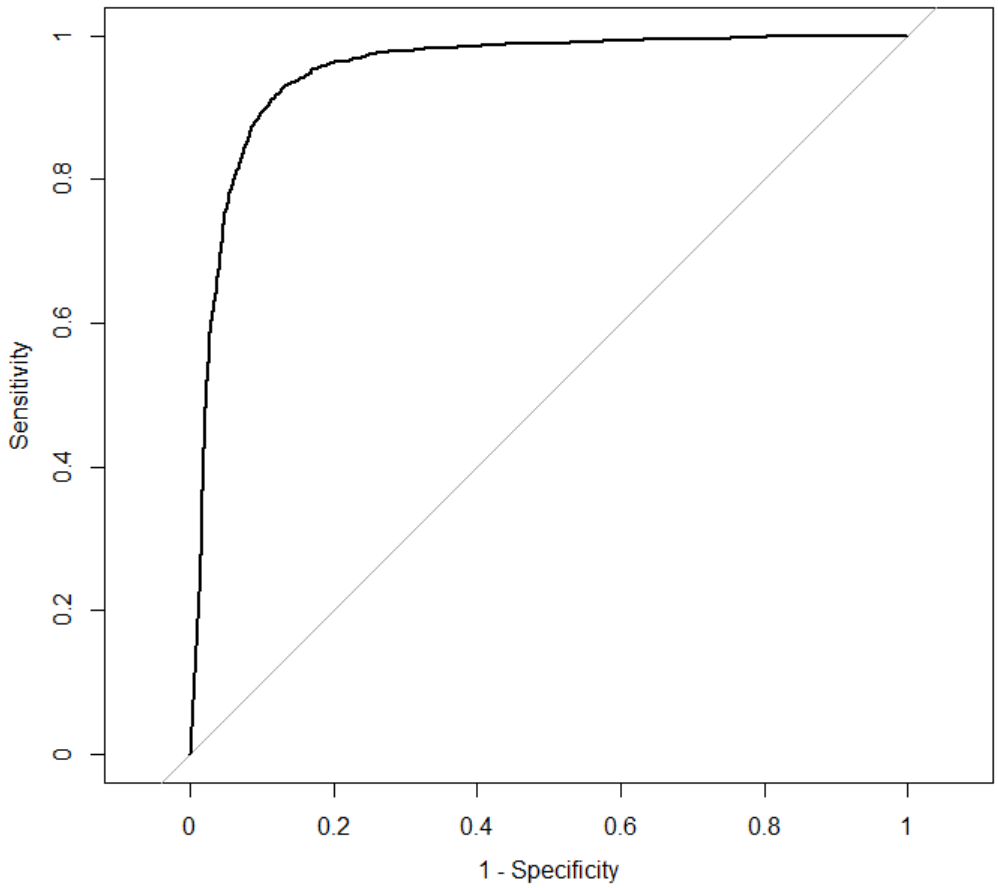

4.8.5 Linear Probability Model

Table 4.17: Linear Probability Model for Concern (Market Level)

	(1) Pre-reform	(2) Post-reform
Joint market share	0.0050***	0.0032***
	(0.0006)	(0.0005)
Post-merger HHI (high)	0.0000**	0.0000*
	(0.0000)	(0.0000)
Entry barriers	0.3927***	0.3148***
	(0.0511)	(0.0865)
Risk of foreclosure	0.1435***	0.0478
	(0.0345)	(0.1715)
Full merger	-0.0019	-0.0319
	(0.0363)	(0.0260)
Joint Venture	-0.0466	-0.0846**
	(0.0320)	(0.0308)
Conglomerate merger	-0.0024	-0.1008**
	(0.0308)	(0.0469)
Vertical merger	-0.0359**	0.0136
	(0.0148)	(0.0326)
National market	-0.0164	-0.1224
	(0.0153)	(0.0914)
EU wide market	-0.0309***	-0.0861
	(0.0099)	(0.0917)
Worldwide market	0.0404	-0.1373
	(0.0337)	(0.1093)
Number of competitors	-0.0042	-0.0008
	(0.0036)	(0.0050)
No competitor information	-0.0408	-0.0637***
	(0.0328)	(0.0213)
EU acquirer	0.0319*	0.0518
	(0.0155)	(0.0347)
EU target	0.0111	-0.0045
	(0.0186)	(0.0288)
Indicator for July/August	0.0533**	-0.0010
	(0.0227)	(0.0396)
Indicator for December	-0.0102	0.1070**
	(0.0412)	(0.0496)
Constant	-0.0445	0.1160
	(0.0590)	(0.1201)
Industry Group FE	Yes	Yes
R2	0.439	0.280
Observations	6,825	11,424

We report heteroskedasticity robust standard errors clustered at the industry group level. Significance at the 1%, 5%, and 10% levels is represented by ***,** and * respectively.

Chapter 5

Estimating Demand with Multi-Homing in Two-Sided Markets [1]

5.1 Introduction

Two-sided markets are markets in which firms sell two products or services to two different types of consumers, taking into account that the two demands are linked by indirect network effects. Examples of such markets are media markets, where demand for advertising is related to the size of the audience, and the market for online social networks, where advertising demand depends on the number of users.

A media firm typically operates in a two-sided market as it sells content to readers/viewers and advertising space to advertisers. Moreover, it knows that the size (and possibly the characteristics) of the audience influences the demand for advertising space and, vice versa, the amount (or concentration) of advertising might influence the audience's demand. In other words, a media company recognizes the existence of indirect network effects between the two sides of the market when making its strategic decisions.

With the emergence of digital technologies, multi-homing has become a widespread phenomenon in media markets. In fact, the cost of multi-homing for consumers of media content has dramatically dropped. For instance, newspaper readers can now access multiple online news outlets with just a few clicks, no longer needing to buy

[1]This chapter is based on joint work with Elena Argentesi and Lapo Filistrucchi. We thank Lukasz Grzybowski, Alessandro Iaria, Leonardo Madio, Andrea Mantovani, Christine Zulehner, as well as participants at the 2018 Media Economics Workshop, the 4th Economics of Media Bias Workshop, the 2019 Annual Scientific Seminar on Media and the Digital Economy, the 2019 MaCCI Annual Conference, and the 11th Paris Conference on Digital Economics for helpful comments.

and carry a pile of different newspapers, then physically leafing through them. Similarly, TV viewers now have access to many more digital channels. Thus, effectively consumers more frequently multi-home.

Hence, advertisers are now able to reach consumers over a greater number of outlets during their preferred time period (whether a day, a week or a month). This has implications for the willingness to pay of advertisers for reaching consumers. For instance, considering the extreme case in which an advertiser wishes to reach its target consumers only once, the value of second impressions would be zero. Thus, a merger between newspapers read by two distinct set of readers (i.e. whose readers single-home) could have very different effects on the prices charged to advertisers than a merger of newspapers with perfectly overlapping readers (i.e. whose readers multi-home). In turn, due to two-sidedness, the cover prices of the newspapers may also be affected differently.

Therefore, allowing for multi-homing is crucial when devising policy decisions, whether in assessing mergers among media outlets, in regulating cross-ownership of media, or in setting advertising limits.

We empirically study the role of multi-homing in two-sided markets. First we build a micro-founded structural econometric model, encompassing the demand for differentiated products on both sides of the market and allowing for multi-homing on each side. We then estimate the model, using data covering Italian daily newspapers, alternatively taking into account and not taking into account information on multi-homing by readers. We show that not accounting for multi-homing leads to a substantial bias in the estimation of own- and cross-price elasticities on the readers' side of the market. In particular, mean own-price elasticities increase from between -1.27 and -1.48 when readers are assumed to single-home to between -1.29 and -4.13 when reader multi-homing is taken into account. Furthermore, while newspapers are assumed to be substitutes in the single-homing model, they can be substitutes or complements when multi-homing by readers is taken into account. We find that while newspapers of the same type are substitutes, newspapers of different types are complements. We also show that, on the advertising side of the market, own-price elasticities decrease with the amount of captive readers while cross-price elasticities increase with the amount of overlapping readers between newspapers.

Hence, disregarding multi-homing is likely to bias the conclusions of exercises like market definition or merger evaluation, in which own- and cross-price elasticities play a crucial role. Lastly, we discuss to what extent disregarding multi-homing information may bias policy decisions, particularly in the field of competition policy. Specifically, we consider the traditional product market definition in markets for daily newspapers, which distinguishes generalist, sport, and financial newspapers.

This market definition is potentially affected by the use of information on multi-homing. While a full-fledged market definition exercise would require performing a two-sided SSNIP test, our findings confirm the importance of incorporating multi-homing in the analysis. We find that, on the readers' side, only newspapers of the same type (generalist, sport, or financial) may be in the same relevant market (although they are not necessarily). On the advertisers' side of the market, instead, it appears that only newspapers of similar advertising importance may be in the same relevant market.

Our paper contributes to the economic literature on two-sided markets, in which empirical work accounting for multi-homing on both sides of the market is still quite scarce (see next section for a discussion). Moreover, our contribution allows for a better understanding of the implications of multi-homing and, therefore, is useful for competition and regulation authorities seeking to improve their quantitative assessment in cases involving two-sided platforms. Although print newspapers are a classic example of an *offline* two-sided market, the empirical part of this paper should be seen as an approach to studying the role of multi-homing in (non-transaction) platform markets. The methodology can also be applied to other two-sided markets for which data on user multi-homing is available. In light of the prevalence and rising importance of multi-sided platforms in digital markets and the pervasiveness of multi-homing by users on online platforms, the results and conclusions from this paper are also relevant in the context of competition policy cases involving online multi-sided platforms.

The paper proceeds as follows. Section 5.2 discusses the literature on multi-homing in two-sided markets. Section 5.3 provides an overview of the market for daily newspapers in Italy and the data we use. Section 5.4 presents a two-sided model of demand allowing for multi-homing on both sides, while we discuss estimation results in Section 5.5. Section 5.6 highlights the importance of accounting for multi-homing in market definition. Section 5.7 concludes.

5.2 Multi-Homing in Two-Sided Markets

Following the seminal works of Caillaud and Jullien (2001, 2003), Rochet and Tirole (2002, 2003, 2006), Parker and van Alstyne (2005), and Armstrong (2006), a growing number of papers has dealt with the theoretical aspects of two-sided markets, such as Anderson and Gabszewicz (2006) on media markets. Some, including Evans (2003), Wright (2004), Evans and Schmalensee (2007), Filistrucchi, Klein, and Michielsen (2012), and Filistrucchi, Geradin, Van Damme, and Affeldt (2014) have focused on competition policy in two-sided markets.

So far, most of the theoretical literature on two-sided markets has assumed single-homing on at least one side of the market, most often on the readers'/viewers' side in media markets. In this context, the competitive bottleneck problem of Armstrong (2006) arises, whereby each media outlet is a monopolist over providing access to its exclusive audience and, thus, advertisers must patronize all of them in order to reach all consumers. Only recently, as discussed in Anderson and Jullien (2015), the theoretical literature has started filling this gap, e.g. Ambrus, Calvano, and Reisinger (2016), Anderson, Foros, and Kind (2018), Athey, Calvano, and Gans (2018), and Jeitschko and Tremblay (2018). The fact that a fraction of consumers patronizes more than one platform changes the model predictions quite dramatically. If advertisers can reach multi-homing consumers on more than one platform, media outlets no longer only compete for consumers on the audience side of the market but also compete for advertisers on the advertising side of the market. In particular, it turns out that "each platform is able to price to advertisers only the value of its exclusive consumers plus the incremental value associated with multi-homing (shared) consumers" (Anderson, Foros, and Kind (2018), p.35). This so-called "principle of incremental pricing" has important implications for platforms' strategies in terms of pricing, reaction to mergers and content provision.

However, empirical work has lagged behind in accounting for multi-homing. Starting from the seminal papers of Rysman (2004) on the market for yellow pages and Kaiser and Wright (2006) on the German magazine market, most empirical contributions have assumed single-homing at least on one side of the market, typically the audience side. For example Filistrucchi, Klein, and Michielsen (2012) and Affeldt, Filistrucchi, and Klein (2013), while using data on the Dutch daily newspaper market to simulate the unilateral effects of mergers, do not allow for multi-homing; similarly for Filistrucchi and Klein (2013). More recently, Ivaldi and Muller-Vibes (2018) estimate a two-sided nested logit model of demand for the print media industry in France, but lack information on readers' multi-homing behavior.

While Rysman (2007) has shown that multi-homing in adoption, but not in usage, is an important feature of the payment card market, only Fan (2013) moved a step forward by allowing each household to buy up to two newspapers in the econometric model. Yet, Fan (2013) lacks information on double-readership of newspapers at the household level and, therefore, cannot estimate a model with multi-homing readers. We do have this information at the individual newspaper level and use it to analyze the impact of allowing for multi-homing readers on the empirical results. Gentzkow (2007) develops a methodology that allows for the consumption of two products in order to study competition between print and online newspapers. The same demand model is also applied by Gentzkow, Shapiro, and Sinkinson (2014), which show that

advertising competition leads to increased ideological diversity. Unlike Gentzkow (2007), which has individual-level data on readership for a small set of newspapers, we have aggregate data for a larger set of newspapers. Thus, we build a nested logit demand model encompassing the multi-homing of readers by allowing them to choose between bundles of newspapers. Also Shi (2015) accounts for readers' multi-homing in the estimation of demand for U.S. magazines, finding that advertising prices are related to the share of exclusive versus overlapping readers. However, he has data on readers' multi-homing just for one period, while we have much richer survey data including bi-annual information from 1992 through 2006. Finally, a recent paper by Liu (2018) estimates the effect of consumer multi-homing on prices in the online advertising market.

This paper builds on Argentesi and Filistrucchi (2007), where the authors test for market power in the national daily newspaper market in Italy. However, that paper, lacking information on multi-homing, assumes that readers do not multi-home, i.e. that they read only one newspaper. Both the structural econometric model and the estimation are conducted under this assumption. Moreover, the analysis is conducted on a smaller sample of newspapers (i.e. only the national generalist newspapers) and over a shorter time. Finally, their dataset on the advertising side of the market is much less detailed than the one we use in this paper.

5.3 Data

The dataset contains information on seven national daily Italian newspapers, belonging to three different categories: general interest, sport, and financial newspapers.[2] The four general interest newspapers are *Corriere della Sera*, *La Repubblica*, *La Stampa*, and *Il Giornale*. The two sport newspapers are *Corriere dello Sport* and *Gazzetta dello Sport*. The financial newspaper is *Il Sole 24 Ore*. In December 2006, these seven newspapers accounted for more than 40% of overall circulation of daily newspapers in Italy. In particular, in the sub-market of general interest newspapers, *Corriere della Sera* and *La Repubblica* were, and still are, the largest players in terms of circulation. Other newspapers are not included in our dataset because their circulation is mainly regional (e.g. *Il Messaggero* or *QN*) or much smaller than those in our sample (e.g. *Avvenire*). As for sport newspapers, *Gazzetta dello Sport* and *Corriere dello Sport* are the largest outlets, with more than 80% of copies in

[2]This segmentation of the newspaper market has been adopted in several antitrust decisions, both in Italy and across the European Union. See, for instance, Italian case 3354/95 *Ballarino vs. Grandi Quotidiani* and the European Commission's decisions M.3817 Wegener/PCM/JV and M.1401 Recoletos/Unedisa.

this segment. Finally, *Il Sole 24 Ore* is by far the main financial paper, retaining more than 80% of the market segment in terms of copies sold.

On the readers' side, the dataset features monthly observations for each newspaper on each day of the week from 1992 through 2006. Market-level data on circulation come from those collected for advertising purposes by Accertamenti Diffusione Stampa (ADS).[3] Specifically, we use monthly average printed copies for each day of the week as a proxy for circulation, since information on the number of copies sold in each day of the week is not available in this dataset. Indeed, it is important to have information disaggregated by day of the week because some weekly supplements are bundled with the newspapers only on some days of the week and cover prices vary by day of the week. We collected information from newspaper publishers on the cover prices of the newspapers and on content characteristics such as the dates regular supplements were introduced, the changes of editors, the presence of local news sections, and the dates newspapers' websites were established.

Information on multi-homing by readers (i.e. on how many readers of a given newspaper also read each of the other newspapers) was collected, for advertising purposes, by Audipress in bi-annual surveys.[4] In particular, the survey asks readers which newspapers they read on an average day. Then, for each newspaper, it computes the number of readers that read only that newspaper as well as the number of readers that also read each of the other newspapers. However, we do not know whether readers of that newspaper, who also read another newspaper, overlap with readers double-homing on a third newspaper. Thus, we only refer to double-homing in the following as we cannot identify those readers who read more than two newspapers from the data. Note that, to the extent that they do not carry out additional surveys themselves, this information comprises all that advertisers know about single-homing or multi-homing by readers.[5]

Table 5.1 shows, by newspaper, the percentage of single- and double-homing readers. Depending on the newspaper, on average between 25% and 62% of the readers single-home, i.e. only buy this specific newspaper. Whether readers single-home or double-home also seems to depend on the type of newspaper: while many readers single-home on a general interest newspaper, only 25% of the readers of the financial newspaper Il Sole 24 Ore single-home. The Table also shows on which newspapers readers double-home. Thus, the second line shows that, on average, 14.9% and 14.8% of Corriere's readers also buy Gazzetta dello Sport and La Repubblica, respectively. The sixth line shows that 21% of Il Sole 24 Ore readers also read

[3]See http://www.adsnotizie.it/.

[4]See http://www.audipress.it/.

[5]In the Audipress survey, readers of a newspaper are defined as those who read or leaf through that newspaper at least once a day.

Corriere della Sera or La Repubblica. Figure 5.1 represents the information on the percentage of readers single- and double-homing graphically. There is one column for each newspaper. The dark-blue area at the bottom of each column represents the percentage of single-homing readers, while all colored areas above it represent the percentage of multi-homing of readers of the given newspapers on each of the other newspapers.

Finally, information on the total population above 14 years of age (considered traditionally as the set of potential readers of newspapers) was obtained from ISTAT, the Italian Statistical Office. We used this to calculate newspapers' market shares.

Table 5.1: Percentage of Readers Single- and Double-Homing by Newspaper

Newspaper	Single-Homing	DH Corriere	DH Corriere Sport	DH Gazzetta Sport	DH Giornale	DH Repubblica	DH Il Sole	DH Stampa
Corriere	45.8	.	4.4	14.9	5.8	14.8	10.2	4.1
Corriere Sport	41.5	7.8	.	29.5	3	10.8	4.5	3.1
Gazzetta Sport	50.1	13.1	14.7	.	3.5	8.7	5	4.9
Giornale	29	20.4	5.9	13.8	.	12	12	6.9
Repubblica	51.4	14.6	6.1	9.7	3.3	.	9.7	5.3
Il Sole	25	21.4	5.3	11.9	7.1	20.6	.	8.6
Stampa	61.6	6.9	2.9	9.5	3.3	8.9	6.9	.

Mean percentage of readers single-homing and double-homing over the years 1992-2006.

On the advertisers' side of the market, the dataset contains market-level data on advertising quantity, prices, and reader characteristics of those same newspapers, with monthly observations for each different day of the week from 1992 through 2006. Data on advertising quantities and advertising prices net of discounts come from the database of Nielsen Media Research, while data on readers' demographics come from those collected by Audipress. The latter data are collected bi-annually. However, to the extent that they do not carry out additional surveys themselves, advertisers can be expected to base their advertising decisions on this demographic information for a period of sixth month.

Information on the total number of (advertising and non-advertising) pages per newspaper also comes from Nielsen Media Research, while information on the size of a newspaper was collected browsing on the internet. In combination with information on the price of the paper used to print daily newspapers, collected from Camera di Commercio di Milano, this data allows us to calculate the paper input cost per page and per printed copy.

Figure 5.1: Single- and Double-Homing by Newspaper

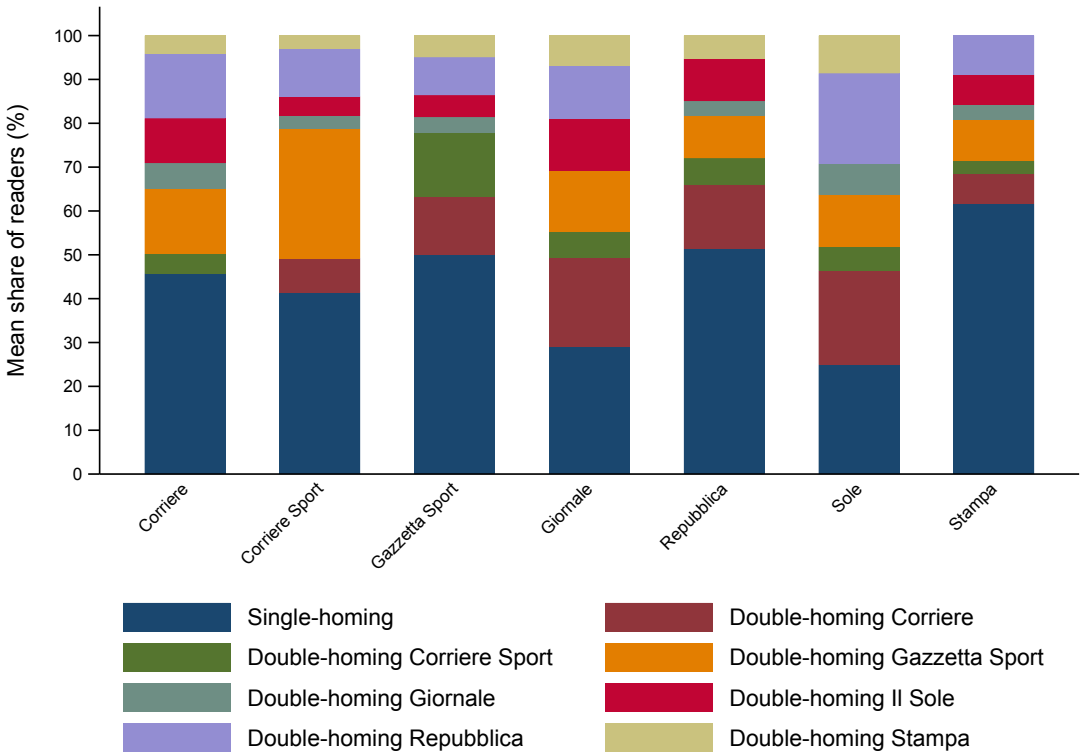

Mean percentage of readers single-homing and double-homing over the years 1992-2006.

Table 5.2 presents summary statistics on the variables we use in the estimation of the readers' side of the market. The average daily circulation of the newspapers included in the dataset is about 560,000 copies, while the mean real cover price over the sample period is about 0.96 Euros per copy.

Table 5.3 presents instead summary statistics for the variables we use in the estimation of the advertising side of the market. While the mean advertising expenditure share is 14%, the average real advertising price is about 180 Euros per advertising slot.

Importantly the variable "captive readers" measures the percentage of single-homing readers for each newspaper. This measure is crucial in order to properly account for multi-homing by readers: the more readers single-home, the higher is the market power of the newspaper on the advertising side of the market as newspapers enjoy monopoly power over providing access to these captive readers (see Armstrong (2006)).

Hence, the final datasets both on the reader side as well as on the advertising side of the market cover monthly observations for each different day of the week for the seven newspaper from January 1992 through December 2006. Appendix 5.8.1

Table 5.2: Summary Statistics Reader Side, 1992-2006

	mean	sd	min	max
Average newspaper's prints (10k)	56.76	20.19	21.12	127.43
Market share	0.01	0.00	0.00	0.03
Real cover price (EUR/copy)	0.96	0.12	0.79	1.60
Number of pages	40.43	13.00	16.50	115.00
Number of advertising slots (k)	3.96	2.19	0.19	16.48
Advertising intensity (slots/pages)	95.12	36.59	7.23	249.87
Generalist magazine	0.56	0.50	0.00	1.00
Generalist magazine (day)	0.07	0.26	0.00	1.00
Women magazine	0.21	0.40	0.00	1.00
Women magazine (day)	0.03	0.17	0.00	1.00
Economic insert	0.29	0.45	0.00	1.00
Economic insert (day)	0.04	0.20	0.00	1.00
Local pages	5.45	5.64	0.00	22.00
Website	0.50	0.50	0.00	1.00
Real paper cost (EUR/copy)	0.11	0.04	0.04	0.32
Observations	8,795			

Table 5.3: Summary Statistics Advertiser Side, 1992-2006

	mean	sd	min	max
Advertising expenditure share	0.14	0.10	0.00	0.51
Real advertising slot price (k EUR)	0.18	0.09	0.03	0.59
Percentage of readers between 14 and 17	4.76	2.88	0.84	12.87
Percentage of readers between 18 and 24	13.04	4.23	5.40	21.68
Percentage of readers between 25 and 34	21.52	3.22	15.90	30.40
Percentage of readers between 35 and 44	19.46	2.78	13.70	26.60
Percentage of readers between 45 and 54	17.36	2.59	11.78	23.40
Percentage of readers between 55 and 64	12.56	2.57	7.30	18.83
Percentage of readers above 65	11.30	4.91	3.20	23.60
Percentage of readers in low income group	12.71	7.26	2.90	32.20
Percentage of readers in middle income group	61.35	4.87	49.25	72.20
Percentage of readers in high income group	25.94	9.91	9.03	46.60
Percentage of female readers	31.66	12.26	9.00	46.50
Percentage of captive readers	43.49	12.80	9.49	66.50
Observations	8,820			

contains a list of all the variables used in our empirical analyses together with the corresponding data sources.

5.4 Demand Model

The structural econometric model encompasses demand for differentiated products on both sides of the market and allows for multi-homing on each side of the market. We estimate both readers' demand and advertisers' demand taking into account the inter-market network effects that characterize two-sided markets.

On the readers' side of the market, demand derives from random utility maximization by readers and is estimated using a nested logit model, as in Berry (1994). The structure of the nests draws on the traditional classification of national daily newspapers between generalist, sport, and financial. On this side of the market, we have information on multi-homing. When taking into account this information, readers are allowed to choose between all possible newspaper-pairings and nests are designed accordingly as combinations of newspapers of the same or of different categories.

On the advertisers' side of the market, demand derives from advertisers' choice to allocate a given advertising budget, which changes with the business cycle, across different newspapers. This is similar to consumers allocating a given budget among different types of beers in Hausman, Leonard, and Zona (1994). Product differentiation is interpreted in the spatial sense proposed by Pinkse, Slade, and Brett (2002), as applied parametrically in Slade (2004) and in Pinkse and Slade (2004). Hence, cross-price elasticities among two products (in our case advertising slots in two different newspapers) are assumed to be a function of the distance among the products in characteristic space, so that elasticities would be higher when products are closer to each other in terms of characteristics. In our application, the distance metrics are derived from differences among newspapers in the demographic characteristics of readers. In addition, own-price effects are allowed to depend on readers' characteristics. While our model also allows for advertisers to multi-home, we do not have, and hence do not use, data on multi-homing by advertisers. However, consistently with the theoretical models of two-sided markets, the information on multi-homing by readers can be used also in the estimation of advertising demand. In particular, we derive distance metrics from the number of overlapping readers between two newspapers and the number of captive readers is considered as an additional newspaper characteristic from the point of view of advertisers. Finally, in applying the distance metrics model to a two-sided market such as the newspaper market, we allow advertising demand on a newspaper to depend on its circulation.

5.4.1 Readers' Demand

On the readers' side of the market, demand derives from random utility maximization by readers and is estimated using a nested logit [6] model as in Berry (1994). Hence, reader i at time t in weekday d chooses one unit of newspaper $j \in J$ to maximize utility

$$u_{ijtd} = \alpha p_{jtd} + \beta x_{jtd} + \gamma a_{jtd} + \xi_{jtd} + \zeta_{gtd} + (1-\sigma)\varepsilon_{ijtd}, \tag{5.1}$$

where p_{jtd} is the cover price of newspaper j at time t in weekday d, x_{jtd} is a set of observed newspaper characteristics, a_{jtd} is the advertising intensity in newspaper j at time t in weekday d, ξ_{jtd} is an unobserved (by the econometrician) product characteristic, ζ_{gtd} represents consumer utility common to all newspapers of nest g at time t in weekday d, and ϵ_{ijtd} is an idiosyncratic error term assumed to be i.i.d. extreme value type 1. σ measures the correlation of unobserved utility between newspapers within nests relative to the between ones. As σ approaches one, newspapers within a nest approach being perfect substitutes, if σ is instead equal to zero, the correlation of unobserved utility within nests is zero and we are back to the simple logit case.

The structure of the nests draws on the traditional classification of national daily newspapers into generalist, sport, and financial newspapers. As discussed in Section 5.3, we have data on four general interest newspapers, two sport newspapers, and one financial newspaper. Figure 5.2 in Appendix 5.8.3 shows the structure of the nests under the single-homing assumption.

The resulting baseline estimating equation of the nested logit model is the following:

$$\ln(s_{jtd}) - \ln(s_{0td}) = \alpha p_{jtd} + \beta x_{jtd} + \gamma a_{jtd} + \sigma \ln(s_{jtd|g}) + \phi_{jd} + \tau_{tg} + \nu_{jtd}, \tag{5.2}$$

where s_{jtd} is the market share of newspaper j at time t in weekday d, s_{0td} is the market share of the outside good, p_{jtd} is the newspaper's cover price, x_{jtd} is a set of observed newspaper characteristics, a_{jtd} is the advertising intensity in newspaper j, and $s_{jtd|g}$ is the share of newspaper j within nest g. We split the unobserved product characteristic ξ_{jtd} into the newspaper-weekday fixed effect ϕ_{jd}, the time fixed effects τ_{tg}, and the i.i.d. error term ν_{jtd}.

The newspaper-weekday fixed effects capture the unobserved product characteris-

[6] We also estimated a random coefficient logit model. However, the random coefficients were not estimated to be significant.

tics that are constant over time. The time fixed effects, specific to each nest, capture variations in time in the relative attractiveness of the outside good with respect to the different categories of newspapers in our sample (for instance because of the appearance of internet, which is here allowed to have differential effects on sales of sport and generalist newspapers). The market shares are defined over the total potential market size, which is considered to be the total population in Italy older than 14 years, as is usual in studies on media markets.

The newspaper cover price, the advertising intensity, as well as the within nest market share are all endogenous as they might be correlated with the unobserved product characteristic ξ_{jtd}. Following Berry, Levinsohn, and Pakes (1995) and Nevo (2000), we use the sum of the characteristics of the other newspapers as instruments for the newspaper cover price. In particular, we use the dummy variables for the weekday when a supplement is bundled to the newspaper, when a women magazine is bundled to the newspaper, as well as the number of local pages to construct the BLP instruments for newspaper cover price. The within nest market share is instrumented with the corresponding BLP instruments within a given nest. As described in Section 5.3, we also construct a marginal cost measure, the real paper cost per copy, which varies over newspapers and over time, that is used as an instrument for cover price and advertising intensity. Lastly, we use Italian gross domestic product (GDP) to instrument advertising intensity, as GDP is a measure of the overall business cycle and advertising expenditures by companies increase with the business cycle, whereas, given the low price for a copy of a newspaper, income effects from the business cycle are not expected to be substantial and to directly affect readers' demand.[7]

We aim to investigate the bias that can take place in the estimation of demand parameters, in particular price elasticities and indirect network effects, when neglecting readers' multi-homing. In order to assess the relevance of this bias, we estimate two different specifications of readers' demand. The traditional demand equation (similar to Argentesi and Filistrucchi (2007)) assumes that readers only read one newspaper in each period, i.e. all readers single-home. Therefore, we estimate the nested logit model *at the newspaper level*: j in equations (5.1) and (5.2) refers to a newspaper. We also estimate a second, alternative demand equation where readers are allowed to read up to two newspapers (which is what we observe in the data). Thus, we model readers as choosing between all possible newspaper-

[7]See for example van der Wurff, Bakker, and Picard (2008) for an empirical study on the link between economic growth, measured by GDP, and advertising intensity and expenditures for different media and in different industrialized countries. In particular, the results show that advertising expenditures in newspapers respond, relatively closely, to economic change, while the link is weaker for TV, radio, and cinema. The paper also contains a comprehensive review of the existing literature establishing the relationship between advertising spending and economic growth.

pairings (including single-homing on a newspaper) and estimate readers' demand *at the bundle level*.[8] This implies that j in equations (5.1) and (5.2) now refers to a bundle of up to two newspapers. Nests are accordingly designed at the bundle level. Consistent with the nests under the single-homing assumption, we distinguish six nests, comprising respectively: general interest newspaper bundles, sport newspaper bundles, financial newspaper bundles, general interest/sport bundles, general interest/financial bundles, and sport/financial bundles. Figure 5.3 in Appendix 5.8.3 shows the structure of the nests under this double-homing assumption.

Estimating the nested logit discrete choice model at the bundle level relaxes the strong assumption of all newspapers being substitutes. While a discrete choice model at the bundle level still assumes that newspaper bundles are substitutes, individual newspapers can be complements rather than substitutes.[9]

The own-price elasticity of demand η_{jj} in the nested logit model is given by (for $\alpha > 0$):

$$\eta_{jj} = \frac{\partial s_{jt}}{\partial p_{jt}} \frac{p_{jt}}{s_{jt}} = -\frac{\alpha}{1-\sigma} p_{jt}[1 - (1-\sigma)s_{jt} - \sigma s_{jt|g}] \tag{5.3}$$

The cross-price elasticities of demand η_{jk} are instead given by (for $\alpha > 0$):

$$\eta_{jk} = \frac{\partial s_{jt}}{\partial p_{kt}} \frac{p_{kt}}{s_{jt}} = \begin{cases} \alpha p_{kt}[s_{kt} + \frac{\sigma}{1-\sigma} s_{kt|g}] & \text{if } k \neq j \text{ and } k \in g \\ \alpha p_{kt} s_{kt} & \text{if } k \neq j \text{ and } k \notin g \end{cases} \tag{5.4}$$

Note that, when readers' demand is estimated at the bundle level, this implies that the above elasticity formulas give the own-price and cross-price elasticities at the bundle rather than at the newspaper level. However, we are ultimately interested in the elasticities at the newspaper level. Hence, when computing elasticities, we first compute the marginal effects at the bundle level, sum up all the relevant marginal effects at the bundle level to get to the marginal effects at the newspaper level and then multiply these marginal effects with the respective newspaper price (own or cross) and divide by the newspaper's market share to obtain the elasticities at the newspaper level.

[8] Thus, we construct a second dataset based on the data on multi-homing behavior of readers, in which the level of observation is no longer a newspaper but a bundle of up to two newspapers for a given weekday and month. Appendix 5.8.2 contains a detailed description of how we construct the dataset at the bundle level based on the newspaper level data and on the survey information on multi-homing by readers.

[9] In particular, as the utility is parametrized at the bundle level, the Γ_{AB} in Gentzkow (2007), which determines whether goods A and B are substitutes or complements, can be negative or positive. However, we do not estimate the Γs explicitly.

5.4.2 Advertisers' Demand

On the advertisers' side of the market, demand derives from advertisers' choice to allocate a given budget, which changes with the business cycle, across different media outlets. This approach follows Hausman, Leonard, and Zona (1994), who model consumer choices among different brands of beer. Product differentiation is interpreted in the spatial sense proposed by Pinkse, Slade, and Brett (2002), as applied parametrically in Slade (2004) and in Pinkse and Slade (2004). The basic idea is that cross-price elasticities among two products (in our case advertising slots in two different newspapers) should be a function of the distance among the products in characteristic space. One would then expect higher cross-price elasticities when products are closer to each other in terms of characteristics. In addition, own-price effects are allowed to depend on characteristics of the newspapers. In our case, the distance metrics are derived from differences among newspapers in the demographic characteristics of readers (e.g. age, gender, income) and from the amount of overlapping readers between the two newspapers, while the own-price elasticities are allowed to depend on the amount of captive readers. As required by two-sidedness of the media market, we allow advertising demand on a newspaper to depend on its circulation. In the demand specification, the circulation of the newspapers is treated as product advertising is treated in Rojas (2008) and Rojas and Peterson (2008), by allowing own-circulation effects to depend on product characteristics and cross-circulation effects to depend on the distance between newspapers in characteristic space.

Following Hausman, Leonard, and Zona (1994), advertising demand is estimated using demand equations at different levels. At the top level, advertisers decide a budget to spend on advertising in national print newspapers (relative to advertising via other channels). The estimation equation of overall demand for advertising space on national newspapers is the following:

$$\ln(q_{td}) = \beta_0 + \beta_1 \ln(y_t) + \beta_2 \ln(P_{td}) + \phi Z_{td} + \varepsilon_{td}, \tag{5.5}$$

where q_{td} is total advertising quantity (measured in advertising slots) at time t in weekday d, y_t is GDP at time t, P_{td} is a deflated price index for advertising in newspapers at time t in weekday d (see explanation on the price index below), Z_{td} is a set of time and seasonal controls, potentially different by weekday, and ε_{td} is an i.i.d. error term varying across time and weekday. The coefficient β_2 on the price index for advertising in newspapers is hence the overall price elasticity of advertising demand in these seven newspapers relative to other media outlets.

The endogenous advertising price index P_{td} is instrumented by printing cost

shifters, in particular the average real paper cost per page (averaged across the seven newspapers), an electricity price index for industrial consumers,[10] and an hourly wage index for the printed media sector.

At the newspaper level, advertisers decide in a second step on how to allocate their newspaper advertising budget chosen at the top level across the seven newspapers. Thus, advertisers are allowed to multi-home and can potentially decide to buy (different amounts of) advertising space in all of the seven newspapers at the same time.

Following Rojas (2008) and Rojas and Peterson (2008), we use a linear approximation of the Almost Ideal Demand System (AIDS) by Deaton and Muellbauer (1980) to model newspaper level advertising demand. The estimation equation of demand for advertising space in a particular newspaper is then:

$$w_{jtd} = f_{jtd} + \sum_{k=1}^{J} b_{jk} \ln(p_{ktd}) + \sum_{k=1}^{J} c_{jk} \ln(circ_{ktd}) + d_j \ln(\frac{x_{td}}{P_{td}}) + \varepsilon_{jtd}, \tag{5.6}$$

where w_{jtd} is the advertising sales share of newspaper j at time t in weekday d (i.e. $w_{jtd} = \frac{p_{jtd}q_{jtd}}{x_{td}}$), p_{ktd} is the real price per slot of advertising in newspaper k at time t in weekday d, $circ_{ktd}$ is the circulation of newspaper k at time t in weekday d, x_{td} is total advertising expenditures at time t in weekday d (i.e. $x_{td} = \sum_{j=1}^{J} p_{jtd}q_{jtd}$), and P_{td} is an overall advertising price index. Rojas and Peterson (2008) use a Laspeyres index for the overall advertising price index defined as $\ln(P_{td}^L) = \sum_{j=1}^{J} w_j^0 \ln(p_{jtd})$, with w_j^0 being the base share of newspaper j, defined as $w_j^0 = \frac{1}{7T} \sum_{t=1}^{T} \sum_{d=1}^{7} w_{jtd}$. However, in our dataset, some of the newspapers are not available on all weekdays in the earlier years of the data. Using the base share w_j^0 of newspapers to compute the overall advertising price index would, thus, artificially understate the price index for those weekday observations when not all seven newspapers are available. Therefore, we use the Stone price index to measure the overall advertising price instead, defined as $\ln(P_{td}) = \sum_{j=1}^{J} w_{jtd} \ln(p_{jtd})$, as has been proposed by Deaton and Muellbauer (1980) and often applied in practice. The term f_{jtd} can incorporate time and newspaper dummy variables, newspaper characteristics and market specific variables such as demographics. The term $\ln(\frac{x_{td}}{P_{td}})$ enters the estimation equation as the advertising sales shares are conditional on the total advertising expenditures x_{td} set at the top level. It is interacted with newspaper dummy variables d_j, as the effect of a change in total advertising expenditures x_{td} can affect the sales shares w_{jtd} differently for each newspapers: for some newspapers, an increase in total advertising expenditures will increase their advertising sales more than proportionally, for other newspapers,

[10] series nrg_pc_205_h, consumption band Ie from Eurostat.

the increase can be less than proportional. However, in sum these effects must add up to zero, i.e. $\sum_{j=1}^{J} d_j = 0$ must hold.

Equation (5.6) is a first-order approximation in prices and circulation to a demand function allowing for unrestricted price and circulation parameters. In order to reduce the number of cross-price and cross-circulation coefficients that need to be estimated, following Slade (2004) and Pinkse and Slade (2004), we model the cross-price and cross-circulation coefficients b_{jk} and c_{jk} as linear functions of distance measures between newspapers j and k. In particular, how close substitutes two newspapers are in the eyes of advertisers depends on how close these two newspapers are in characteristics space. The closeness metrics are derived from differences among newspapers in the demographic characteristics of readers (age, gender, income) and the amount of overlapping readers between the two newspapers.

The estimation equation thus becomes:

$$\begin{aligned} w_{jtd} = f_{jtd} + b_{jj}\ln(p_{jtd}) + c_{jj}\ln(circ_{jtd}) + \sum_{k\neq j}^{J} g(\delta_{jk})\ln(p_{ktd}) \\ + \sum_{k\neq j}^{J} h(\mu_{jk})\ln(circ_{ktd}) + d_j \ln(\frac{x_{td}}{P_{td}}) + \varepsilon_{jtd}, \end{aligned} \tag{5.7}$$

with

$$b_{jk} = g(\delta_{jk}) = \sum_{l=1}^{L} \lambda_l \delta_{jk}^l \tag{5.8}$$

$$c_{jk} = h(\mu_{jk}) = \sum_{m=1}^{M} \tau_m \mu_{jk}^m \tag{5.9}$$

Where δ_{jk} and μ_{jk} are the L and M closeness measures that determine the cross-price and cross-circulation effects respectively and λ_l and τ_m are parameters to be estimated.

Substituting (5.8) and (5.9) into (5.7) and regrouping terms, the estimation equation is given by:

$$\begin{aligned} w_{jtd} = f_{jtd} + b_{jj}\ln(p_{jtd}) + c_{jj}\ln(circ_{jtd}) + \sum_{l=1}^{L}\lambda_l \sum_{k\neq j}^{J} \delta_{jk}^l \ln(p_{ktd}) \\ + \sum_{m=1}^{M} \tau_m \sum_{k\neq j}^{J} \mu_{jk}^m \ln(circ_{ktd}) + d_j \ln(\frac{x_{td}}{P_{td}}) + \varepsilon_{jtd} \end{aligned} \tag{5.10}$$

The closeness measures between newspapers for continuous product characteristics use an inverse measure of Euclidean distance.[11] These closeness measures

[11] In particular, the closeness between newspapers j and k in terms of product characteristic x is defined as $\frac{1}{1+2\sqrt{(x_j - x_k)^2}}$.

between two newspapers vary between zero and one, so that a one implies that the two newspapers are at the same location in characteristic space. For discrete closeness measures (for example the type of newspaper: generalist, sport, financial), the closeness between two newspapers is equal to one if they belong to the same group (i.e. are of the same type) and zero otherwise. The cross-price and cross-circulation coefficients b_{jk} and c_{jk} are then recovered by replacing the estimated coefficients λ_l and τ_m and the closeness measures δ_{jk} and μ_{jk} into (5.8) and (5.9), respectively.

Note that also the own-price and own-circulation coefficients b_{jj} and c_{jj} can be modelled as functions of newspaper j's own product characteristics. For example, using the percentage of female readers as the relevant product characteristic, the own-price coefficient in (5.8) would be defined as $b_{jj} = b_1 + b_2 femalereaders_j$, where $femalereaders_j$ is the percentage of female readers of newspaper j.

Since we specifically aim to investigate the effect of reader multi-homing on the estimation of demand parameters on both sides of the market, we model the own-price effects as a function of the percentage of captive readers a newspaper has and the cross-price effects as a function of the overlap in readers between two newspapers. Similar to the estimation of readers' demand, we estimate two specifications, one disregarding multi-homing information, such as the percentage of captive readers or the percentage of overlapping readers, and one using this information. The objective is, as in the estimation of readers' demand, to compare estimates of own- and cross-price effects and own- and cross-circulation effects when information on multi-homing is either disregarded or considered. Note that we treat the number of captive and joint readers as exogenous when estimating advertising demand. The reason for this is that, first, the survey information on reader demographics, including on multi-homing behavior, is collected bi-annually. To the extent that they do not carry out additional surveys themselves, advertisers can be expected to base their advertising decisions on this demographic information for a period of six month. Thus, in the estimation, this information is predetermined. Secondly, we do not estimate random effects on the reader side of the market. Hence, in a potential simulation exercise, there will be no changes in the composition of readers following price increases.

Advertising prices as well as newspaper circulation are endogenous and must be instrumented for in the estimation. Following Slade (2004), we use product characteristics of competing newspapers as instrument for price (i.e. BLP instruments). In particular, we use the sum across competitors of the percentage of female readers as well as the mean age of readers to instrument for own advertising price. In addition, we use the cost shifter real paper cost per page to instrument own advertising price. Newspaper circulation is instrumented with the real paper cost per issue as well

as the dummy for the day of the week when a magazine of general information is bundled to the newspaper. Increases in the printing costs should increase the newspaper price and, hence, decrease reader demand, while adding a magazine to the newspaper increases reader demand (see estimation results on readers' demand in Section 5.5). Using the dummy for a magazine of general information as instrument for circulation relies on the assumption that the presence of the magazine does not directly influence the demand for advertising space *on the newspaper* (other than through its effect on newspaper circulation). Lastly, as total advertising expenditures x_{td} are constructed from prices and quantity variables, they are also treated as endogenous and instrumented with GDP.

The price elasticities of advertising demand $\tilde{\eta}_{jk}$ *conditional on total advertising expenditures* x_{td} are given by:

$$\tilde{\eta}_{jk} = \begin{cases} -1 + \frac{b_{jj}}{w_{jt}} - d_j & \text{if } k = j \\ \frac{b_{jk}}{w_{jt}} - d_j \frac{w_{kt}}{w_{jt}} & \text{if } k \neq j \end{cases} \tag{5.11}$$

with b_{jj} potentially being a function of own product characteristics and b_{jk} being a function of the closeness measures.

Unconditional advertising price elasticities also need to take into account how advertising price increases by one newspaper change the overall price index for advertising in newspapers, how this in turn then changes the overall demand for advertising in print newspapers at the top level (relative to other media outlets), and how this change in total advertising expenditures x_{td} then affects the advertising quantity demanded on newspaper j. The *unconditional* price elasticities for advertising demand η_{jk} take these effects into account and are given by:

$$\eta_{jk} = \begin{cases} \tilde{\eta}_{jj} + (1 + \frac{d_j}{w_{jt}})(\beta_2 + 1)w_{jt} & \text{if } k = j \\ \tilde{\eta}_{jk} + (1 + \frac{d_j}{w_{jt}})(\beta_2 + 1)w_{kt} & \text{if } k \neq j \end{cases} \tag{5.12}$$

where β_2 is the overall price elasticity of advertising demand in the seven newspapers relative to other media outlets estimated at the top level (see equation (5.5)).

The circulation elasticities of advertising demand ρ_{jk}, which do not depend on total advertising expenditures x_{td}, are given by:

$$\rho_{jk} = \begin{cases} \frac{c_{jj}}{w_{jt}} & \text{if } k = j \\ \frac{c_{jk}}{w_{jt}} & \text{if } k \neq j \end{cases} \tag{5.13}$$

with c_{jj} potentially being a function of own product characteristics and c_{jk} being a function of the closeness measures. Appendix 5.8.4 contains the derivation of both the conditional and the unconditional price as well as the circulation elasticities.

5.5 Estimation Results

5.5.1 Estimation Results Readers' Demand

We present here the results on the estimation of readers' demand, once under the assumption of single-homing readers (nested logit at newspaper level) and once under the assumption of double-homing readers (nested logit at bundle level).

Table 5.4 shows results for readers' demand assuming single-homing by readers. Both specifications contain nest-specific (i.e. newspaper type or bundle type) year fixed effects to account for the time trend as well as month fixed effects to account for seasonality. While the first specification contains month fixed effects that are common across nests, the second specification allows seasonality to differ across nests by including nest-specific month fixed effects.

The price coefficient is negative and statistically significant at the 1% significance level in both specifications and varies between -0.86 and -1.32 depending on how we model seasonality. The advertising intensity coefficient is positive and significant. This result may seem counter-intuitive at first as Argentesi and Filistrucchi (2007) find insignificant effects of advertising quantity on reader demand with similar data. However, the positive effect of advertising intensity on reader demand is very small: increasing the advertising quantity by one slot increases circulation on average by between 15 and 33 readers depending on the newspaper. Hence, it may be that those who are not interested in ads in the newspaper can easily skip them while those who are interested can enjoy them, such that overall demand is affected little but positively.

The estimated σ is positive and statistically significant at the 1% significance level, thus showing that the chosen nesting structure matters.

The other coefficients have the expected signs and are mostly consistent with Argentesi and Filistrucchi (2007): both the coefficients for the dummy variables for the day of issue of a magazine of general information or the day of issue of a women magazine are positive and statistically significant at the 1% level. The coefficient for the dummy variable for the day of issue of an economic insert has a negative but mostly statistically insignificant effect. The number of local pages in a newspaper also impacts readers' demand negatively (but insignificantly) in the two specifications. For the website dummy variables interacted with internet penetration rate, which account for the launch of a website by a given newspaper, the coefficients are often statistically insignificant but mostly negative in the cases where they are statistically significant. Thus, introducing websites seems to negatively impact demand for printed newspapers, as in Filistrucchi (2005). We also include a set of editor

Table 5.4: Readers' Demand - Single-Homing

VARIABLES	(1)	(2)
Real cover price	-0.862***	-1.318***
	(-2.882)	(-4.291)
σ	0.231***	0.169***
	(3.460)	(2.736)
Advertising intensity (slots/pages)	0.003***	0.002**
	(6.558)	(2.492)
Generalist magazine (day)	0.218***	0.325***
	(3.503)	(5.012)
Women magazine (day)	0.300***	0.462***
	(2.918)	(4.369)
Economic insert (day)	-0.025	-0.052***
	(-1.402)	(-3.218)
Local pages	-0.005	-0.003
	(-1.473)	(-0.745)
Observations	8,795	8,795
R-squared	0.292	0.420
Number of Newspaper/Weekday FE	49	49
Website opening	YES	YES
Director dummy variables	YES	YES
Newspaper/Weekday Fixed Effects	YES	YES
Time trend	Year/Nest Fixed Effects	Year/Nest Fixed Effects
Seasonality	Month Fixed Effects	Month/Nest Fixed Effects
Adjusted R-squared	0.280	0.408
Kleibergen Paap stat.	59.94	50.46
p-value KP	0	3.81e-09
Chi-squared quadratic web	14.18	6.277
p-value web	0.048	0.508

Robust z-statistics in parentheses, *** p<0.01, ** p<0.05, * p<0.1

dummy variables in the estimation. Standard errors are heteroskedasticity robust and corrected for autocorrelation of order 1. The Kleibergen-Paap test statistic for weak instruments indicates that there is no problem of weak instruments in the first stages.

Table 5.5 shows the estimation results for readers' demand accounting for double-homing by readers. The two specifications are analogous to the two specifications in the estimations assuming single-homing by readers. However, the structure of the nests is richer. Since demand is now modelled as demand for bundles of one or two newspapers, in addition to the outside good, there are now not only the nests of generalist, of sport, and of financial newspapers, but also the nests of generalists and sport, of generalist and financial, and of sport and financial bundles. See Figure 5.3 in Appendix 5.8.3 for the structure of the nests.

The estimated price coefficient is negative and significant at the 1% significance level and varies between -1.08 and -1.34 depending on how we account for seasonality.

Table 5.5: Readers' Demand - Double-Homing

VARIABLES	(1)	(2)
Real cover price	-1.079***	-1.339***
	(-9.555)	(-11.342)
σ	0.730***	0.745***
	(12.785)	(12.877)
Advertising intensity (slots/pages)	0.001***	-0.001***
	(3.634)	(-2.765)
Generalist magazine (day)	0.241***	0.304***
	(9.989)	(11.910)
Generalist magazine plus (day)	0.203***	0.248***
	(4.846)	(5.344)
Women magazine (day)	0.371***	0.461***
	(9.952)	(11.879)
Women magazine plus (day)	0.337***	0.416***
	(10.088)	(11.838)
Economic insert (day)	-0.003	-0.016*
	(-0.331)	(-1.847)
Economic insert plus (day)	-0.035***	-0.032***
	(-2.871)	(-2.674)
Local pages	-0.002*	-0.003*
	(-1.859)	(-1.878)
Observations	35,105	35,105
R-squared	0.727	0.703
Number of Bundle/Weekday FE	196	196
Website opening	YES	YES
Director dummy variables	YES	YES
Bundle/Weekday Fixed Effects	YES	YES
Time trend	Year/Nest Fixed Effects	Year/Nest Fixed Effects
Seasonality	Month Fixed Effects	Month/Nest Fixed Effects
Adjusted R-squared	0.724	0.700
Kleibergen Paap stat.	298.2	288.5
p-value KP	0	0
Chi-squared quadratic web	50.38	82
p-value web	1.21e-08	0

Robust z-statistics in parentheses, *** p<0.01, ** p<0.05, * p<0.1

The magnitude of the price coefficient is slightly larger than in the estimations under the assumption of single-homing by readers. The advertising intensity coefficient is now negative and significant in specification (2), where we allow seasonality to differ across nests. In any case, the impact of advertising on reader demand is small. The estimated σ is again positive and significant at the 1% significance level and much larger than the nesting parameter in the estimations under the assumption of single-homing by readers.

The other coefficients have the expected signs: both the coefficients for the dummy variables for the day of issue of a magazine of general information or the day of issue of a women magazine are positive and statistically significant in both specifications. Again, the coefficient of the dummy variable for the issue of an economic insert is negative and statistically significant in the last specification. These dummy variables for the day of issue of a specific magazine or insert now measure whether there is at least one of these magazines/inserts in the bundle on a particular day. Those variables marked "plus" measure in addition whether there was a magazine/women magazine/economic insert issued on the same day in both of the newspapers in the respective bundle. The effect of a second magazine or women's magazine in the bundle is positive and significant while the effect of a second economic insert is negative and significant. The number of local pages in the bundle again has a negative impact on readers' demand for the bundle. Again, the effect of the launch of a website is mostly negative, suggesting that the print version and the online version of a newspaper are substitutes from the readers' perspective. The Kleibergen-Paap test statistic for weak instruments indicates that there is no problem of weak instruments in the first stages.

Table 5.6 shows the resulting mean own- and cross-price elasticities for the seven newspapers based on the estimation results of specification (2) in Table 5.4, i.e. under the assumption of single-homing readers.[12] While the mean own-price elasticity varies between -1.27 and -1.48, the cross-price elasticities are small and vary between 0.008 and 0.15.

While the dataset allowing for double-homing is at the bundle level, the price and network effect elasticities should still be at the product, i.e. newspaper, level. Thus, we sum over all relevant marginal effects at the bundle level to obtain the marginal effects at the newspaper level when we account for double-homing by readers.[13] We

[12]We prefer specification (2) as it allows seasonality to differ across nests.

[13]For example, for the own-price effect of newspaper A, we take into account the effects of a price increase of all bundles containing newspaper A on all bundles containing newspaper A. For, e.g., the cross-price effect on newspaper A of a price increase of newspaper B, we take into account the effect of a price increase of all bundles containing newspaper B on all bundles containing newspaper A. Note that this also includes the own-price effect of bundle AB increasing its price, which is negative.

Table 5.6: Mean Own- and Cross-Price Elasticities - Readers' Demand - Single-Homing

	Corriere	Corriere Sport	Gazzetta Sport	Giornale	Repubblica	Il Sole	Stampa
Corriere	-1.479	0.010	0.014	0.040	0.107	0.012	0.067
	(0.229)	(0.002)	(0.003)	(0.006)	(0.022)	(0.002)	(0.015)
Corriere Sport	0.023	-1.313	0.152	0.008	0.021	0.012	0.013
	(0.005)	(0.079)	(0.010)	(0.001)	(0.005)	(0.002)	(0.003)
Gazzetta Sport	0.023	0.112	-1.268	0.008	0.021	0.012	0.013
	(0.005)	(0.009)	(0.070)	(0.001)	(0.005)	(0.002)	(0.003)
Giornale	0.117	0.010	0.014	-1.461	0.107	0.012	0.067
	(0.023)	(0.002)	(0.003)	(0.088)	(0.022)	(0.002)	(0.015)
Repubblica	0.117	0.010	0.014	0.040	-1.478	0.012	0.067
	(0.022)	(0.002)	(0.003)	(0.006)	(0.230)	(0.002)	(0.015)
Il Sole	0.023	0.010	0.014	0.008	0.021	-1.323	0.013
	(0.005)	(0.002)	(0.003)	(0.001)	(0.005)	(0.084)	(0.003)
Stampa	0.117	0.010	0.014	0.040	0.107	0.012	-1.444
	(0.023)	(0.002)	(0.003)	(0.006)	(0.022)	(0.002)	(0.207)

Mean elasticities over the years 1992-2006. Standard deviations are reported in parentheses.

Table 5.7: Mean Own- and Cross-Price Elasticities - Readers' Demand - Double-Homing

	Corriere	Corriere Sport	Gazzetta Sport	Giornale	Repubblica	Il Sole	Stampa
Corriere	-3.621	-0.004	-0.144	0.144	0.710	-0.145	0.686
	(0.574)	(0.021)	(0.032)	(0.062)	(0.144)	(0.019)	(0.155)
Corriere Sport	-0.013	-3.043	1.246	-0.065	-0.240	-0.064	-0.008
	(0.050)	(0.240)	(0.126)	(0.049)	(0.084)	(0.018)	(0.032)
Gazzetta Sport	-0.241	0.917	-2.498	-0.032	-0.019	-0.069	-0.067
	(0.057)	(0.096)	(0.170)	(0.049)	(0.028)	(0.021)	(0.020)
Giornale	0.417	-0.082	-0.054	-4.127	0.818	-0.152	0.502
	(0.192)	(0.056)	(0.074)	(0.267)	(0.204)	(0.027)	(0.136)
Repubblica	0.775	-0.112	-0.015	0.313	-3.778	-0.141	0.653
	(0.157)	(0.023)	(0.018)	(0.081)	(0.555)	(0.022)	(0.140)
Il Sole	-0.268	-0.051	-0.076	-0.096	-0.240	-1.288	-0.099
	(0.048)	(0.011)	(0.022)	(0.017)	(0.050)	(0.089)	(0.021)
Stampa	1.197	-0.006	-0.072	0.301	1.047	-0.094	-3.959
	(0.228)	(0.021)	(0.026)	(0.059)	(0.198)	(0.015)	(0.549)

Mean elasticities over the years 1992-2006. Standard deviations are reported in parentheses.

then multiply these marginal effects with the respective price or advertising intensity (own or cross) and divide by the newspaper's market share to obtain the elasticities at the product level.[14]

Table 5.7 shows the resulting mean own- and cross-price elasticities for the seven newspapers based on the estimation results of specification (2) in Table 5.5, i.e. under the assumption of double-homing readers. Firstly, the mean own-price elasticities now vary between -1.29 and -4.13, which is much larger than the estimated mean own-price elasticities based on the assumption of single-homing readers. In particular, demand for the generalist newspapers is more elastic if double-homing is taken into account in the estimation.

Secondly, note that cross-price elasticities can now be positive or negative. In particular, we find that cross-price elasticities between newspapers of the same type (i.e. generalist, sport, financial) are positive, while cross-price elasticities between newspapers of different types are negative. This implies that newspapers of the same newspaper type are substitutes while newspapers of different types are complements.[15] Additionally, the magnitude of the cross-price elasticities is mostly larger than those based on the assumption of single-homing by readers.

Tables 5.8 and 5.9 show the mean own- and cross-network effects elasticities based on the same estimation results under the assumption of single-homing and double-homing readers, respectively.

As the estimated coefficient on the advertising intensity is small but positive in specification (2) in Table 5.4, the estimated own-network effect elasticities are small and positive, while the cross-network effect elasticities are negative but small. The own-network effect elasticities vary between 0.13 and 0.22 while the cross-network effect elasticities vary between -0.001 and -0.021.

As the estimated coefficient on the advertising elasticity is negative in specification (2) in Table 5.5, when we account for double-homing by readers, the estimated own-network effect elasticities are now negative and vary between -0.10 and -0.42. As for the cross-price elasticities, also the cross-network effect elasticities suggest that newspapers of the same type are substitutes while newspapers of different types are complements.

[14] We multiply here with the market share based on the actual circulation of the newspaper contained in the newspaper level data. However, as a robustness check, we also computed the elasticities based on the newspaper market shares that are implied by the estimated newspaper circulation resulting from the procedure to create the bundle-level dataset. Qualitative results did not change when we used these alternative market shares in the computation of elasticities.

[15] We also estimated a nested logit model allowing for different σs for each nest. We still find that newspapers of the same type are substitutes while newspapers of different types are complements.

Table 5.8: Mean Own- and Cross-Network Effect Elasticities - Readers' Demand - Single-Homing

	Corriere	Corriere Sport	Gazzetta Sport	Giornale	Repubblica	Il Sole	Stampa
Corriere	0.216	-0.001	-0.002	-0.004	-0.012	-0.001	-0.007
	(0.066)	(0.001)	(0.001)	(0.001)	(0.004)	(0.001)	(0.002)
Corriere Sport	-0.003	0.151	-0.021	-0.001	-0.002	-0.001	-0.001
	(0.001)	(0.062)	(0.007)	(0.000)	(0.001)	(0.001)	(0.000)
Gazzetta Sport	-0.003	-0.013	0.173	-0.001	-0.002	-0.001	-0.001
	(0.001)	(0.005)	(0.058)	(0.000)	(0.001)	(0.001)	(0.000)
Giornale	-0.017	-0.001	-0.002	0.144	-0.012	-0.001	-0.007
	(0.006)	(0.001)	(0.001)	(0.047)	(0.004)	(0.001)	(0.002)
Repubblica	-0.017	-0.001	-0.002	-0.004	0.167	-0.001	-0.007
	(0.006)	(0.001)	(0.001)	(0.001)	(0.054)	(0.001)	(0.002)
Il Sole	-0.003	-0.001	-0.002	-0.001	-0.002	0.132	-0.001
	(0.001)	(0.001)	(0.001)	(0.000)	(0.001)	(0.060)	(0.000)
Stampa	-0.017	-0.001	-0.002	-0.004	-0.012	-0.001	0.155
	(0.006)	(0.001)	(0.001)	(0.001)	(0.004)	(0.001)	(0.051)

Mean elasticities over the years 1992-2006. Standard deviations are reported in parentheses.

Table 5.9: Mean Own- and Cross-Network Effect Elasticities - Readers' Demand - Double-Homing

	Corriere	Corriere Sport	Gazzetta Sport	Giornale	Repubblica	Il Sole	Stampa
Corriere	-0.4239	-0.0005	-0.0156	0.0117	0.0646	-0.0115	0.0579
	(0.1286)	(0.0019)	(0.0059)	(0.0062)	(0.0216)	(0.0052)	(0.0185)
Corriere Sport	-0.0018	-0.2808	0.1364	-0.0051	-0.0223	-0.0054	-0.0004
	(0.0057)	(0.1138)	(0.0463)	(0.0045)	(0.0110)	(0.0033)	(0.0025)
Gazzetta Sport	-0.0279	0.0850	-0.2734	-0.0027	-0.0016	-0.0056	-0.0059
	(0.0092)	(0.0357)	(0.0907)	(0.0045)	(0.0026)	(0.0034)	(0.0028)
Giornale	0.0506	-0.0076	-0.0059	-0.3281	0.0752	-0.0119	0.0426
	(0.0275)	(0.0068)	(0.0093)	(0.1070)	(0.0287)	(0.0054)	(0.0154)
Repubblica	0.0912	-0.0105	-0.0016	0.0251	-0.3441	-0.0114	0.0555
	(0.0307)	(0.0053)	(0.0021)	(0.0101)	(0.1077)	(0.0054)	(0.0173)
Il Sole	-0.0313	-0.0047	-0.0084	-0.0075	-0.0218	-0.1029	-0.0084
	(0.0093)	(0.0021)	(0.0038)	(0.0026)	(0.0075)	(0.0469)	(0.0028)
Stampa	0.1413	-0.0003	-0.0083	0.0242	0.0953	-0.0077	-0.3426
	(0.0481)	(0.0020)	(0.0046)	(0.0098)	(0.0317)	(0.0041)	(0.1164)

Mean elasticities over the years 1992-2006. Standard deviations are reported in parentheses.

Table 5.10: Advertisers' Demand - Top Level

VARIABLES	(1)
Log(advertising price index)	-1.856***
	(-3.764)
Log(GDP)	12.365***
	(10.266)
Constant	492.864***
	(5.556)
Observations	1,260
Time trend	Weekday specific quadratic yearly trend
Seasonality	Month Fixed Effects
Kleibergen Paap stat.	24.01
p-value KP	2.48e-05

Robust z-statistics in parentheses, *** p<0.01, ** p<0.05, * p<0.1

5.5.2 Estimation Results Advertisers' Demand

Here, we present the results on the estimation of advertisers' demand, first at the top level, where advertisers decide on the budget to spend on advertising in print newspapers and, second, at the newspaper level, where advertisers decide on how to split this advertising budget across the different newspapers. Since we aim to highlight the possible bias due to the omission of information on multi-homing by readers, at the newspaper level, we present results from two specifications: one disregarding information on double-homing by readers; the other making use of the information on captive readers and overlapping readers between newspapers.

Table 5.10 shows results for top level advertisers' demand. The coefficient on the price index is negative and statistically significant at the 1% significance level. GDP, as an indicator for the overall business cycle, has a positive and statistically significant effect on overall advertising demand. We account for the time trend by including a weekday specific quadratic yearly time trend and allow for seasonality by including month fixed effects.

Specification (1) in Table 5.11 shows estimation results for advertisers' demand at the newspaper level when information on single-homing and multi-homing by readers is omitted. The specification contains newspaper and weekday fixed effects as well as a newspaper type-specific (generalist, sport, business) quadratic yearly time trend. Seasonality is accounted for by month fixed effects.

Specification (2) in Table 5.11 shows instead estimation results for advertisers' demand at the newspaper level when information on the percentage of captive readers and the percentage of overlapping readers is used.

The coefficient on advertising price is negative and significant at the 5% signifi-

Table 5.11: Advertisers' Demand - Newspaper Level

VARIABLES	(1) No Info on DH Readers	(2) With Info on DH Readers
Log(real price per ad slot)	-0.040***	-0.536***
	(-3.049)	(-7.883)
Log(real price per ad slot)*Log(captive readers)		0.139***
		(8.023)
Log(real price per ad slot)*Log(female readers)	0.011***	0.185***
	(4.532)	(9.683)
Log(real price per ad slot)*Log(captive readers)		-0.046***
*Log(female readers)		(-9.085)
Log(circulation)	0.307***	0.307***
	(17.622)	(18.126)
Readers' income cross-price measure	0.009***	-0.001
	(4.434)	(-0.457)
Joint readers cross-price measure		0.018***
		(2.780)
Same type cross-circulation measure	-0.005***	-0.005***
	(-9.464)	(-8.485)
Log(xt/Pt)*Corriere	-0.091***	-0.164***
	(-5.802)	(-9.208)
Log(xt/Pt)*Corriere Sport	0.150***	0.071***
	(10.365)	(5.786)
Log(xt/Pt)*Gazzetta Sport	0.147***	0.074***
	(10.214)	(5.767)
Log(xt/Pt)*Giornale	0.000	0.057***
	(0.005)	(3.433)
Log(xt/Pt)*Repubblica	-0.065***	-0.154***
	(-4.182)	(-8.992)
Log(xt/Pt)*Il Sole	0.198***	0.123***
	(11.007)	(6.808)
Log(xt/Pt)*Stampa	0.187***	0.020
	(12.571)	(1.272)
Observations	8,795	8,795
Number of id	7	7
Newspaper Fixed Effects	YES	YES
Weekday Fixed Effects	YES	YES
Time trend	Type specific quadratic yearly trend	Type specific quadratic yearly trend
Seasonality	Month Fixed Effects	Month Fixed Effects
Kleibergen Paap statistic	382.4	343.5
p-value KP	0	0
Chi-squared $\sum_{j=1}^{J} d_j = 0$	38.35	0.107
p-value	0.000	0.743

Robust z-statistics in parentheses, *** p<0.01, ** p<0.05, * p<0.1

cance level in specification (1). The interaction term between the advertising slot price and the percentage of female readers is positive and significant. This implies that advertisers' demand is less elastic for newspapers that offer access to more female readers.

In specification (2), we allow the own advertising price elasticity in addition to depend on the percentage of captive readers of the newspaper. The coefficient on advertising price is negative and significant at the 1% level and much larger in magnitude than in specification (1), when information on multi-homing by readers is ignored in the estimation of advertising demand. The interaction term between the advertising slot price and the percentage of female readers is not just still positive and significant, but also larger in magnitude than in specification (1). Importantly, the interaction term between the advertising slot price and the percentage of captive readers is also positive and significant, while the interaction term between price, the percentage of captive readers, and the percentage of female readers is negative and significant. This implies that advertisers' demand is less elastic for newspapers that offer exclusive access to more female and more captive readers, particularly if the latter are men.

The coefficient on the own circulation is positive and significant at the 1% significance level in both specifications, highlighting that advertisers value a newspaper more the more readers this newspaper reaches (all else equal).

Cross-price effects are modelled as a function of the closeness in the income of newspaper readers and, in specification (2) when we take into account multi-homing reader information, also as a function of the overlap in readers between two newspapers. The reader income based cross-price measure is positive and significant in specification (1), implying that higher prices of competing newspapers increase own advertising demand and that newspapers are closer substitutes for advertisers if they reach readers that are similar/close in terms of socio-economic status/income. The coefficient on the joint readers cross-price measure in specification (2) is positive and significant, showing that newspapers are closer substitutes for advertisers if they have a higher share of readers in common; i.e. if these readers can be reached on either of the two newspapers. The coefficient is also much larger in magnitude than the one on the reader income based cross-price measure in specification (1), thus the overlap in readers seems to be a more important determinant of substitutability of newspapers for advertisers than readers' income. In particular, the reader income based cross-price measure is no longer significant in specification (2) where we allow cross-price effects to depend on the share of common readers.

Cross-circulation effects are modelled as a function of the discrete closeness measure of newspaper type in both specifications. The negative and statistically sig-

nificant coefficient on the same type cross-circulation measure shows that higher circulation of competing newspapers of the same newspaper type decreases own advertising demand.

Lastly, not all $\ln(\frac{x_{td}}{P_{td}})$ terms are statistically significant, indicating that, for some newspapers, an increase in the overall print advertising budget does not affect their advertising sales share, while other newspapers gain or lose sales share with increasing overall print newspaper advertising expenditures. For the computation of price elasticities, we set the statistically insignificant $\ln(\frac{x_{td}}{P_{td}})$ terms to zero. As discussed in Section 5.4, the theoretical model implies that $\sum_{j=1}^{J} d_j = 0$ must hold. We do not impose this constraint *ex ante*, but estimate the unrestricted model and test whether the constraint holds *ex post*. The H0 hypothesis of all d_j adding up to zero is clearly rejected in specification (1), in which information on single-homing and multi-homing readers is omitted. However, in specification (2), when multi-homing reader information is taken into account, H0 cannot be rejected. This implies that advertiser demand in specification (1) might be misspecified and that taking into account information on single-homing and multi-homing readers is crucial for correctly estimating advertising demand.[16] Therefore, we present price elasticities as well as network effect elasticities based only on the specification that considers information on single-homing and double-homing readers.

Table 5.12 shows the resulting mean *conditional* own- and cross-price elasticities for advertising demand of the seven newspapers based on specification (2) in Table 5.11, i.e. using the information on single-homing and multi-homing readers in the demand estimation. Note that the conditional cross-price elasticities can be positive or negative depending on the sign (and significance) of the estimated d_j and the magnitude of the $\frac{w_{kt}}{w_{jt}}$ term, which capture the effect of a price change on P_{td} (see equation (5.11)). The mean conditional own-price elasticity varies between -0.33 and -0.92. The conditional cross-price elasticities vary between -0.6 and 0.19.

Looking back at the formula for the conditional own-price elasticity in equation (5.11), note firstly that the b_{jj} will increase, the higher the percentage of female readers and captive readers of newspaper j. This implies that newspapers with a high percentage of captive and female readers should have a smaller own-price elasticity, *ceteris paribus*. Secondly, note that a positive (negative) d_j implies a higher (lower) conditional own-price elasticity (in absolute value) via its effect on the price index P_{td}, which in turn has an effect on the expenditure share w_{jtd} (income effect). This implies that a newspaper for which the estimated d_j is positive should have a larger conditional own-price elasticity than a newspaper with a negative estimate

[16] We also tried using different time trends in specification (1), however the restriction was never satisfied when information on single-homing and double-homing readers was ignored.

Table 5.12: Mean Conditional Own- and Cross-Price Elasticities - Advertisers' Demand - Including DH Readers

	Corriere	Corriere Sport	Gazzetta Sport	Giornale	Repubblica	Il Sole	Stampa
Corriere	-0.756	0.027	0.048	0.039	0.173	0.114	0.092
	(0.018)	(0.014)	(0.018)	(0.016)	(0.059)	(0.066)	(0.034)
Corriere Sport	-0.601	-0.712	0.057	-0.108	-0.563	-0.353	-0.307
	(0.382)	(0.196)	(0.079)	(0.058)	(0.383)	(0.244)	(0.174)
Gazzetta Sport	-0.345	0.006	-0.795	-0.066	-0.353	-0.225	-0.184
	(0.156)	(0.026)	(0.104)	(0.032)	(0.181)	(0.149)	(0.081)
Giornale	-0.246	-0.023	-0.020	-0.325	-0.253	-0.133	-0.131
	(0.124)	(0.027)	(0.039)	(0.352)	(0.113)	(0.075)	(0.054)
Repubblica	0.186	0.030	0.047	0.038	-0.776	0.110	0.094
	(0.085)	(0.027)	(0.045)	(0.018)	(0.023)	(0.053)	(0.040)
Il Sole	-0.250	-0.031	-0.047	-0.042	-0.215	-0.839	-0.122
	(0.192)	(0.036)	(0.055)	(0.035)	(0.148)	(0.233)	(0.091)
Stampa	0.010	0.004	0.013	0.005	0.012	0.010	-0.915
	(0.003)	(0.001)	(0.003)	(0.001)	(0.004)	(0.003)	(0.024)

Mean elasticities over the years 1992-2006. Standard deviations are reported in parentheses.

of d_j, *ceteris paribus*. Lastly, the conditional own-price elasticity also depends on the advertising sales share w_{jtd}. However, whether a high advertising sales share increases or decreases the conditional own-price elasticity depends on the sign of the marginal effect of price on the sales share w_{jtd}, i.e. the sign of the b_{jj}. The average marginal effect of own price on the expenditure share w_{jtd} is positive for all newspapers based on the results of specification (2), i.e. the quantity decreases by less than the price increases, implying inelastic demand. Secondly, given that the b_{jj} is positive, a high advertising sales share w_{jtd} actually increases the own-price elasticity (in absolute value), implying more elastic demand. The interaction between all of these effects determines the conditional own-price elasticity.

Il Sole 24 Ore and La Stampa have the highest conditional own-price elasticities, while Il Giornale has the lowest conditional own-price elasticity (in absolute value) in Table 5.12. The relatively high own-price elasticity of Il Sole 24 Ore can be explained by its relatively low share of captive readers as well as the positive estimate of the d_j term that increases the conditional own-price elasticity. The relatively low own-price elasticity of Il Giornale on the other hand, can be explained by the relatively large positive average marginal effect driven by the high percentage of female readers, implying less elastic demand. This effect is then further magnified by the low advertising sales share, decreasing the conditional own-price elasticity (in absolute value). Lastly, note the relatively high own-price elasticities of the

two sport newspapers, Corriere dello Sport and Gazzetta dello Sport. These two newspapers have relatively many single-homing readers as well as low advertising sales shares (4% and 6% compared to a mean across newspapers of 14%), which decreases the conditional own-price elasticity. However this effect is compensated by their extremely low percentage of female readers (on average 13% for both newspapers compared to a mean across newspapers of 32%) as well as the positive estimates of the d_j terms that increase the conditional own-price elasticities.[17]

As explained above, conditional cross-price elasticities can be positive or negative depending on the sign of the estimated d_j and on the magnitude of the $\frac{w_{kt}}{w_{jt}}$ term, which captures the effect of a price change on P_{td}. Looking back at the formula for the conditional cross-price elasticity (see equation (5.11)), the cross-price elasticity can become negative only if d_j is positive. If newspaper k increases its price, this increases demand for newspaper j via the substitution effect. However, if d_j is positive, the increase in the price of newspaper k also leads to an increase in the price index P_{td}, which decreases demand for newspaper j (income effect). If the conditional cross-price elasticity is negative, this second effect dominates the positive substitution effect. This is the case for most of the conditional cross-price elasticities of newspapers where the estimated d_j is positive in specification (2).

The cross-price elasticities between two newspapers will be higher the more joint readers two newspapers have, *ceteris paribus*. Note in particular the cross-price elasticities between the two sport newspapers, Corriere dello Sport and Gazzetta dello Sport, in Table 5.12. Even though cross-price elasticities towards all other newspapers are actually negative as the estimated d_js for both sport newspapers are positive, the cross-price elasticities between each other are positive. This is the case because these two newspapers have a high share of readers in common (see Table 5.1), implying that the positive substitution effect dominates the negative income effect via the price index P_{td}.

Unconditional advertising price elasticities also need to take into account how advertising price increases by one newspaper change the overall price index for advertising in newspapers, how this in turn then changes the overall demand for advertising in print newspapers at the top level (relative to other media outlets), and how this change in total advertising expenditures x_{td} then affects the advertising quantity demanded on newspaper j. The unconditional price-elasticities for advertising demand take these effects into account and are given in Table 5.13 based on the estimation results of specification (2). Given that a price increase by one newspaper will result in a price increase in the price index P_{td}, which subsequently leads to a decrease in the overall budget spend on advertising in national print newspapers

[17]See Appendix 5.8.5 for the relevant newspaper level summary statistics.

Table 5.13: Mean Unconditional Own- and Cross-Price Elasticities - Advertisers' Demand - Including DH Readers

	Corriere	Corriere Sport	Gazzetta Sport	Giornale	Repubblica	Il Sole	Stampa
Corriere	-0.856	0.014	0.027	0.021	0.089	0.066	0.046
	(0.072)	(0.013)	(0.019)	(0.019)	(0.083)	(0.072)	(0.043)
Corriere Sport	-1.397	-0.806	-0.099	-0.264	-1.323	-0.817	-0.708
	(0.756)	(0.210)	(0.096)	(0.121)	(0.769)	(0.500)	(0.344)
Gazzetta Sport	-0.920	-0.068	-0.911	-0.182	-0.904	-0.571	-0.478
	(0.330)	(0.048)	(0.126)	(0.072)	(0.381)	(0.328)	(0.174)
Giornale	-0.761	-0.093	-0.133	-0.423	-0.729	-0.423	-0.388
	(0.295)	(0.069)	(0.099)	(0.363)	(0.256)	(0.190)	(0.123)
Repubblica	0.098	0.019	0.030	0.021	-0.867	0.058	0.049
	(0.130)	(0.040)	(0.070)	(0.024)	(0.062)	(0.058)	(0.056)
Il Sole	-0.730	-0.098	-0.154	-0.135	-0.646	-1.083	-0.360
	(0.414)	(0.085)	(0.130)	(0.077)	(0.315)	(0.275)	(0.190)
Stampa	-0.231	-0.030	-0.040	-0.044	-0.210	-0.130	-1.037
	(0.056)	(0.018)	(0.025)	(0.016)	(0.046)	(0.056)	(0.052)

Mean elasticities over the years 1992-2006. Standard deviations are reported in parentheses.

at the top-level, we expect own-price elasticities to increase (in absolute value) and cross-price elasticities to decrease. This is reflected in the unconditional price elasticities in Table 5.13. All own-price elasticities increase in absolute value and now vary between -0.42 and -1.08 in Table 5.13. All cross-price elasticities decrease to the point where some positive conditional cross-price elasticities turn to negative unconditional cross-price elasticities.

Lastly, Table 5.14 shows the own- and cross-network effect elasticities based on the newspaper level advertising demand estimation results of specification (2). As the cross-circulation effects are modelled as a function of the discrete closeness measure of newspaper type, cross-circulation effects are zero by construction between newspapers of different newspaper types. Own-network effect elasticities vary between 1.16 and 9.98, while the cross-network effect elasticities are small and vary between -0.16 and -0.02.

Note that the own-network effect elasticities are given by the coefficient on the own circulation divided by the advertising sales share w_{jt} of newspaper j (see equation (5.13)). This implies that newspapers with a low advertising sales share will have larger own-network effect elasticities, which is the case for Corriere dello Sport, Gazzetta dello Sport, and Il Giornale.

Table 5.14: Mean Own- and Cross-Circulation Elasticities - Advertisers' Demand - Including DH Readers

	Corriere	Corriere Sport	Gazzetta Sport	Giornale	Repubblica	Il Sole	Stampa
Corriere	1.156	0.000	0.000	-0.018	-0.018	0.000	-0.018
	(0.282)			(0.004)	(0.004)		(0.004)
Corriere Sport	0.000	9.975	-0.156	0.000	0.000	0.000	0.000
		(5.288)	(0.083)				
Gazzetta Sport	0.000	-0.092	5.857	0.000	0.000	0.000	0.000
		(0.035)	(2.267)				
Giornale	-0.094	0.000	0.000	5.983	-0.094	0.000	-0.094
	(0.030)			(1.940)	(0.030)		(0.030)
Repubblica	-0.019	0.000	0.000	-0.019	1.245	0.000	-0.019
	(0.006)			(0.006)	(0.401)		(0.006)
Il Sole	0.000	0.000	0.000	0.000	0.000	2.325	0.000
						(1.371)	
Stampa	-0.036	0.000	0.000	-0.036	-0.036	0.000	2.290
	(0.009)			(0.009)	(0.009)		(0.605)

Mean elasticities over the years 1992-2006. Standard deviations are reported in parentheses. Cross-circulation elasticities across different newspaper types are zero by construction.

5.6 Impact of Multi-Homing on Market Definition

A natural exercise to evaluate to what extent disregarding multi-homing in the estimation of demand in a two-sided market may affect policy decisions is to assess how, in competition policy, the definition of the relevant market for a two-sided platform changes. The issue is relevant for competition policy because market definition is often the first step in a competition policy case. In particular, the definition of the relevant market is crucial for cases of abuse of dominance: a wrong definition might lead to sanctioning behaviors that should not be sanctioned or to allow abuses that should be sanctioned. In regulated sectors dominance of a firm, assessed with respect to the identified relevant market, is often a necessary requirement to impose constraints on the firm's behavior. In fact, a wrong definition of the relevant market is sufficient for the courts to rule in favor of the parties irrespective of any other argument brought up by the antitrust or regulatory authorities.

The objective of market definition is to define a set of products that are substitute enough to a given product to pose a competitive constraint to the firm that produces it. While in theoretical economic models, the relevant market in which firms compete is often assumed as a starting point of the analysis, in reality identifying

the competitors of a given product sold by a firm is far less obvious. Most products are differentiated and one needs to identify the degree of differentiation below which products should be considered competitors (thus included in the relevant market) and above which they should not (thus excluded from the relevant market).

Since own-price elasticities increase with the number of available substitutes and cross-price elasticities measure the degree of substitution among products and decrease with differentiation, the correct estimation of own- and cross-price elasticities is crucial in order to define the relevant market correctly. This is true in a one-sided market as well as in a two-sided market. However, in a two-sided market own- and cross-network elasticities are also important (Filistrucchi, Geradin, Van Damme, and Affeldt, 2014).

In this section, we thus use the different estimates, based on the specifications accounting for readers multi-homing, presented in Tables 5.7 and 5.9 for the reader side and Tables 5.13 and 5.14 for the advertiser side to assess the impact of disregarding readers multi-homing on market definition.

Traditional newspapers are two-sided non-transaction platforms according to the definition given in Filistrucchi, Geradin, and van Damme (2013). Hence, following Filistrucchi, Geradin, Van Damme, and Affeldt (2014), one needs to define two separate but inter-related markets: one on the readers' and one on the advertisers' side. Starting from a given newspaper, in order to define the relevant market on the readers' side, the question to be addressed is which other newspapers are substitutes enough for a given newspaper that they impose a competitive constraint on its publisher. The equivalent question has to be answered for advertisers in order to define the relevant market on the advertisers' side.

While a complete answer to this question would require performing a SSNIP test, looking at the estimated own- and cross-price elasticities may already provide some partial answer. Similar to a one-sided market, also in a two-sided market, the sign of cross-price elasticities is enough to give an indication regarding the degree of substitutability or complementarity between products: when the cross-price elasticity is positive, products are substitutes and may be in the same relevant market; when the cross-price elasticity is instead negative, products are complements and they should not be included in the same relevant market. The same also holds true in two-sided markets. However, the price elasticities that need to be taken into consideration in a two-sided markets are not the partial price elasticities reported in Table 5.7 for the readers' side and Table 5.13 for the advertisers' side. Instead, these are the total price elasticities that one obtains when taking into account the indirect network effects reported in Table 5.9 for the readers' side and Table 5.14 for the advertisers' side. Such total price elasticities can be obtained by first calculating

the total marginal effects of a price increase, i.e. the effect that a price increase has on the quantity demanded, taking into account also the feedback loops between the two sides of the market.

Following Filistrucchi and Klein (2013), the matrix of total marginal effects of prices $\widehat{S}$ may be obtained as follows:

$$\widehat{S}=\begin{pmatrix} \widehat{S}^{rr} & \widehat{S}^{ar} \\ \widehat{S}^{ra} & \widehat{S}^{aa} \end{pmatrix} = -N^{-1}S = -\begin{pmatrix} -I & N^{ar} \\ N^{ra} & -I \end{pmatrix}^{-1} \begin{pmatrix} S^{r} & 0 \\ 0 & S^{a} \end{pmatrix}, \tag{5.14}$$

where the matrices N and S are respectively the matrix of the network effects and the matrix of the partial marginal effects of price.

More precisely, the block $\widehat{S}^{aa}$ is the matrix of total marginal effects of advertising prices on advertising demand, $\widehat{S}^{ar}$ is the matrix of total marginal effects of advertising prices on newspaper demand, $\widehat{S}^{rr}$ is the matrix of total marginal effects of cover prices on newspaper demand and $\widehat{S}^{ra}$ is the matrix of total marginal effects of cover prices on advertising demand, whereas N^{ra} is a matrix of externalities of readership on advertising and N^{ar} is a matrix of externalities of advertising on readership, and, finally, S^{a}, is a matrix of the marginal effects of advertising prices on advertising demand and S^{r} is a matrix of marginal effects of cover prices on newspaper demand.

The resulting total price elasticities are reported in Tables 5.15, 5.16, 5.17 and 5.18.

Table 5.15: Mean Total Own- and Cross-Price Elasticities - Readers' Demand - Double-Homing

	Corriere	Corriere Sport	Gazzetta Sport	Giornale	Repubblica	Il Sole	Stampa
Corriere	-2.303	-0.010	-0.041	0.041	0.381	-0.091	0.273
Corriere Sport	-0.023	-0.805	0.127	-0.009	-0.054	-0.019	-0.010
Gazzetta Sport	-0.070	0.093	-0.890	-0.008	-0.030	-0.025	-0.022
Giornale	0.118	-0.011	-0.014	-1.406	0.211	-0.052	0.099
Repubblica	0.419	-0.026	-0.020	0.083	-2.527	-0.098	0.288
Il Sole	-0.169	-0.016	-0.028	-0.032	-0.167	-1.062	-0.067
Stampa	0.482	-0.007	-0.023	0.062	0.462	-0.063	-2.154

Mean total elasticities based on the mean marginal effects over the years 1992-2006.

Table 5.16: Mean Total Own- and Cross-Price Elasticities - Advertisers' Demand - Double-Homing

	Corriere	Corriere Sport	Gazzetta Sport	Giornale	Repubblica	Il Sole	Stampa
Corriere	-0.652	0.013	0.030	0.014	0.024	0.056	0.005
Corriere Sport	-0.500	-0.314	-0.204	-0.096	-0.434	-0.302	-0.254
Gazzetta Sport	-0.407	-0.096	-0.464	-0.086	-0.411	-0.266	-0.218
Giornale	-0.328	-0.022	-0.028	-0.170	-0.350	-0.112	-0.174
Repubblica	0.004	0.016	0.025	0.010	-0.669	0.051	0.002
Il Sole	-0.467	-0.067	-0.102	-0.093	-0.442	-0.982	-0.239
Stampa	-0.266	-0.013	-0.011	-0.026	-0.211	-0.053	-0.606

Mean total elasticities based on the mean marginal effects over the years 1992-2006.

Table 5.17: Mean Total Own- and Cross-Price Elasticities - Cover Price on Advertisers' Demand - Double-Homing

	Corriere	Corriere Sport	Gazzetta Sport	Giornale	Repubblica	Il Sole	Stampa
Corriere	-2.705	-0.011	-0.047	0.071	0.479	-0.102	0.352
Corriere Sport	-0.194	-7.271	1.274	-0.076	-0.486	-0.170	-0.083
Gazzetta Sport	-0.399	0.609	-5.131	-0.043	-0.170	-0.144	-0.125
Giornale	0.792	-0.060	-0.070	-7.984	1.347	-0.270	0.706
Repubblica	0.538	-0.031	-0.023	0.125	-3.072	-0.114	0.382
Il Sole	-0.302	-0.028	-0.051	-0.057	-0.298	-1.893	-0.120
Stampa	1.143	-0.014	-0.049	0.184	1.106	-0.133	-4.861

Mean total elasticities based on the mean marginal effects over the years 1992-2006.

Table 5.18: Mean Total Own- and Cross-Price Elasticities - Advertising Price on Readers' Demand - Double-Homing

	Corriere	Corriere Sport	Gazzetta Sport	Giornale	Repubblica	Il Sole	Stampa
Corriere	0.278	-0.003	-0.003	-0.006	-0.059	-0.008	-0.033
Corriere Sport	0.094	0.078	-0.007	0.017	0.087	0.055	0.046
Gazzetta Sport	0.098	0.000	0.116	0.017	0.086	0.055	0.046
Giornale	0.079	0.013	0.018	0.060	0.070	0.060	0.039
Repubblica	-0.073	-0.001	-0.003	-0.006	0.229	-0.003	-0.033
Il Sole	0.085	0.009	0.015	0.013	0.075	0.113	0.037
Stampa	-0.003	0.009	0.015	0.009	0.010	0.039	0.210

Mean total elasticities based on the mean marginal effects over the years 1992-2006.

Estimated cross-price elasticities in Table 5.15 imply that newspapers of the same type are substitutes while newspapers of different types are complements in the eyes of readers. Estimated cross-price cross-side elasticities in Table 5.17 also confirm this finding. Hence, on the readers' side, only newspapers of the same type (generalist, sport, or financial) may be in the same relevant market (although they not necessarily are). Such a finding, which stems directly from the estimated partial price elasticities, would not have been possible in the absence of information on double-homing.

Turning to the relevant market on the advertising side, estimated cross-price elasticities in Table 5.16 show that there is substantial asymmetry in the cross-price elasticities, not only with regard to size but also with regard to signs. In fact, advertising price changes by big newspapers (i.e. newspapers that account for a great part of the advertising expenses, namely Corriere and La Repubblica) cause a loss in advertising sales to all other newspapers. In other words, there is a complementarity between big and small newspapers and substitutability only within. This is mainly due to the income effect and to the high aggregate price elasticity of newspaper advertising, as previously discussed. Hence, on the advertisers' side of the market it would appear that only newspapers of similar advertising importance may be in the same relevant market.

However, the picture is a bit more complex. In fact, in a two-sided market, a price increase has two effects on profits: one on the side of the market whose price has been raised and one on the other side of the market. On the readers' side of the market these two effects, which measure the competitive constraint faced by the given newspapers, move in the same direction. However, on the advertising side of the market, due to the finding that readers dislike advertising, the effects often move in opposite directions, as shown by the estimated cross-price cross-side elasticities reported in Table 5.18. Hence, the overall effect on quantities sold (and thus on profits) of a price increase is sometimes uncertain as it depends on the relative size of two, often opposite, effects. For instance, an increase in the advertising price of Corriere is estimated to lower the sales of advertising in that newspaper, but increases the number of readers of that newspaper. Such an increase is also estimated to have, on the one hand, a negative effect on the advertising sales of all other newspapers except La Repubblica, and, on the other hand, a positive effect on the circulation of all other newspapers but La Repubblica and La Stampa. Hence, in order to define the precise boundaries of the relevant market, a two-sided SSNIP test may be needed. Whereas such a test would require a more complex analysis, the evidence provided above is enough to suggest that accounting for multi-homing would indeed matter for the definition of the relevant market for antitrust and

regulatory purposes.

Finally, since in mergers, competition concerns arise only among firms producing substitute products, the evidence just presented suggests that accounting for multi-homing is also relevant for the assessment of unilateral effects of mergers, even though a more complete assessment would require performing a full merger simulation.

5.7 Conclusion

We study the role of multi-homing in the newspaper market. We first build a micro-founded structural econometric model that encompasses the demand for differentiated products on both sides of the market and allows for multi-homing on each side of the market. We then estimate the model above alternatively taking into account and not taking into account information on multi-homing by readers. Finally, we discuss to what extent disregarding multi-homing information may bias policy decisions, particularly in the field of competition policy.

Results for the readers' side of the market show that not accounting for multi-homing leads to a substantial bias in the estimation of own- and cross-price elasticities as well as own- and cross-network effect elasticities. In particular, mean own-price elasticities increase substantially when multi-homing by readers is taken into account in estimating demand. Furthermore, while we find that newspapers of the same type are substitutes, newspapers of different types are found to be complements in the model accounting for multi-homing. However, a discrete choice model at the newspaper level assumes that all readers single-home and that all newspapers are substitutes. Similarly, disregarding reader multi-homing information may bias the estimation of own- and cross-price as well as network effects elasticities also on the advertising side of the market. In particular, we find that own-price elasticities on the advertising side decrease with the number of captive readers while cross-price elasticities increase with the number of overlapping readers between newspapers.

Lastly, we look at the Italian market for national daily newspapers and the traditional market definition, distinguishing generalist, sport, and financial newspapers, that is potentially affected by the use of information on multi-homing. While a full-fledged market definition exercise would require performing a two-sided SSNIP test, our findings confirm the importance of incorporating multi-homing in the analysis. We find that, on the readers' side, only newspapers of the same type (generalist, sport, or financial) may be in the same relevant market (although they not necessarily are). On the advertisers' side of the market, instead, it appears that only newspapers of similar advertising importance may be in the same relevant market.

Our paper contributes to the economic literature on two-sided markets, in which empirical work accounting for multi-homing is still quite scarce. Moreover, our contribution allows for a better understanding of how multi-homing by users in platform markets matters and how it influences price elasticities on both sides of the market. This is likely to bias the conclusions of such exercises as market definition or merger evaluation in which both own- and cross-price elasticities and own- and cross-network effect elasticities play a crucial role. Thus, it can be useful for competition and regulation authorities to improve their quantitative assessment in cases involving two-sided platforms. While print newspapers are a classical example of an *offline* two-sided market, the empirical part of this paper should be seen as an application allowing for studying the role of multi-homing in (non-transaction) platform markets. The methodology can also be applied to other two-sided markets for which data on user multi-homing is available. Especially in light of the prevalence and rising importance of multi-sided platforms in digital markets and the relevance of multi-homing by users, the results and conclusions from this paper are also relevant in the context of competition policy cases involving online multi-sided platform markets.

5.8 Appendix

5.8.1 Data Sources

Table 5.19: Used Data and Corresponding Data Sources

Data	Source
Newspaper Data	
Average prints	Accertamenti Diffusione Stampa
Cover prices	Newspaper publishers
Magazines, inserts, local pages, websites, editors	Newspaper publishers
Number of pages	Nielsen Media Research
Advertising quantities and revenues	Nielsen Media Research
Data on multi-homing readers	Audipress
Reader demographics (age groups, gender, socio-economic status)	Audipress
Additional Data	
Italian consumer price index	OECD
Italian population above 14	ISTAT
Italian GDP	OECD
Paper price	Camera di Commercio di Milano
Hourly wage index printed media sector	ISTAT
Journalist hourly wage index	ISTAT
Electricity prices for industrial consumers	Eurostat
Italian internet penetration rate	ISTAT

5.8.2 Construction of Bundle Level Dataset

In order to estimate readers' demand accounting for multi-homing by readers, we estimate a demand equation where consumers are allowed to read up to two newspapers (which is what we observe in the data). Thus, readers are allowed to choose between all possible pairs of newspapers, including single-homing on a newspaper.

Hence, we construct a dataset in which the level of observation is no longer a newspaper at a particular point in time but a bundle of up to two newspapers for a given average weekday within a month.

We construct the monthly average circulation for each day of the week for a given bundle based on the survey data on multi-homing behavior of readers. First, for each newspaper we compute the share of single-homing (captive) readers, which also includes readers of other newspapers not included in our sample. Multiplying

this share of single-homing readers with the circulation of the newspaper gives the circulation of the single-homing "bundle" at a particular point in time.

Constructing the circulation of a bundle of two newspapers is more complicated as double-homing reader percentages are not symmetric. For example, the percentage of readers of newspaper A who also read newspaper B will not be the same as the percentage of readers of newspaper B who also read newspaper A. However, multiplying these percentages with the respective newspaper circulation provides a lower and an upper bound for the circulation of the bundle of the two newspapers at a particular point in time.

We then must decide on the optimal point within the interval between lower and upper bound bundle circulation. Of course, setting bundle circulation to any value within this interval implies that, if we calculate newspaper circulation as the sum of the circulation of all bundles containing that particular newspaper, overall newspaper level circulation will not be equal to the actual circulation of that newspaper in the original data.

Therefore, we choose the optimal point in the interval between lower and upper bound bundle circulation so as to minimize the difference between actual and estimated newspaper circulation for the seven newspapers at each point in time. We do so by running constrained regressions for each half-year period (the interval at which we have information on reader multi-homing). The dependent variable in this regression is the circulation of the respective newspaper minus the single-homing circulation minus the lower bound circulations of all bundles including this newspaper. Hence, the dependent variable is the part of the total newspaper circulation that we still need to distribute across bundles. The independent variables are then the differences between upper and lower bound circulation for all possible bundles of two newspapers, where this difference is set to zero for all bundles that do not include the respective newspaper. We then run constrained OLS regressions, in which all the coefficients are constrained to lie between zero and one - thus, the estimated coefficients give us the optimal points in the respective intervals between lower and upper bound bundle circulation for all bundles of two newspapers. We repeat this constrained regression for every half year period in the dataset and construct the bundle circulation for all bundles of two newspapers based on the estimated optimal points in the interval between lower and upper bound bundle circulation.

As mentioned before, this procedure implies that if we collapse the bundle level dataset back to the newspaper level, circulation at the newspaper level will not be equal to the actual newspaper circulation of the original dataset. Table 5.20 shows the mean actual and mean estimated newspaper circulation as well as the mean percentage difference between the two. The percentage difference between actual

and estimated newspaper circulation based on the constrained regression procedure is always lower than 8%.

Table 5.20: Difference between Actual and Estimated Circulation by Newspaper

Newspaper	Mean Actual Circulation	Mean Estimated Circulation	Mean Percentage Difference
Corriere	866,130	850,469	1.52
Corriere Sport	428,295	414,402	2.59
Gazzetta Sport	584,629	608,265	-4.28
Giornale	315,328	292,570	7.24
Repubblica	799,423	808,024	-1.48
Il Sole	462,816	443,114	4.22
Stampa	521,152	511,133	1.81

The construction of the other variables in the bundle level dataset is straightforward: the bundle price is the sum of the two newspaper prices, advertising intensity is the sum of advertising slots in both newspapers divided by the total number of pages of the two newspapers, the paper cost is the sum of the two newspaper costs, and the number of local pages is the sum of the local pages of the two newspapers. For the other newspaper characteristics, which are dummy variables (for example the point in time when the newspaper introduced a website), we also calculate the sum of the two dummy variables. However, this variable will then only capture when, for example, the first website in the bundle was introduced. Therefore, we define additional variables for those product characteristics that capture the change from zero to one for the second newspaper included in the bundle.

5.8.3 Structure of Nests

Figure 5.2: Structure of Nests - Single-Homing

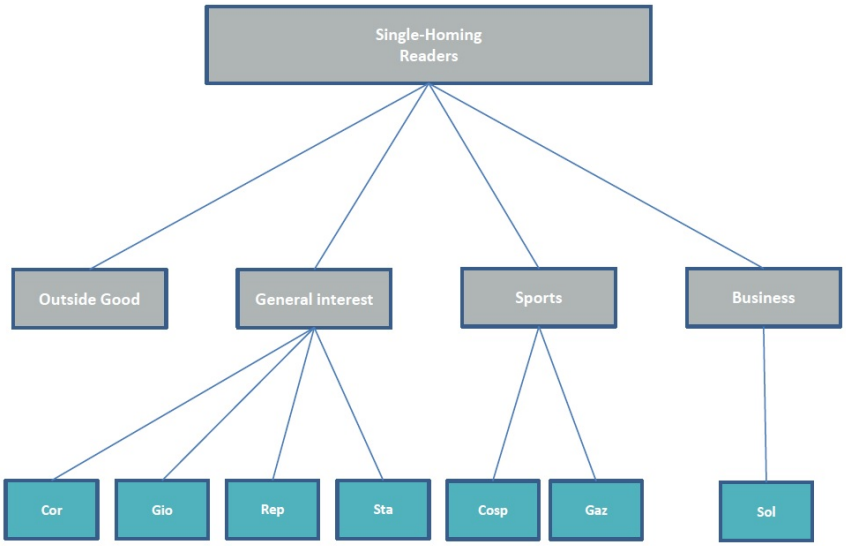

Figure 5.3: Structure of Nests - Double-Homing

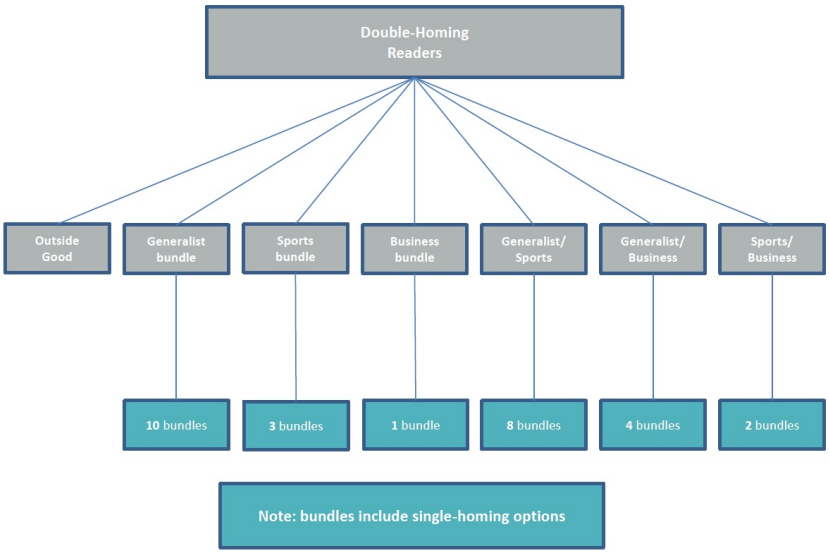

5.8.4 Elasticities Advertisers' Demand

5.8.4.1 Conditional Price Elasticities

Differentiating the advertising sales share w_{jt} of newspaper j at time t with respect to own- and cross-prices (p_{jt} and p_{kt}), holding x_t constant, gives:

$$\frac{\partial w_{jt}}{\partial p_{kt}} = \begin{cases} \frac{\partial(p_{jt}q_{jt}/x_t)}{\partial p_{jt}} = \frac{q_{jt}}{x_t} + \frac{p_{jt}}{x_t}\frac{\partial q_{jt}}{\partial p_{jt}} = \frac{q_{jt}}{x_t} + \frac{p_{jt}}{x_t}\frac{\partial q_{jt}}{\partial p_{jt}}\frac{q_{jt}}{q_{jt}} = \frac{q_{jt}}{x_t} + \frac{q_{jt}}{x_t}\tilde{\eta}_{jj} & \text{if } k = j \\ \frac{\partial(p_{jt}q_{jt}/x_t)}{\partial p_{kt}} = \frac{p_{jt}}{x_t}\frac{\partial q_{jt}}{\partial p_{kt}} = \frac{p_{jt}}{x_t}\frac{\partial q_{jt}}{\partial p_{kt}}\frac{p_{kt}}{q_{jt}}\frac{q_{jt}}{p_{kt}} = \frac{w_{jt}}{p_{kt}}\tilde{\eta}_{jk} & \text{if } k \neq j \end{cases} \tag{5.15}$$

Solving for $\tilde{\eta}_{jj}$ and $\tilde{\eta}_{jk}$ gives:

$$\tilde{\eta}_{jk} = \begin{cases} \frac{x_t}{q_{jt}}\frac{\partial w_{jt}}{\partial p_{jt}} - 1 & \text{if } k = j \\ \frac{\partial w_{jt}}{\partial p_{kt}}\frac{p_{kt}}{w_{jt}} & \text{if } k \neq j \end{cases} \tag{5.16}$$

Now, given the functional form of the estimation equation for w_{jt} in equation (5.10), the own- and cross-price derivatives, $\frac{\partial w_{jt}}{\partial p_{jt}}$ and $\frac{\partial w_{jt}}{\partial p_{kt}}$, are the following:

$$\frac{\partial w_{jt}}{\partial p_{kt}} = \begin{cases} \frac{\partial w_{jt}}{\partial \ln(p_{jt})}\frac{\partial \ln(p_{jt})}{\partial p_{jt}} = \frac{1}{p_{jt}}\frac{\partial w_{jt}}{\partial \ln(p_{jt})} = \frac{1}{p_{jt}}[b_{jj} - d_j w_{jt}] & \text{if } k = j \\ \frac{\partial w_{jt}}{\partial \ln(p_{kt})}\frac{\partial \ln(p_{kt})}{\partial p_{kt}} = \frac{1}{p_{kt}}\frac{\partial w_{jt}}{\partial \ln(p_{kt})} = \frac{1}{p_{kt}}[b_{jk} - d_j w_{kt}] & \text{if } k \neq j \end{cases} \tag{5.17}$$

Lastly, the conditional own- and cross-price elasticities are obtained by replacing (5.17) into (5.16):

$$\tilde{\eta}_{jk} = \begin{cases} -1 + \frac{b_{jj}}{w_{jt}} - d_j & \text{if } k = j \\ \frac{b_{jk}}{w_{jt}} - d_j\frac{w_{kt}}{w_{jt}} & \text{if } k \neq j \end{cases} \tag{5.18}$$

Note that the b_{jj} potentially is a function of own product characteristic while b_{jk} is a function of the closeness measures.

5.8.4.2 Unconditional Price Elasticities

Compared to the price elasticities of advertising demand $\tilde{\eta}_{jk}$ *conditional on total advertising expenditures* x_t, the *unconditional* advertising price elasticities η_{jk} also need to take into account how advertising price increases by one newspaper change the overall price index for advertising in newspapers, how this in turn then changes the overall demand for advertising in print newspapers at the top level (relative to other media outlets), and how this change in total advertising expenditures x_t then affects the advertising quantity demanded on newspaper j. See Heien and Pompelli (1988) for the general formula of the total own-price elasticity in the AIDS model.

Given the functional form of the estimation equations in equations (5.5) and (5.10) and the choice of the price index P_t, the unconditional price-elasticities for advertising demand η_{jk} are given by:

$$\eta_{jk} = \begin{cases} \tilde{\eta}_{jj} + (1 + \frac{d_j}{w_{jt}})(\beta_2 + 1)w_{jt} & \text{if } k = j \\ \tilde{\eta}_{jk} + (1 + \frac{d_j}{w_{jt}})(\beta_2 + 1)w_{kt} & \text{if } k \neq j \end{cases} \tag{5.19}$$

The additional term in comparison to the conditional price elasticities of advertising demand $\tilde{\eta}_{jk}$ is the percentage change in advertising quantity q_{jt} demanded on newspaper j following an advertising price increase of newspaper k that goes via its effect on total advertising expenditure x_t:

- w_{kt} is the elasticity of the price index P_t with respect to a change in the advertising price of newspaper k, p_{kt}.
- $(\beta_2 + 1)$ is the elasticity of total advertising expenditure x_t with respect to a change in the price index P_t (Heien and Pompelli, 1988, p.40), where β_2 is the overall price elasticity of advertising demand in the seven newspapers relative to other media outlets estimated at the top level (see top level estimation equation (5.5)).
- $(1 + \frac{d_j}{w_{jt}})$ is the expenditure elasticity, i.e. by how much a change in total advertising expenditure x_t changes the advertising quantity q_{jt} demanded on newspaper j (Alston, Foster, and Green, 1994, p.352). If a change in total advertising expenditure x_t leaves the advertising sales share of newspaper j, w_{jt}, unchanged (i.e. the estimated d_j in equation (5.10) is zero), the expenditure elasticity is 1. If, instead, w_{jt} changes with a change in x_t, the expenditure elasticity of newspaper j depends on the estimated d_j.

5.8.4.3 Circulation Elasticities

Circulation elasticities, which do not depend on total advertising expenditures x_t, are derived in a similar way as the conditional price elasticities. Differentiating the advertising sales share w_{jt} of newspaper j at time t with respect to own- and cross-circulation ($circ_{jt}$ and $circ_{kt}$) gives:

$$\frac{\partial w_{jt}}{\partial circ_{kt}} = \begin{cases} \frac{\partial(p_{jt}q_{jt}/x_t)}{\partial circ_{jt}} = \frac{p_{jt}}{x_t}\frac{\partial q_{jt}}{\partial circ_{jt}} = \frac{p_{jt}}{x_t}\frac{\partial q_{jt}}{\partial circ_{jt}}\frac{circ_{jt}}{q_{jt}}\frac{q_{jt}}{circ_{jt}} = \frac{w_{jt}}{circ_{jt}}\rho_{jj} & \text{if } k = j \\ \frac{\partial(p_{jt}q_{jt}/x_t)}{\partial circ_{kt}} = \frac{p_{jt}}{x_t}\frac{\partial q_{jt}}{\partial circ_{kt}} = \frac{p_{jt}}{x_t}\frac{\partial q_{jt}}{\partial circ_{kt}}\frac{circ_{kt}}{q_{jt}}\frac{q_{jt}}{circ_{kt}} = \frac{w_{jt}}{circ_{kt}}\rho_{jk} & \text{if } k \neq j \end{cases} \tag{5.20}$$

Solving for ρ_{jj} and ρ_{jk} gives:

$$\rho_{jk} = \begin{cases} \frac{\partial w_{jt}}{\partial circ_{jt}} \frac{circ_{jt}}{w_{jt}} & \text{if } k = j \\ \frac{\partial w_{jt}}{\partial circ_{kt}} \frac{circ_{kt}}{w_{jt}} & \text{if } k \neq j \end{cases} \tag{5.21}$$

Given the functional form of the estimation equation for w_{jt} in equation (5.10), the own- and cross-circulation derivatives, $\frac{\partial w_{jt}}{\partial circ_{jt}}$ and $\frac{\partial w_{jt}}{\partial circ_{kt}}$, are the following:

$$\frac{\partial w_{jt}}{\partial circ_{kt}} = \begin{cases} \frac{\partial w_{jt}}{\partial \ln(circ_{jt})} \frac{\partial \ln(circ_{jt})}{\partial circ_{jt}} = \frac{1}{circ_{jt}} \frac{\partial w_{jt}}{\partial \ln(circ_{jt})} = \frac{c_{jj}}{circ_{jt}} & \text{if } k = j \\ \frac{\partial w_{jt}}{\partial \ln(circ_{kt})} \frac{\partial \ln(circ_{kt})}{\partial circ_{kt}} = \frac{1}{circ_{kt}} \frac{\partial w_{jt}}{\partial \ln(circ_{kt})} = \frac{c_{jk}}{circ_{kt}} & \text{if } k \neq j \end{cases} \tag{5.22}$$

Lastly the own- and cross-circulation elasticities are obtained by replacing (5.22) into (5.21):

$$\rho_{jk} = \begin{cases} \frac{c_{jj}}{w_{jt}} & \text{if } k = j \\ \frac{c_{jk}}{w_{jt}} & \text{if } k \neq j \end{cases} \tag{5.23}$$

with c_{jj} potentially being a function of own product characteristics and c_{jk} being a function of the closeness measures.

5.8.5 Determinants of Advertisers' Demand Elasticities

Table 5.21: Mean Characteristics by Newspaper, 1992-2006

	Advertising expenditure share	Percentage of female readers	Percentage of single-homing readers
Corriere	28.09	40.97	45.81
Corriere Sport	3.92	13.08	41.45
Gazzetta Sport	6.21	12.67	50.08
Giornale	5.65	38.44	29.02
Repubblica	26.03	41.21	51.42
Il Sole	16.25	33.55	25.03
Stampa	14.27	41.7	61.63

Chapter 6

Concluding Remarks

This dissertation consists of three essays that empirically study European merger control, while Chapter 2 documents the data used in the first two essays. These two essays study the determinants and predictability of DG Comp's merger decisions.

The first essay studies the time-dynamics of the EC's merger decision procedure over the first 25 years of European merger control using a relevant market level dataset containing all merger cases with an official decision documented by DG Comp. Specifically, we evaluate how consistently different arguments related to the structural market parameters – market shares, concentration, likelihood of entry, and foreclosure – put forward to motivate a particular decision were applied over time. Using LPM models and non-parametric machine learning techniques, we find that the importance of market share and concentration measures has declined over time while the importance of barriers to entry and the risk of foreclosure has increase in DG Comp's merger assessment following the 2004 merger policy reform.

In the second essay, I analyze the predictability of the EC's merger decision procedure before and after the 2004 merger policy reform. Using the highly flexible, non-parametric random forest algorithm to predict DG Comp's assessment of competitive concerns in markets affected by a merger, I find that the predictive performance of the random forests is much better than the performance of simple linear models. In particular, the random forests do much better in predicting the rare event of competitive concerns. Secondly, post-reform, DG Comp seems to base its assessment on a more complex interaction of merger and market characteristics than pre-reform. The highly flexible random forest algorithm is able to detect these potentially complex interactions and, therefore, still allows for high prediction precision.

Unlike the macro perspective of the first two essays, the third essay looks at one particular market. Specifically, we investigate empirically the impact of multi-homing in two-sided markets. We first build a micro-founded structural econometric

model, which encompasses demand for differentiated products and allows for multi-homing on both sides of the market. We then use an original dataset on the Italian daily newspaper market that includes information on double-readership of newspapers to estimate demand. We estimate the model above and compare the estimation results obtained in the presence of information on multi-homing to those obtained in its absence. The results show that an econometric model that does not allow for multi-homing is likely to produce biased estimates of own- and cross-price elasticities on both sides of the market. Finally, we discuss to what extent disregarding multi-homing information may distort policy exercises, such as market definition.

In conclusion, this dissertation provides an in-depth analysis of European merger control. While the first two essays analyze the determinants and predictability of DG Comp's merger decisions at the aggregate and over the past 25 years, the last essay studies the particular issue of multi-homing in two-sided markets and how it matters for the assessment of antitrust cases. Therefore, the dissertation relates to two currently widely debated topics. The first is the current discussion about the global rise of concentration, profits, mark-ups, and market power across many markets and industries as well as the role and importance of competition policy as one tool to foster competition and prevent abusive behavior. The second is the discussion about whether and how the competition policy toolkit might need to be re-designed in order to provide competition authorities with the analytical tools required to analyze multi-sided markets in the digital economy.

Bibliography

Affeldt, P., T. Duso, and F. Szücs (2018): "EU Merger Control Database: 1990-2014," *DIW Berlin Data Documentation 95*.

——— (2019): "25 Years of European Merger Control," *DIW Discussion Paper No. 1797*.

Affeldt, P., L. Filistrucchi, and T. J. Klein (2013): "Upward Pricing Pressure in Two-sided Markets," *The Economic Journal*, 123(572), F505–F523.

Alston, J. M., K. A. Foster, and R. D. Green (1994): "Estimating Elasticities with the Linear Approximate Almost Ideal Demand System: Some Monte Carlo Results," *Review of Economics and Statistics*, 76(2), 351–356.

Ambrus, A., E. Calvano, and M. Reisinger (2016): "Either or Both Competition: A 'Two-Sided' Theory of Advertising with Overlapping Viewerships," *American Economic Journal: Microeconomics*, 8(3), 189–222.

Anderson, S. P., Ø. Foros, and H. J. Kind (2018): "Competition for Advertisers and for Viewers in Media Markets," *The Economic Journal*, 128(608), 34–54.

Anderson, S. P., and J. J. Gabszewicz (2006): "The Media and Advertising: A Tale of Two-Sided Markets," in *Handbook of the Economics of Art and Culture*, ed. by V. A. Ginsburgh, and D. Throsby, vol. 1, pp. 567–614. Elsevier.

Anderson, S. P., and B. Jullien (2015): "The Advertising-Financed Business Model in Two-Sided Media Markets," in *Handbook of Media Economics*, ed. by S. P. Anderson, J. Waldvogel, and D. Stromberg, vol. 1, pp. 41–90. NorthHolland.

Argentesi, E., and L. Filistrucchi (2007): "Estimating Market Power in a Two-Sided Market: The Case of Newspapers," *Journal of Applied Econometrics*, 22(7), 1247–1266.

Armstrong, M. (2006): "Competition in Two-Sided Markets," *RAND Journal of Economics*, 37(3), 668–691.

Athey, S. (2018): "The Impact of Machine Learning on Economics," Mimeographed.

Athey, S., E. Calvano, and J. S. Gans (2018): "The Impact of Consumer Multi-homing on Advertising Markets and Media Competition," *Management Science*, 64(4), 1574–1590.

Athey, S., and G. Imbens (2016): "Recursive partitioning for heterogeneous causal effects," *Proceedings of the National Academy of Sciences of the United States of America*, 113(27), 7353–7360.

Athey, S., J. Tibshirani, and S. Wager (2017): "Generalized Random Forests," Mimeographed.

Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen (2017): "The Fall of the Labor Share and the Rise of Superstar Firms," *NBER Working Paper No. 23396*.

Bergman, M., M. B. Coate, A. Mai, and S. W. Ulrick (2016): "Does Merger Policy Converge after the 2004 European Union Reforms?," Mimeographed.

Bergman, M. A., M. B. Coate, M. Jakobsson, and S. W. Ulrick (2007): "Comparing Merger Policies: The European Union versus the United States," *Potomac Papers in Law and Economics*, (07-01).

Bergman, M. A., M. B. Coate, M. Jakobsson, and S. W. Ulrick (2010): "Comparing Merger Policies in the European Union and the United States," *Review of Industrial Organization*, 36(4), 305–331.

Bergman, M. A., M. Jakobsson, and C. Razo (2005): "An econometric analysis of the European Commission's merger decisions," *International Journal of Industrial Organization*, 23(9-10), 717–737.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63(4), 841–890.

Berry, S. T. (1994): "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, 25(2), 242–262.

Björkegren, D., and D. Grissen (2018): "Behavior Revealed in Mobile Phone Usage Predicts Credit Repayment," Mimeographed.

Bradford, A., R. J. Jackson, and J. Zytnick (2018): "Is E.U. Merger Control Used for Protectionism? An Empirical Analysis," *Journal of Empirical Legal Studies*, 15(1), 165–191.

Breiman, L. (2001): "Random Forests," *Machine Learning*, 45(1), 5–32.

Caillaud, B., and B. Jullien (2001): "Competing Cybermediaries," *European Economic Review*, 45(4-6), 797–808.

——— (2003): "Chicken and Egg: Competition Among Intermediation Service Providers," *RAND Journal of Economics*, 34(2), 309–328.

Chalfin, A., O. Danieli, A. Hillis, Z. Jelveh, M. Luca, J. Ludwig, and S. Mullainathan (2016): "Productivity and Selection of Human Capital with Machine Learning," *American Economic Review: Papers and Proceedings*, 106(5), 124–127.

Chen, C., A. Liaw, and L. Breiman (2004): "Using Random Forests to Learn Imbalanced Data," Mimeographed.

Clougherty, J., T. Duso, M. Lee, and J. Seldeslachts (2016): "Effective European Antitrust: Does EC Merger Policy Generate Deterrence?," *Economic Inquiry*, 54(4), 1884–1903.

Clougherty, J., and J. Seldeslachts (2013): "The Deterrence Effects of US Merger Policy Instruments," *Journal of Law, Economics and Organization*, 29, 1114–1144.

Crémer, J., Y.-A. De Montjoye, and H. Schweitzer (2019): "Competition policy for the digital era," Report, European Commission.

De Loecker, J., J. Eeckhout, and G. Unger (2018): "The Rise of Market Power and the Macroeconomic Implications," Mimeographed.

Deaton, A., and J. Muellbauer (1980): "An Almost Ideal Demand System," *American Economic Review*, 70(3), 312–326.

Dewenter, R., U. Heimeshoff, and F. Löw (2017): "Market Definition of Platform Markets," *Working Paper No. 176, Helmut Schmidt University Hamburg.*

Duso, T. (2012): "A Decade of Ex-post Merger Policy Evaluations: A Progress Report," in *The Pros and Cons of Merger Control*, ed. by D. Sjöblom, pp. 125–188. Swedish Competition Authority.

Duso, T., K. Gugler, and F. Szücs (2012): "An Empirical Assessment of the 2004 EU Merger Policy Reform," *DICE Discussion Paper No.58*.

——— (2013): "An Empirical Assessment of the 2004 EU Merger Policy Reform," *Economic Journal*, 123(572), F596–F619.

Duso, T., K. Gugler, and B. B. Yurtoglu (2011): "How Effective is European Merger Control?," *European Economic Review*, 55(7), 980–1006.

Duso, T., D. J. Neven, and L.-H. Röller (2007): "The Political Economy of European Merger Control: Evidence using Stock Market Data," *Journal of Law & Economics*, 50(3), 455–489.

Elman, P. (1965): "The Need for Certainty and Predictability in the Application of the Merger Law," *New York University Law Review*, 40, 613–627.

Evans, D. S. (2003): "The Antitrust Economics of Multi-Sided Platform Markets," *Yale Journal on Regulation*, 20(2), 325–382.

Evans, D. S., and R. Schmalensee (2007): "The Industrial Organization of Markets with Two-Sided Platforms," *Competition Policy International*, 3(1), 151–179.

Fan, Y. (2013): "Ownership Consolidation and Product Characteristics: A Study of the U.S. Daily Newspaper Market," *American Economic Review*, 103(5), 1598–1628.

Filistrucchi, L. (2005): "The Impact of Internet on the Market for Daily Newspapers in Italy," *EUI Working Paper ECO No. 2005/12*.

Filistrucchi, L., D. Geradin, and E. van Damme (2013): "Identifying Two-Sided Markets," *World Competition*, 36(1), 33–59.

Filistrucchi, L., D. Geradin, E. Van Damme, and P. Affeldt (2014): "Market Definition in Two-Sided Markets: Theory and Practice," *Journal of Competition Law and Economics*, 10(2), 293–339.

Filistrucchi, L., and T. J. Klein (2013): "Price Competition in Two-Sided Markets with Hetergoenous Consumers and Network Effetcs," *NET Institute Working Paper No. 13-20*.

Filistrucchi, L., T. J. Klein, and T. Michielsen (2012): "Assessing Unilateral Merger Effects in a Two-Sided Market: An Application to the Dutch Daily Newspaper Market," *Journal of Competition Law and Economics*, 8(2), 297–329.

Gentzkow, M. (2007): "Valuing New Goods in a Model with Complementarity: Online Newspapers," *American Economic Review*, 97(3), 713–744.

Gentzkow, M., J. M. Shapiro, and M. Sinkinson (2014): "Competition and Ideological Diversity: Historical Evidence from US Newspapers," *American Economic Review*, 104(10), 3073–3114.

Gerber, D. J. (2014): "Searching for a Modernized Voice: Economics, Institutions, and Predictability in European Competition Law," *Fordham International Law Journal*, 37(5), 1421–1450.

Grullon, G., Y. Larkin, and R. Michaely (2018): "Are U.S. Industries Becoming More Concentrated?," *Forthcoming in Review of Finance.*

Gutiérrez, G., and T. Philippon (2017): "Declining Competition and Investment in the U.S.," *NBER Working Paper No. 23583.*

——— (2018): "How EU Markets Became More Competitive Than US Markets: A Study of Institutional Drift," *NBER Working Paper No. 24700.*

Hartman-Glaser, B., H. N. Lustig, and M. Z. Xiaolan (2018): "Capital Share Dynamics When Firms Insure Workers," *Forthcoming in Journal of Finance.*

Hastie, T., R. Tibshirani, and J. Friedman (2008): *The Elements of Statistical Learning - Data Mining, Inference, and Prediction*, Springer Series in Statistics. Springer New York Inc.

Hausman, J., G. Leonard, and J. D. Zona (1994): "Competitive Analysis with Differentiated Products," *Annales d'Economie et de Statistique*, 34, 159–180.

Heien, D., and G. Pompelli (1988): "The Demand for Beef Products: Cross-Section Estimation of Demographic and Economic Effects," *Western Journal of Agricultural Economics*, 13(1), 37–44.

Ivaldi, M., and C. Muller-Vibes (2018): "The Differentiated Effect of Advertising on Readership: Evidence from a Two-Sided Market Approach," *Toulouse School of Economics Working Paper Number 18-900.*

Jeitschko, T. D., and M. J. Tremblay (2018): "Platform Competition with Endogenous Homing," Mimeographed.

Kaiser, U., and J. Wright (2006): "Price Structure in Two-Sided Markets: Evidence from the Magazine Industry," *International Journal of Industrial Organization*, 24(1), 1–28.

Kleinberg, J., H. Lakkaraju, J. Leskovec, J. Ludwig, and S. Mullainathan (2018): "Human Decisions and Machine Prediction," *Quarterly Journal of Economics*, 133(1), 237–293.

Kleinberg, J., J. Ludwig, S. Mullainathan, and Z. Obermeyer (2015): "Prediction Policy Problems," *American Economic Review*, 105(5), 491–495.

Kuhn, M. (2008): "Building Predictive Models in R Using the caret Package," *Journal of Statistical Software, Articles*, 28(5), 1–26.

Kwoka, J. E. (2013): "Does Merger Control Work? A Retrospective on U.S. Enforcement Actions and Merger Outcomes," *Antitrust Law Journal*, 78(3), 619–650.

Liaw, A., and M. Wiener (2002): "Classification and Regression by randomForest," *R News*, 2(3), 18–22.

Liu, Y.-H. (2018): "The Impact of Consumer Multi-homing Behavior on Ad Prices: Evidence from an Online Marketplace," Mimeographed.

Lyons, B. R. (2004): "Reform of European Merger Policy," *Review of International Economics*, 12(2), 246–261.

Mai, A. T. V. (2016): "Is EU Merger Policy Less Stringent After Its 2004 Reform?," *PESO Working Papers 2016:1*.

McAfee, P. R. (2010): "Transparency and Antitrust Policy," Mimeographed.

Mini, F. (2018): "Fifty is the New Forty: EU Merger Policy Permits Higher Market Shares After the 2004 Reform," *Review of Industrial Organization*, 53(3), 535–561.

Monopolkommission (2015): "Competition policy: The challenge of digital markets," Special Report No 68, Monopolkommission.

Motta, M. (2004): *Competition Policy - Theory and Practice*. Cambridge University Press.

Nevo, A. (2000): "A Practioner's Guide to Estimation of Random Coefficients Logit Models of Demand," *Journal of Economics & Management Strategy*, 9(4), 513–548.

Noel, M. D., and D. S. Evans (2005): "Analyzing Market Definition and Power in Multi-Sided Platform Markets," SSRN Electronic Journal.

OECD (2018): "Rethinking Antitrust Tools for Multi-Sided Platforms," Report, OECD.

Parker, G. G., and M. W. van Alstyne (2005): "Two-Sided Network Effects: A Theory of Information Product Design," *Management Science*, 51(10), 1494–1504.

Patterson, D., and C. Shapiro (2001): "Transatlantic Divergence in GE/Honeywell: Causes and Lessons," *Antitrust Magazine*, 16, 18–26.

Pinkse, J., and M. Slade (2004): "Mergers, Brand Competition and the Price of a Pint," *European Economic Review*, 48(3), 617–643.

Pinkse, J., M. Slade, and C. Brett (2002): "Spatial Price Competition: a Semiparametric Approach," *Econometrica*, 70(3), 1111–1155.

Ribers, M., and H. Ullrich (2018): "Battling Resistance: Using Machine Prediction to Improve Antibiotic Prescribing," Mimeographed.

Rochet, J.-C., and J. Tirole (2002): "Cooperation among Competitors: Some Economics of Payment Card Associations," *RAND Journal of Economics*, 33(4), 549–570.

——— (2003): "Platform Competition in Two-Sided Markets," *Journal of the European Economic Association*, 1(4), 990–1029.

——— (2006): "Two-sided Markets: A Progress Report," *RAND Journal of Economics*, 37(3), 645–667.

Rojas, C. (2008): "Price Competition in U.S. Brewing," *The Journal of Industrial Economics*, 56(1), 1–31.

Rojas, C., and E. B. Peterson (2008): "Demand for differentiated products: Price and advertising evidence from the U.S. beer market," *International Journal of Industrial Organization*, 26(1), 288–307.

Rubin, D. B. (1974): "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies," *Journal of Educational Psychology*, 66(5), 688–701.

Rysman, M. (2004): "Competition between Networks: A Study of the Market for Yellow Pages," *Review of Economic Studies*, 71(2), 483–512.

——— (2007): "An Empirical Analysis of Payment Card Usage," *The Journal of Industrial Economics*, 55(1), 1–36.

Shi, M. C. (2015): "Catching (Exclusive) Eyeballs: Multi-Homing and Platform Competition in the Magazine Industry," Mimeographed.

Slade, M. (2004): "Market Power and Joint Dominance in UK Brewing," *The Journal of Industrial Economics*, 52(1), 133–163.

Smith, B. (1958): "Precedent, Public Policy and Predictability," *Georgetown Law Journal*, 46, 632–645.

Szücs, F. (2012): "Investigating transatlantic merger policy convergence," *International Journal of Industrial Organization*, 30(6), 654–662.

Tuszynski, J. (2014): *caTools: Tools: moving window statistics, GIF, Base64, ROC AUC, etc.*R package version 1.17.1.

van der Wurff, R., P. Bakker, and R. G. Picard (2008): "Economic Growth and Advertising Expenditures in Different Media in Different Countries," *Journal of Media Economics*, 21(1), 28–52.

Wager, S., and S. Athey (2017): "Estimation and Inference of Heterogeneous Treatment Effects using Random Forests," Mimeographed.

Wismer, S., and A. Rasek (2018): "Market Definition in Multi-Sided Markets," in *Rethinking Antitrust Tools for Multi-Sided Platforms*, pp. 37–54. OECD.

Wright, J. (2004): "One-Sided Logic in Two-Sided Markets," *Review of Network Economics*, 3(1), 44–64.