

Essays on industrial organisation

BY
TIMOTHY GUY ALARIC WYNDHAM

NHH



PhD THESIS

NO.2020/03

APRIL 2020

ISBN/978-82-405-0406-9



CEMS



Department of Economics



Dissertation Abstract

This dissertation examines the implications of changes in technology and government regulations for the prices consumers pay and the volumes they consume in two important markets: the market for news and the market for alcohol.

Digital technology has transformed the newspaper market, a market where preferential tax treatment of printed newspapers has been widespread although this has not always been extended to digital newspapers. Given the fundamental similarity of content between digital and printed newspapers this digital divide is under pressure. But, what are the implications of lowering the tax rate on digital news? Will it lead to lower prices and more consumption? In a model which allows consumers to buy more than one source of news, and for newspapers to have different cost structures the first paper, published in the *International Journal of Industrial Organisation*, finds that reducing VAT rates on digital newspapers leads to higher prices for readers.

The second paper, published in *European State Aid Quarterly*, describes and applies the findings of the first paper and associated literature to a decision made by the European Free Trade Area Surveillance Authority (ESA). The decision permitted a temporary reduction of VAT on digital newspapers in Norway. We confirm that even a tax reduction which equalises tax conditions between printed and digital newspapers qualifies as State Aid. However, the Norwegian authorities argued that the State aid should be allowed as the tax reduction would reduce prices and increase consumption. This finding goes against the findings of Chapter 1, and as such we argue that the ESA did not have grounds to accept the aid as being compatible with the single market.

The third paper uses the recent expansion of Vinmonopolet to examine the relationship between alcohol consumption and sick leave. Our core finding is that an increase in alcohol consumption of 1 percent leads to an increase of sick leave in men of around 0.3 per 10,000 men, an increase of around 0.16 percent, at the mean.

The fourth paper returns to the market for newspapers and investigates the impact of digital newspapers imposing paywalls on news consumption. This is a central question for digital newspapers, as they need to carefully weigh the increased revenue from

charging readers for news against lost advertising revenues if they receive less views. When newspapers introduce a paywall we find that short run consumption of news falls by 3-4 percent, and continues to fall by between 9-11 percent in the long run.

Timothy Wyndham

Essays on industrial organisation

NHH



June 2, 2019

Dissertation for the Ph.D. degree
Department of Economics
Norwegian School of Economics

Acknowledgements

Writing a PhD thesis is hard: many people have helped me along the way.

First and foremost, I would like to thank my supervisor Hans Jarle Kind, without whom this thesis would not have been possible. Hans represents everything you might want from a supervisor. Fast, firm, fair and with enough room for fun too. I have been spoiled by the speed and quality of his feedback, as well as his eye for detail. It can be hard to identify which ideas need more work and which should be dropped. From a student perspective it is particularly useful to understand when something is unlikely to be a paper, even if you've just spent the last two weeks burrowing into it. I expect it is difficult to find the optimum balance between allowing the student the freedom to fail, and not letting them losing too much time. I am grateful for all of his feedback and support. I was also lucky to have Lars Sørgard as my supervisor in my first year, before his return to Konkurransetilsynet. It was a big jump from policy to academia and Lars made the transition much smoother with sage advice and plenty of patience! My loss was definitely competition policy's gain. I am also grateful for the discussions I've had with Marko Koethenbuerger my co-supervisor.

I've been lucky to work with great co-authors. The paper with Hans Jarle and Øystein contained a lot of firsts; first moment of enlightenment, first conference invite, first desk rejection, first double guessing of what a reviewer comment really meant and, finally, first acceptance. Perhaps Øystein is one of the few who can match Hans Jarle for both the speed and quality of comments. And generosity of spirit.

My other co-authors have also taught me a lot. The paper with Oddmund Berg has been a fantastic experience. I doubt I will ever work on any project with such a good moment to bad moment ratio. I suspect Oddmund will; you make your own luck, after all, even if modesty prevents you from admitting it. Malgorzata Cyndecka's patience in explaining the subtleties of State aid law to me (and with a tardy Editor, for other reasons) was exemplary. As was her

resilience. I'm so glad that paper found its home where, and when, it did.

Good research should have a chance of failing. Looking back, I'm not quite sure how I only share one paper with Frode Steen and none with Simen Ulsaker (and that one with Mathias would have been really, really fun!). I am grateful to both Frodes for keeping the faith when I lost mine; to Frode Skjeret for his speedy coding and to Frode Steen for, in essence, becoming an unofficial deputy supervisor and a bit more, too. Simen was a guide throughout the PhD, friendly modest and reassuringly good at theory.

For me, work-duty was a misnomer. Graced by a linguistic restriction I avoided the bulky gradings during exam time that the Scandinavians received. I will always be particularly grateful to Linda Nstbakken for her guidance, trust and willingness to challenge me as an instructor. Her impact on my career will be longstanding, even if Kazakhstan slipped through the net!

Given my background I had hoped to find a part-time home in BECCLE. I could not have wished for better colleagues. Teis, in particular, was a wonderful (and humble) support throughout and I am grateful to Tommy for his leadership of BECCLE during my period of deepest involvement. More widely, I am grateful to Ingrid, Hvard, Hvard, Ignacio and Øyvind for fun evenings in Bergen and beyond.

Closer to home, at NHH I received a lot of support from colleagues in what became the Center for Business Economics. Steffen and Chiara were always very generous with written comments, whilst Malin and Chang-Koo did a lot of the heavy lifting to ensure we had a steady stream of visitors that are so important for young researchers. A softly spoken word of support from Dirk was always well received. It was great that Charlotte (eventually) decided to join the PhD programme and CBE, I expect she will move up very quickly.

Within the department, I received a lot of day-to-day support from Dagny, Irene and Christel. Friendly, efficient and patient. Perfect. More broadly, Anne-Liv provided excellent drafting advice. So many colleagues were willing to spare time around the coffee machines, it seems unfair to single out any, but thank you to Liam and Roger.

Sometimes it can be hard to do enough sports and social activities during the PhD. On the sporting side, a big thank you goes out to: Maja, Ingar and Lars my spiritual Stoltzen guides; Rolf and Mads my spinning buddies; Felix, Mario and the NHH squash crew and to Ole Andreas for the runs. I won't be the last foreigner to be grateful that Klara and Norunn work at NHH. On

the social side there were some fantastic Friday nights of poker, board games, and the odd party and nachspill. Thomas, Lassi, Thor-Andreas, Luca, Ingrid, Jack, Anja, Pat and Ingeborg all played their part in lightening the load. NHH has the best people and the best parties.

I felt a really close bond with all who started in SAM with me. Ingar and I will always have Brussels, the pillars and my daughter's invigilation! Ranveig's warmth away from work is only matched by her determination within it. Xiaogeng, thank you for the dumplings, your laugh and the endless spoiling of Isabella. The SAM stipendiat group more widely is a special group.

Looking further backwards, I am grateful to three people in particular for their contribution to me as an economist. I am grateful to "Mr" Tidball for showing me the subject and believing in my capabilities. I am also grateful to Rich Dubourg who showed me how to use economics in the real world. Basically, it was great. I was lucky to be taught by Jaap Bos, the first to show me how academic economics should really be done.

Mum and Dad have always been there. Believing and supporting. Thank you. I could not have asked for a more caring and level headed sister than Ginny.

I am not an easy man to live with. Tidyness is not an attribute I am blessed with or, indeed, much care for. Living with a PhD student in a foreign country is not straightforward. I am grateful to Victoria and Isabella for sharing this crazy adventure.

Contents

Chapter I	Introduction	01
Chapter II	Tax-free digital news?	08
Chapter III	Price Increasing Tax Reductions for Electronic News- papers: Implications for State Aid Policy	40
Chapter IV	Outlet proximity, alcohol sales and sick leave: Evidence from Norway	62
Chapter V	Paywalls and the demand for online news	87

CHAPTER I

Introduction

Introduction

The joy of economics is its ability to shed light on the human world around us. Economics helps to answer interesting and important questions.

My economic awakening was 16 September 1992, also known as Black Wednesday. In the space of a day, the UK's base interest rate was increased from 10 percent to 12 percent and then to 15 percent as the UK Government desperately tried to remain within the Exchange Rate Mechanism before falling back the next day after the interventions had failed. So many questions! Why did the UK have to raise its interest rate? Why did the intervention fail? Why did the UK want to be in an Exchange Rate Mechanism?

Over time the interest in economics has been maintained and increased, but the questions have changed. What are the economic costs and benefits of migration? To what extent can we use economic instruments to reduce crime? How can the UK best support economic reform in Nigeria? How do economics and law interact? How best to incentivise monopolists with complicated services to deliver? How much can we trust competition in oligopolistic markets? What do we know about competitive market outcomes in multi-sided markets?

This final question, spurred by my involvement from the sidelines of the UK's domestic and international intervention on Interchange Fees, drove my interest in doing economic research. Some nuances and static/dynamic questions aside, economics had always been resolutely clear, lower prices in a market are likely to increase allocative efficiency; consumers gain more than firms lose. But digital markets and network effects complicate this picture. In these markets, when should we want more competition?

The economics of three of the papers in this dissertation have been driven by the insights of Harold Hotelling. Hotelling had an unlikely start to his academic career. His undergraduate degree was a BA in Journalism, hardly a common predictor of a career in statistics and mathematical economics. A Masters and Doctorate in Mathematics followed by 1924 and in the following years

he published two highly cited papers in economics.

“Stability in Competition” (Hotelling, 1929) was published in *The Economics Journal*. The intuition of the paper is a thing of great aesthetic beauty, a jewel set in a glass case;¹ and belies the author’s analytical background. Over time it has been reduced to an analogy about ice-cream. What is the equilibrium location of two ice-cream vendors on a beach? Since beachgoers prefer a short walk to a long walk, if the vendors locate at either end of the beach, the vendors could charge a price that extracts some of the benefit of the walking time saved by its nearest consumers: the competitive pressure between the two vendors would be weak, and profits plentiful. However, holding the other vendor’s location and price constant, each vendor has an incentive to move a little towards the middle. That vendor would capture some of customers in the middle of the beach whilst still attracting consumers on its own extremity. When a consumer’s travel cost is linear with respect to the distance travelled, the equilibrium outcome, according to the paper, is that they each locate in the middle selling to either side. Hotelling was clear that his mathematical analysis had far wider applications than just distance. Indeed, he observed if consumers had different preferences for the sourness of cider, then his approach could be used for product differentiation as well as spatial competition. Since I like both Hotelling and alcohol, I will come back to these subjects in Chapter 4.

Unfortunately, the (pure-strategy) equilibrium does not exist. Fifty years later d’Aspremont, Gabszewicz, and Thisse (1979) proved as much, and then showed that if transport costs are quadratic in distance and consumers are homogenous apart from their location, the conclusion inverts. The demand expansion effect that pulls firms to the centre is always dominated by the price increasing effect that pushes firms to the poles. However, Böckem (1994) rejects this maximum differentiation result; she shows that their result is not robust to allowing consumers to vary in their valuation of an outside option, as well as by location. Instead, she argues, we should expect neither minimal nor maximal differentiation.

Despite its shortcoming, Hotelling’s framework has become the standard unit of analysis for horizontally differentiated products. For questions where we can assume locations are fixed we use a linear transport cost, as I use in Chapters 2 and 3. For questions where locations vary, a quadratic cost is normally used. The more literal use of the framework has remained a relevant

¹Language credit: (Lancaster, 1966) used this phrase to describe the theory of consumer behaviour.

tool in many retail markets, where consumers need to physically travel for purchases. This insight drove the approach in Chapter 4.

In Chapter 2, forthcoming in the *International Journal of Industrial Organisation*, Øystein Foros, Hans Jarle Kind and I study the implications of reducing VAT on digital newspapers. Preferential tax treatment of printed newspapers is widespread, although this has not always been extended to digital newspapers. Given the fundamental similarity of content between digital and printed newspapers this digital divide is under pressure. But, what are the implications of lowering the tax rate on digital news? Will it lead to lower prices and more consumption? In a model which allows consumers to buy more than one source of news, and for newspapers to have different cost structures we find that reducing VAT rates on digital newspapers leads to higher prices for readers.

The intuition for this result stems from the balancing act that newspapers have to perform between their two sources of profits: readers and advertisers. Increasing reader prices leads to higher profits from readers, but reduces advertising profits as the higher price implies fewer readers. Reducing VAT on reader subscriptions makes the reader market more attractive whilst the advertising market is unaffected. Thus, the digital newspaper will put more weight on revenue from the reader than before the change in VAT. The higher profits from readers now offset the reduced profit from advertisers.

Previous papers looking at the impact of indirect taxes on newspapers assumed that consumers buy only one paper and that rivals had symmetric cost structures. We relax both these assumptions.

Relaxing the assumption that consumers only buy one newspaper is important for two key reasons. First, it has long been the case that some consumers have bought more than one newspaper. Survey evidence from a century ago ([Gentzkow, Shapiro, and Sinkinson, 2014](#)) shows that 15 percent of American consumers who reported reading a newspaper read two or more. Second, multi-homing has important and interesting implications for the advertiser market. If newspapers do not share the same consumers because everyone buys, at most, one product, we do not really observe competition for advertisers by newspapers. We have the “competitive bottleneck” of [Armstrong \(2006\)](#). However, when some consumers are available on both newspapers, an advertiser considering an ad with a particular newspaper will know that even if it does not choose to advertise there, it can still reach some its readers by advertising on another newspaper. Thus,

multi-homing consumers loosen the grip of newspapers on advertisers.

In Chapter 3, published in *European State Aid Quarterly*, Malgorzata Cyndecka and I describe and apply the findings of Chapter 1 and associated literature to a decision made by the European Free Trade Area Surveillance Authority (ESA). The decision permitted a temporary reduction of VAT on digital newspapers in Norway.

We confirm that even a tax reduction which equalises tax conditions between printed and digital newspapers qualifies as State Aid. However, the Norwegian authorities argued that the State aid should be allowed as the tax reduction would reduce prices and increase consumption. This finding goes against the findings of Chapter 1, and as such we argue that the ESA did not have grounds to accept the aid as being compatible with the single market.

In Chapter 4 I live up to the earlier promise to combine two preferences of mine: Harold Hotelling and alcohol. In this chapter Oddmund Berg and I took inspiration from the standard models of spatial competition to consider a recent expansion of the state owned monopolist of high strength alcohol, Vinmonopolet. Models with disutility from travel costs predict that if a region receives a new store, then travel distances should reduce and consumption should increase. As the great philosopher Homer (Simpson) once said, alcohol is the cause of (and solution to) all of life's problems. Norway's highly interventionist approach to restrict alcohol availability and consumption recognises the downside of alcohol consumption. Would any increased consumption cause wider costs to society, such as higher sick leave?

We are not aware of any previous papers that have been able to attribute a causal relationship between alcohol consumption and recorded sick leave. Previous correlations studies might overestimate the causal relationship (Norström, 2006; Norström and Moan, 2009). Prior to being signed off with stress, for example, employees might self medicate with alcohol. Furthermore, being on sick leave might lead to feelings of isolation and depression which might also lead to increased consumption of alcohol.

We find that a reduction in the average driving distance increases alcohol sales. We then use driving distance to examine the causal relationship between alcohol consumption and recorded sick leave. Our main source of variation comes from the opening of new stores reducing driving distances in some areas in some periods. Our core identifying assumption is that the only way reduced driving distances to new stores affects sick leave is through alcohol consumption.

We find that an increase in alcohol consumption of 1 percent leads to an increase of sick leave in men of around 0.3 per 10,000 men, an increase of around 0.16 percent, at the mean.

In Chapter 5, I return to digital newspapers. With Frode Skjeret and Frode Steen, I investigate the impact of paywalls on consumption of digital news. All companies have a difficult decision when deciding on the price, or prices for a product. But newspapers, particularly digital newspapers have an even more fundamental question. Should they charge readers a price at all?

A key factor will be how many views, and therefore how much advertising revenue, will the newspaper lose? In this chapter we estimate the short and long run reductions in consumption following the introductions of paywalls. We use weekly consumption data for 4 years from 122 Norwegian news producers, of which 69 introduce paywalls.

We find that short run consumption of news falls by 3-4 percent, and continues to fall to between 9-11 percent in the long run. Larger media outlets tend to be punished more by readers with regional market leaders seeing falls of 13-15 percent, compared to 8-11 percent for the others.

Economics has allowed me to ask and answer many interesting questions. Now, my role is to show the joy of economics to others and, hopefully, to inspire them to ask and answer some questions of their own.

References

- ARMSTRONG, M. (2006): "Competition in two-sided markets," *The RAND Journal of Economics*, 37(3), 668–691.
- BÖCKEM, S. (1994): "A generalized model of horizontal product differentiation," *The Journal of Industrial Economics*, pp. 287–298.
- D'ASPROMONT, C., J. J. GABSZEWICZ, AND J.-F. THISSE (1979): "On Hotelling's "Stability in competition",," *Econometrica: Journal of the Econometric Society*, pp. 1145–1150.
- 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.
- HOTELLING, H. (1929): "Stability in Competition," *The Economic Journal*, 39(153), 41–57.
- LANCASTER, K. J. (1966): "A new approach to consumer theory," *Journal of Political Economy*, 74(2), 132–157.
- NORSTRÖM, T. (2006): "Per capita alcohol consumption and sickness absence," *Addiction*, 101(10), 1421–1427.
- NORSTRÖM, T., AND I. S. MOAN (2009): "Per capita alcohol consumption and sickness absence in Norway," *The European Journal of Public Health*, 19(4), 383–388.
- PIDD, K. J., J. G. BERRY, A. M. ROCHE, AND J. E. HARRISON (2006): "Estimating the cost of alcohol-related absenteeism in the Australian workforce: the importance of consumption patterns," *Med J Aust*, 185(11-12), 637–41.

CHAPTER II

Tax-free digital news?

Tax-free digital news?

Øystein Foros

NHH Norwegian School of Economics

oystein.foros@nhh.no

Hans Jarle Kind

NHH Norwegian School of Economics and CESifo

hans.kind@nhh.no

Tim Wyndham

NHH Norwegian School of Economics

timothy.wyndham@nhh.no

March 2018

Abstract: News platforms are struggling. Their printed readership is predominantly old, and their digital product struggles to win the attention of the young. For several decades tax reductions have been used in Europe to increase the circulation of printed newspapers. Would extending these reductions to digital platforms stimulate digital consumption? Using a two-sided pricing model where a print platform and a digital platform compete for multi-homing consumers and advertisers we show that the answer is no. The two-sidedness of the market means that the digital price would increase. Not only would digital circulation decrease but so too would the fraction of consumers that access news from both platforms. Key media policy goals of reach (circulation) and pluralism (multi-homing) would be harmed.

Keywords: two-sided markets, multi-homing, digital media, incremental pricing, value added tax.

JEL Classification: D11, D21, L13, L82.

1 Introduction

For several hundred years newspapers have operated in a two-sided market selling news and commentary to readers and eyeballs to advertisers. The United Kingdom’s first regular newspaper, The Daily Courant, was launched in 1702 and consisted of a single page of news with adverts on the back (Williams, 2009). For most of the time afterwards, and certainly since the early 19th century, newspapers have been seen as important for the health of democracy (Gentzkow et al., 2011).

Preferential tax treatment of printed newspapers is widespread, particularly in Europe as illustrated in Table 1.¹ In the UK and Norway, printed newspapers do not pay value added tax (VAT) on sales to readers.^{2,3} The preferential treatment aims to increase circulation and ensure that people get information from several different sources (multi-homing in the two-sided market terminology).

	Standard rate	Printed Newspapers	e-Newspapers
Austria	20%	10%	20%
Belgium	21%	6%	21%
France	20%	2.1%	2.1%
Germany	19%	7%	19%
Spain	21%	4%	21%
Norway	25%	0%	25% (pre 1/3/2016) 0%
UK	20%	0%	20%

Table 1: *VAT Rates in selected European countries. Source: European Commission (2016c) and Statsministerens Kontor (2015).*

The circulation of printed newspapers is shrinking rapidly, leading to questions over their future. For example, the reach of national printed newspapers in the UK decreased

¹In the US "Federal, state, and local governments have traditionally provided a variety of special economic supports to the industry, including exemptions from newspaper and advertising sales taxes and excise taxes on telecommunications equipment used for information gathering" (Picard, 2004).

²Value added refers to "the value that a producer... ..adds to his raw material or purchases (other than labor) before selling the new or improved product or service" and can be calculated as the sum of wages and profits (Tait, 1988). Developed in France in the middle of the 20th Century VAT had become widespread by the turn of the millennium, particularly in Europe. For more on the theory, history and practical details of VAT, see Tait (1988).

³Since 1 March 2016 digital newspapers in Norway also do not pay VAT.

by more than a quarter between 2005 and 2015 (Ofcom, 2015). Circulation in Norway also fell by more than a quarter between 1999 and 2013.⁴ In the UK The Independent has already ended its print edition, believing that it "will be the first of many leading newspapers to embrace a wholly digital future" (Lebvedev, 2016).

A key driver of the fall in the circulation of printed newspapers is the news consumption habits of the young. The young have always been less likely to buy a newspaper than the old, but readership of newspapers has fallen faster for the young. In 2005, national printed newspapers in the UK reached around 75 percent of those aged 65 or above and slightly under 70 percent of those aged between 15 and 34. By 2015, reach had fallen to around two thirds and one third respectively (Ofcom, 2015). The current situation in Norway is even more stark: In 2014, 82 percent of 67-79 year olds read a newspaper on an average day compared to 26 percent of those aged 16-24 (EFTA Surveillance Authority, 2016).

Given the frequent use of VAT policy to stimulate consumption of printed news and the increasing proportion of digital content that is placed behind a paywall, it is natural for policy makers to ask whether the VAT exemption should be extended to digital news. Member States of the European Union are prohibited by the VAT Directive from applying a beneficial rate to digital news.⁵ However, the European Commission and the European Council are seeking to amend the VAT Directive and allow Member States the option to offer e-publications in each country the same tax rate as their printed counterparts (See European Commission 2016a, 2016b, 2016c, European Council 2017). Norway, a member of the European Economic Area (EEA)⁶ but not of the European Union, is not constrained by the VAT Directive and was the first within the EEA to implement a zero tax regime for digital as well as printed news.⁷

In the State aid approval of the Norwegian zero-tax regime towards digital news, EFTA Surveillance Authority (2016) noted (page 13): "*The main objective of the proposed zero VAT rate is to support the demand and use of news and current affairs content among con-*

⁴Calculation by authors based on data presented in Statministerens Kontor (2015).

⁵This restriction has not prevented some Member States from reducing VAT on electronic newspapers. As can be seen in Table 1 France uses a lower rate. Belgium, Luxembourg and Italy have all offered lower than standard VAT rates on electronic newspapers at some point.

⁶The EEA comprises the Member States of the European Union and three non members: Iceland, Liechtenstein, and Norway.

⁷In Iceland newspapers are taxed at a reduced rate of 11%, whilst the rate in Liechtenstein is 2.5%.

sumers, thereby also promoting media pluralism and diversity. This requires instruments aimed at consumers. Reducing the cost of electronic news services is a direct and effective means of ensuring high news consumption and thereby a broad and enlightened public discourse." Furthermore, the following concern of the Norwegian government was accentuated (EFTA Surveillance Authority, 2016, page 12): *"the existing zero VAT rate for newspapers primarily supports the media consumption of the middle-aged or elderly."*

Against this backdrop, our research question is:

- Does reducing VAT on digital news reduce the price of digital news and stimulate digital news consumption?

Surprisingly, the answer is no. If the VAT exemption is extended to digital platforms then the price towards digital readers increases. Demand for the digital platform decreases, as does the fraction of consumers getting information from different sources (multi-homing).

The departure from the standard intuition is driven by the presence of the ad market. A lower tax rate on the digital product increases the profitability of the reader market for the platform, but has no direct effect on the advertising market. This means that the digital platform will place more weight on reader market profits and less weight on advertising market profits; it becomes more important to set a relatively high subscription price and boost reader margins despite the consequent reductions in readership and ad revenues.

Naturally, the profit of the digital platform increases following a reduction of VAT on digital subscriptions. But, surprisingly, the profit of the printed platform also increases. So, tax-free digital news may help the survival of print newspapers, albeit at the expense of reduced online readership.

Our model is based upon Hotelling (1929); the dominant workhorse for analysing two-sided media markets (the seminal contribution is Anderson and Coate, 2005). We allow for asymmetric platforms with respect to tax rates and marginal costs. The digital platform has lower marginal costs and, at the outset, a higher (standard) tax rate than the print platform. Our model incorporates two-sided pricing; both platforms sell subscriptions to consumers and eyeballs to advertisers. We start by recognising the reality that consumers have long bought more than one newspaper.⁸ Thus we allow multi-homing by both consumers and

⁸American survey data from 1917-1919 showed that 15 percent of households who reported reading a daily newspaper reported reading two or more (Gentzkow et al., 2014), but the digitisation of news has increased the prevalence of multi-homing (see Athey et al., forthcoming and Peitz and Reisinger, 2015).

advertisers.

To our knowledge, we are the first to assess the reaction to tax changes by two-sided duopolists facing multi-homing consumers (consumers that are "shared" between each platform). As well as being a significant measure of media pluralism, multi-homing by consumers has important implications for competition between platforms. When all consumers buy only one product (they each single-home or are "exclusive" to a platform), the "competitive bottleneck" problem of Armstrong (2002, 2006) exists: since an exclusive reader can only be accessed through the platform from which she purchases there is no direct competition for advertisers among platforms. Recent contributions by Athey et al. (forthcoming), Ambrus et al. (2016) and Anderson et al. (2018) introduce competition for advertisers by allowing consumers to multi-home.

We combine ingredients from Anderson et al. (2018), which considers multi-homing consumers in a pure ad-financed two-sided market, and Anderson et al. (2017), which considers multi-homing consumers in a one-sided user financed market. These ingredients are used to extend the simple single-homing model with dual source financing (two-sided pricing) of Anderson and Gabszewicz (2006). Despite the many components we construct a simple model to highlight the core mechanism driving our results: the two-sidedness of the market.

There are papers assessing the impact of VAT on price in two-sided markets with a monopolist platform (Kind et al., 2008), with duopolists (Kind et al., 2013) and with perfect competition (Kind et al., 2008). An important limitation of these models is that, even when there is more than one platform, consumers are assumed to single-home. Platforms are also assumed to have symmetric cost structures. We relax both these assumptions.

We also extend our model to the pure single-homing consumer case, for two reasons. First, we show that this might be the outcome of a VAT reduction in our model. It is possible that reducing VAT on digital news could increase the digital price to the extent that no consumer buys both products. Second, previous investigations of the reaction to tax changes by two-sided duopolists facing single-homing consumers have been location games. In those papers reduced tax rates have increased reader prices through increased horizontal differentiation. In our pure single-homing model we find the same inverse relationship between tax rates and prices without any change in horizontal differentiation.

In Sections 2 and 3 we present the foundations and findings of our model with shared

consumers (multi-homing). In Section 4 we investigate the case when there are only exclusive consumers (single-homing). In Section 5 we discuss the robustness of our main results. We summarise and discuss our results in Section 6.

2 The model

Consider two competing media platforms; one producing a printed newspaper, and the other a digital newspaper. The digital platform (D) has marginal costs equal to $c_D \geq 0$, while the print platform (P) faces a marginal cost equal to $c_P > 0$. Throughout we assume $c_P > c_D$.⁹ In the basic model we set $c_D = 0$, while in Section 5 we allow for $c_D > 0$. We abstract from any fixed costs. The tax rate (VAT) on subscriptions for each platform is τ_i , where $i = D, P$. The platforms are located at the extremes of a “Hotelling line” with length 1. Platform D is at the far left and platform P at the far right. Platforms sell subscriptions at price p_i to consumers and eyeballs to advertisers. This implies that the consumers pay p_i and platforms receive $\frac{p_i}{1+\tau_i}$. We specify ad prices below.

Consumer (reader) tastes are uniformly distributed along the line. We may interpret the horizontal differentiation as age. Young people to the left, old people to the right. Consistent with empirical studies of the US newspaper market (Gentzkow, 2007, Fan, 2013, and Gentzkow et al., 2014) we assume that consumers are ad-neutral.¹⁰ In Section 5 we consider the outcome when consumers dislike ads. The distance disutility (transportation costs) is given by t .

Remark (one-sided market): Marginal costs are approximately equal to zero for digital goods (e.g. e-books). It is well known from the tax literature that such a cost structure implies that VAT acts as a pure surplus tax with no impact on consumer prices in one-sided markets. To see this, consider the zero marginal cost profit function $\pi = \frac{p}{1+\tau}x(p)$, where $x(p)$ is the demand function. The tax rate τ clearly drops out of the first-order condition $\partial\pi/\partial p = 0$, so that it only affects the ‘profit-split’ between the firm and the government. This insight provides us with a clear benchmark in the two-sided markets we analyse.

⁹We also use D and P to refer to the digital and printed product, respectively.

¹⁰See Chandra and Kaiser (2015) for a comprehensive survey of the literature on consumers’ attitude towards ads in newspaper markets.

2.1 Consumer demand

Let the consumer utility of buying from only platform D or only platform P be

$$u_D = v_D - tx - p_D \text{ and} \quad (1)$$

$$u_P = v_P - t(1 - x) - p_P, \quad (2)$$

respectively, where v_i is the vertical quality of platform i , p_i is the subscription fee and x is the location of the consumer.

The utility of a consumer who buys both products is the sum of individual utilities less any utility loss due to overlap:¹¹

$$u_{(D+P)} = u_D + u_P - d. \quad (3)$$

We follow the essence of Anderson et al. (2017) and say that consumers with $u_D > u_P$ perceive D as their primary good and P as their secondary good. A consumer will buy both products if the incremental utility of multi-homing is positive,¹² where her incremental utility of multi-homing is specified as $u_{(D+P)} - \max[u_D, u_P]$. To find the location of the consumer indifferent between buying only D and buying both D and P we set $u_{(D+P)} - u_D = 0$ and solve for x , yielding

$$x_{DP} = 1 - \frac{v_P - p_P - d}{t} \quad (4)$$

where we implicitly assume that $t > v_P - p_P - d > 0$.

Likewise the location of the consumer who is indifferent between buying only P and buying both P and D is

$$x_{PD} = \frac{v_D - p_D - d}{t}. \quad (5)$$

Figure 1 identifies those consumers who buy only D , those who buy both D and P and those who buy only P .

¹¹We are grateful to Paul Belleflamme and the Editor for specifying this formulation. Our qualitative findings hold for other formulations, including when the utility "loss" from the incremental purchase is proportional to the sum of vertical qualities and when the "loss" is proportional to the vertical quality of the secondary product (as in Anderson et al., 2017).

¹²In our model consumers who buy both products are those in the middle of the Hotelling line, their distaste for either publication is not "too" strong.




Figure 1: *Shared and exclusive readers.*

Under multi-homing, consumer demand at each platform is

$$X_D^{MH} = \underbrace{x_{DP}}_{\text{D's exclusive readers}} + \underbrace{(x_{PD} - x_{DP})}_{\text{shared readers}} = x_{PD} \quad (6)$$

$$X_P^{MH} = \underbrace{1 - x_{PD}}_{\text{P's exclusive readers}} + \underbrace{(x_{PD} - x_{DP})}_{\text{shared readers}} = 1 - x_{DP}. \quad (7)$$

Demand for each newspaper is strictly decreasing in own price ($\partial X_i^{MH} / \partial p_i = -1/t$).¹³ The number of exclusive readers for each newspaper is, however, independent of the price it charges. This is because we have assumed that all consumers read at least one newspaper. A newspaper's number of exclusive readers is therefore determined by, and more precisely, is inversely related to, the demand for the other newspaper. Since an increase in p_i does not affect the incremental value of newspaper j (leaving demand for that newspaper unchanged), it cannot affect newspaper i 's number of exclusive readers either. This partly reflects the peculiarities of the Hotelling model and our specific assumptions, but does not qualitatively affect the results we derive below (see Section 5.1 for a discussion).

2.2 Platforms and advertisers

Both platforms can costlessly place commercials in their newspaper. As in Anderson et al. (2018) we assume that platforms set a price per ad, and that advertisers only place one advert per platform. We assume a perfectly elastic demand curve for ads, with a mass A of homogenous advertisers.

The expected value for an advertiser of reaching a reader who sees the ad on one and only one of the platforms is equal to β . This value reflects the advertiser's profit margin

¹³For instance, a higher price on the digital newspaper moves x_{PD} to the left (which implies that a larger share of the consumers will read only the printed newspaper).

and the proportion of consumers who, on seeing the ad for the first time, buy the product. As in Anderson and Coate (2005) the platforms are able to extract all the advertisers' surplus from exclusive consumers. We allow for the second impression to be worth less than the first. Thus the expected value to an advertiser of a consumer seeing his advert twice is $\beta(1 + \sigma)$, with $\sigma \in (0, 1]$.

To specify the equilibrium price per advert we use the principle of incremental pricing as developed in Anderson et al. (2018), i.e. prices at a given platform will be determined by the incremental value to an advertiser of advertising on that platform. This prevents either platform from charging more than σb for its shared consumers. The price per ad at platform i will therefore be $a_i = \beta X_i^e + \sigma \beta X^s$, where X_i^e is platform i 's exclusive consumers and X^s represents the number of consumers that purchase both products.¹⁴ Total ad revenues at platform i will be $A(\beta X_i^e + \sigma \beta X^s)$ or $bX_i^e + \sigma bX^s$, using the identity $b \equiv A\beta$.

Table 2 compares the values, incremental values and prices for a platform's single-homing and multi-homing consumers.

	Exclusive consumer	Shared consumer
Value to advertiser	β	$\beta(1+\sigma)$
Incremental value to advertiser	β	$\sigma\beta$
Incremental price	β	$\sigma\beta$
Advertiser Surplus	0	$\beta(1+\sigma)-2\sigma\beta= \beta(1-\sigma)$

Table 2: *Advertising values for exclusive and shared consumers.*

Our model set-up allows for multi-homing consumers, asymmetric platforms with respect to marginal costs and tax rates, as well as two-sided pricing (platforms charge both consumers and advertisers). Given this complex set up, for the sake of simplicity, we search for Nash equilibria in a setting where platforms set prices for readers and for ads simultaneously.

¹⁴To fix ideas consider a numerical example. Let the value of a first impression (β) be 0.9 and the value of a second impression ($\sigma\beta$) be 0.36. The mass of consumers is normalised to 1. Furthermore, let 20% of consumers buy exclusively from the printed firm, 30% exclusively from the digital firm and 50% from both. Then the per advert prices under incremental pricing will be $a_D = 0.9 * 0.3 + 0.36 * 0.5 = 0.45$ and $a_P = 0.9 * 0.2 + 0.36 * 0.5 = 0.36$. A subscription fee decrease at the printed firm that led to 60% of consumers buying from both firms and the remaining 40% exclusive consumers split equally would lead to the symmetric per advert prices: $a_D = a_P = 0.9 * 0.2 + 0.36 * 0.6 = 0.396$.

3 Multi-homing consumers

When some consumers buy both products, the platforms face the following profit functions:

$$\pi_D = \frac{p_D}{1 + \tau_D} X_D^{MH} + bX_D^e + \sigma bX^s \text{ and} \quad (8)$$

$$\pi_P = \left(\frac{p_P}{1 + \tau_P} - c_P \right) X_P^{MH} + bX_P^e + \sigma bX^s, \quad (9)$$

where the first terms represent reader market profit and the second and third terms represent ad market profits. Note that in a one-sided market ($b = 0$), τ_D would act as a pure surplus tax since marginal costs on the digital newspaper are zero (see Remark above).

Using (6) we can rewrite (8) to separate profit from exclusive and shared consumers


$$\pi_D = \left(\frac{p_D}{1 + \tau_D} + b \right) X_D^e + \left(\frac{p_D}{1 + \tau_D} + \sigma b \right) X^s,$$

and write the first order condition for the digital platform as

$$\frac{d\pi_D}{dp_D} = \left[\frac{X_D^e + X^s}{1 + \tau_D} + \frac{p_D}{1 + \tau_D} \frac{\partial X^s}{\partial p_D} \right] + \sigma b \frac{\partial X^s}{\partial p_D} = 0 \quad (10)$$

Raising p_D has the standard effect on reader market profitability. It increases the profit margin but reduces sales. If there were no ads in D (in which case we would have a one-sided product), profit maximization dictates that the term in the square bracket of (10) should be set to zero (marginal revenue equal to marginal cost, which is zero). However, the term outside the bracket is negative when second impressions have a positive value, showing that D 's optimal two-sided price is lower than its optimal one-sided price.¹⁵ This is due to the fact that the increased advertising profits gained from selling eyeballs (alongside the additional reader sales) exceed the foregone margin on existing consumers.

Note that D 's marginal consumers are the ones it shares with P . While a small price rise will reduce the surplus of infra-marginal customers (including exclusive consumers) the only consumers that will stop buying D are those buying it as a secondary product, in other words for its incremental value, as Figure 2 shows. Mathematically, raising p_D reduces x_{PD} but has no effect on x_{DP} .



¹⁵Since $\frac{\partial X^s}{\partial p_D} < 0$ and $\sigma b > 0$. To see the former substitute (4) and (5) into $X^s = x_{PD} - x_{DP}$.

Figure 2: *Effect of a digital price rise on reader demands.*

A lower tax rate on the digital product increases the profitability of the reader market for the platform, but has no direct effect on the advertising market. This means that the digital platform will place more weight on the term in the square bracket of (10) compared to the term outside; it becomes more important to set a relatively high subscription price and boost reader margins despite the consequent reductions in readership and ad revenues. In contrast to typical results in one-sided markets, we might therefore expect the consumer price to be decreasing in the tax rate, other things being equal. This is confirmed by solving (10) to find the digital platform's reaction function

$$p_D(\circ) = \frac{v_D - d - \sigma b(1 + \tau_D)}{2} \quad (11)$$

from which we immediately see that $\frac{\partial p_D(\circ)}{\partial \tau_D} = -\frac{\sigma b}{2}$. Furthermore, we observe that the size of the price change depends on the (incremental) value of the shared consumers on the advertising market. The reason for this is, as noted above, that the platform can only affect the number of shared readers - and not the number of exclusive readers - through its pricing behaviour.

Note also that $p_D(\circ)$ is independent of p_P . The intuition is that a consumer who is considering purchasing D as a secondary product, will only consider the price of D . Prices are thus strategically independent. See Anderson et al. (2017) for a further discussion of this issue.

The reaction function of the printed platform is qualitatively similar:

$$\frac{d\pi_P}{dp_P} = \left[\frac{X_P^e + X^s}{1 + \tau_P} + \left(\frac{p_P}{1 + \tau_P} - c_P \right) \frac{\partial X^s}{\partial p_P} \right] + \sigma b \frac{\partial X^s}{\partial p_P} = 0$$

yielding the best response function

$$p_P(\circ) = \frac{v_P - d - (b\sigma - c_P)(1 + \tau_P)}{2}. \quad (12)$$

This shows that the subscription price of platform P is decreasing in its own tax rate if $\sigma b > c_P$ or when the value on the ad market of reaching a multi-homing consumer exceeds the marginal cost of producing an extra copy.

Since prices are strategically independent, reaction functions (11) and (12) are also equilibrium values.¹⁶ A further important implication of price independence is that the tax rate on one platform has no effect on the price of the other platform.

¹⁶The second-order condition is $\frac{d^2 \pi_i}{dp_i^2} = -\frac{2}{t(1+\tau_i)} < 0$.

It is noteworthy that for sufficiently valuable ad markets, equilibrium subscription prices at either or both platforms could in principle be negative. We restrict our attention to cases where prices are positive.¹⁷ An important reason to abstract from negative prices is that pure negative prices are rarely observed in practice, although often there may be complimentary gifts or other exclusive offers for subscribers. It is possible that the platforms would prefer to have negative prices irrespective of the VAT rate they face but are unable to feasibly implement this. In this situation the price would remain stable at zero.

Summing up the above analysis we can state:

Proposition 1: *The price of the print platform is independent of the tax rate on the digital platform, and vice versa. Suppose that there is a decrease in the tax rate on*

a) the digital platform. Then own price will increase if second impressions have any incremental value ($\sigma b > 0$).

b) the print platform. Then own price will increase if the incremental value of second impressions is worth more per consumer than the marginal cost of printing an extra copy ($\sigma b - c_P > 0$).

This proposition could also be worded that own prices decrease in own tax rates as long as the advertising value to the platform of a shared consumer exceeds that platform's marginal cost.

One might expect that due to price independence, the tax rate on one platform does not affect the profits of the other. Interestingly, this is not true. Suppose that τ_D increases. This will not affect the price (p_P) or total demand ($X_P^{MH} = 1 - x_{DP}$) for the print platform, but will affect its composition of exclusive ($X_P^e = 1 - x_{PD}$) and shared readers ($X^s = x_{PD} - x_{DP}$). Inserting for (4), (5) and (12) into (8) and differentiating π_P with respect to τ_D yields

$$\frac{d\pi_P}{d\tau_D} = b \frac{dX_P^e}{d\tau_D} + \sigma b \frac{dX^s}{d\tau_D}.$$

The digital platform will charge a higher subscription price if its tax rate, τ_D , decreases. As Figure 2 illustrates, the total demand for P is unaffected, but some consumers who previously bought both products will now only buy P . The increase in the digital price has "converted" some of P 's shared consumers into exclusive consumers.¹⁸ This conversion will

¹⁷Specifically we assume $v_D > d + \sigma b(1 + \tau_D)$ and $v_P > d + (b\sigma - c_P)(1 + \tau_P)$.

¹⁸The increase in exclusive consumers is $\frac{dX_P^e}{d\tau_D} = \frac{\sigma b}{2t}$.

not affect the print platform's reader market profit, as its reader price and total demand are unchanged but its advertising market profit will increase. Exclusive consumers are worth more on the ad market than shared consumers so the print platform will increase the price of its ads, a_P . We consequently find that $\frac{d\pi_P}{d\tau_D} = -\frac{\sigma b}{2t}(b - \sigma b) = -\frac{b^2\sigma(1-\sigma)}{2t} < 0$. A lower digital tax rate increases the printed platform's profit.

For the digital platform we likewise find $\frac{d\pi_D}{d\tau_P} = -\frac{b(b\sigma - c_P)(1-\sigma)}{2t}$ which is negative if $b\sigma > c_P$. Under this condition a lower tax rate τ_P increases p_P and we have the same mechanism. We can state:

Proposition 2: *The print platform's profit decreases in the tax rate of its rival ($\frac{d\pi_P^{MH}}{d\tau_D} < 0$). The digital platform's profit decreases in the rival's tax rate if $b\sigma > c_P$.*

From a media pluralism perspective, a major rationale for preferential tax treatment of newspapers has been to increase their circulation and to ensure that people get information from several different sources (multi-homing, in our terminology). This may be important for e.g. democratic processes, knowledge spillovers and anti-bias measures. We will not go into these rationales, but note that the number of multi-homers is equal to

$$X^s = x_{PD} - x_{DP} = \frac{(v_P + v_D) + b\sigma(1 + \tau_D) - (c_P - b\sigma)(1 + \tau_P) - 2d}{2t} - 1, \quad (13)$$

from which it immediately follows that:

Proposition 3: *Reducing the tax rate on the digital platform (τ_D) decreases the number of multi-homing consumers. Reducing the tax rate on the print platform (τ_P) decreases the number of multi-homers if the incremental value of a multi-homing consumer is larger than the print platform's marginal cost ($b\sigma > c_P$).*

We also observe from (13) that the comparative statics of the number of shared readers are intuitively reasonable. The number of shared readers is increasing in the value of second impressions ($b\sigma$) and is decreasing in the strength of horizontal preferences (t), the amount of overlap (d) and the printed platform's marginal cost (c_P).

Figure 3 shows a numerical example where we set $\tau_P = 0$ and vary τ_D .¹⁹ With a tax rate of 25 percent, as Norway used to have, 2.0 percent of readers are shared. A tax rate of 20 percent as in the UK, implies 1.3 percent of readers are shared. Reducing the tax

¹⁹The other parameters are $v_D = v_P = 0.9$, $c_P = 0.3$, $t = 0.55$, $b = 0.55$, $\sigma = 0.26$, $d = 0.35$.

rate for the digital platform below 10 percent implies that there would only be exclusive consumers. In the absence of shared consumers, the nature of competition between the two platforms changes significantly. We investigate this in Section 4.




Figure 3: *Impact of the digital tax rate on the number of shared consumers.*

The logic of two-sided markets, as described above, clearly indicates that subsidising newspapers through reduced value-added taxes might be an ineffective or even counter-productive means to increase newspaper circulation.

Before we proceed to a single-homing environment, we note a more positive insight from the analysis above. A public policy which contributes to higher media quality (an increase in v_i) could be an effective way to increase multi-homing (despite higher newspaper prices) as well as being a political goal in its own right. More precisely, from equations (11), (12) and (13) we observe:

Proposition 4: *Assume that media quality improves (v_D and v_P increase). Then subscription prices and the extent of multi-homing increase.*

Rather than lowering the tax rate on digital platforms (which would lead to a higher digital price but lower circulation), governments could for instance subsidize journalism to ensure both higher media quality and larger newspaper circulation.

4 Single-homing consumers

In the previous section we observed the possibility that no consumers multi-home (e.g. due to low VAT rates, as illustrated in Figure 3). Furthermore, in the introduction we noted that the previous literature on tax in two-sided markets has followed the Hotelling convention of assuming that each consumer buys a maximum of one product. We now extend the "pure single-homing" literature by assessing asymmetric platforms.

Suppose that the market is shared and each consumer buys one and only one of the media products.²⁰ Consumer demands resemble the standard Hotelling set-up: $X_i^{SH} = \frac{1}{2} + \frac{v_i - v_j}{2t} - \frac{p_i - p_j}{2t}$.²¹ Using a similar methodology to that used in Section 3 we can derive the digital platform's reaction function:

$$p_D(\circ) = \frac{t + v_D - v_P - b(1 + \tau_D)}{2} + \frac{p_P}{2}. \quad (14)$$

Equation (14) shows that D 's reaction function shifts up if its tax rate is reduced. As in the multi-homing case it is optimal for D to shift profit from the advertising side to the consumer side by increasing the reader price. As in the multi-homing case, the price increase will be greater the greater the per-reader advertising revenue.²²

Following the same process for the printed platform yields the best response function

$$p_P(\circ) = \frac{t + v_P - v_D - (b - c_P)(1 + \tau_P)}{2} + \frac{p_D}{2}. \quad (15)$$

Also the print platform will respond to a low tax rate on its reader revenues with high reader prices if the value on the advertising market of an extra reader is greater than its marginal cost.

From the response functions we note that prices are strategic complements, so that they tend to move in tandem in response to changes in exogenous variables (e.g. in tax rates). Combining (14) and (15) we find the equilibrium prices:

$$p_D^{SH*} = t + \frac{(v_D - v_P) - 2b(1 + \tau_D) - (1 + \tau_P)(b - c_P)}{3} \text{ and} \quad (16)$$

$$p_P^{SH*} = t + \frac{(v_P - v_D) - 2(1 + \tau_P)(b - c_P) - b(1 + \tau_D)}{3}. \quad (17)$$

²⁰Consumers will choose D or P to maximise (1) or (2). The market is covered iff $t \leq \frac{1}{3}(v_D + v_P + 2b + b(\tau_D + \tau_P) - c_P(1 + \tau_P))$. The market sharing condition is $t > \max\left[\frac{1}{3}(v_D - v_P + b(\tau_D - \tau_P) + c_P(1 + \tau_P)), \frac{1}{3}(v_P - v_D - b(\tau_D - \tau_P) - c_P(1 + \tau_P))\right]$.

²¹See the Appendix for full details.

²²For clarity, note that under pure single-homing all readers are exclusive and so all readers have the same value on the ad market.

Summing up:

Proposition 5: *Single-homing. Suppose that there is a reduction in the tax rate on a) the digital platform. Then both platforms will increase consumer prices.*

b) the print platform. Then both platforms will increase consumer prices if $b > c_P$. Both will decrease consumer prices if $b < c_P$.

It is straight forward to see that profit is strictly decreasing in own tax rate (see Appendix). Interestingly, the multi-homing result in Proposition 2, that even a completely non-altruistic newspaper might find it optimal to lobby for a reduction of the tax rate paid by its rival, survives also under single-homing. In the multi-homing case this was due to competition in the advertising market, while it is due to competition in the reader market under single-homing. More precisely, if platform i responds to a tax reduction by increasing its price under single-homing, platform j will capture a larger number of readers (and charge a higher price, since prices are strategic complements) and thus make higher profits from both reader and advertiser markets.

We can state:

Proposition 6: *Single-homing. Profit is decreasing in own tax rate. The profit of the print platform is, moreover, decreasing in the tax rate of its digital rival. The profit of the digital platform is decreasing in the print platform's tax rate if both prices decrease in that tax rate.*

There have been some concerns that differences in tax rates between print and digital platforms have led to artificial differences in circulation. This might be correct, but perhaps not in the generally perceived direction. Equations (16) and (17) show that reducing the digital tax rate leads to higher prices at each platform but that the price increase is larger at the digital platform. The relative price increase has the intuitive effect of increasing printed sales at the expense of digital sales:²³

Proposition 7: *Single-homing. Reducing the tax rate on the digital platform (τ_D) will increase sales of printed newspapers and reduce sales of digital newspapers.*

²³Full detail in the Appendix. As the profitability of digital newspapers increases with a digital tax reduction, it is possible that some of the reduction in digital circulation would be offset by new entry of digital newspapers. This is not considered in our model.

5 Robustness

5.1 Uncovered markets

For simplicity, we have chosen a framework such that the number of exclusive readers in the multi-homing case is independent of own price (c.f. the discussion below Figure 1). In a more general model, where the supply of exclusive readers is elastic, a price decrease at firm i could win it some exclusive consumers in addition to converting some of j 's exclusive consumers into shared ones. Our qualitative result of a negative relationship between VAT rates and reader prices would be unaffected by such a change. The core requirement for this finding is that a platform's ad revenues are increasing in readership; whether readership increases come from exclusive or shared consumers does not matter. If this positive indirect network externality is present, then a firm can respond to a VAT increase on the reader market by reducing the subscription fee, increasing readership and increasing ad market profits.

5.2 Disutility of ads

Above we assumed that consumer utility is unaffected by the volume of ads. Now we allow for consumer disutility of ads; platforms need to weigh the ad market benefits of an additional advert against the negative reader market impact of a less attractive product. We normalise the mass of advertisers to 1, and platform i chooses an ad level $A_i \in [0, 1]$.

We provide full details of the model in the Appendix. The key difference is that now the utility of buying only from platform i for a consumer located at point x is given by $u_i = v_i - t|x - x_i| - p_i - \gamma A_i^2$, where $x_D = 0$, $x_P = 1$, and $\gamma > 0$. The utility of buying from both platforms is still $u_{(D+P)} = u_D + u_P - b$ and the incremental utility of multi-homing is specified as $u_{(D+P)} - \max[u_D, u_P]$. We start with the multi-homing model and follow the same methodology as in Section 3.²⁴ We identify the following equilibrium outcomes at the

²⁴We focus here on the case when $\sigma = 1$ to achieve tractable solutions. In the Appendix we use numerical methods to relax this assumption. Our principle finding (Proposition 1) is maintained, although our findings on the relation between digital tax rates and the extent of multi-homing and printed profitability are more ambiguous. Full results and intuition are in the Appendix.

digital firm

$$p_D = \frac{v_D - d}{2} - \frac{3b^2(1 + \tau_D)^2}{8\gamma} \text{ and } A_D = \frac{b(1 + \tau_D)}{2\gamma}.^{25}$$

The ad level is increasing in the value of the ad market and decreasing in consumer's distaste for ads ($dA_D/d\gamma < 0$).²⁶ The reader price is decreasing in the value of the ad market ($dp_D/db < 0$) and increasing in consumer's distaste for ads ($dp_D/d\gamma > 0$). The intuition for the latter is that since higher disutility of ads reduces the ad volume, the willingness to pay for the newspaper increases. More interesting for our point of view, is the fact that we still have that the reader price is decreasing in own VAT rate ($dp_D/d\tau_D < 0$).

For the printed platform we find

$$p_P = \frac{v_P - d}{2} - \frac{3b^2(1 + \tau_P)^2}{8\gamma} + \frac{c_P(1 + \tau_P)}{2} \text{ and } A_P = \frac{b(1 + \tau_P)}{2\gamma}.^{27}$$

As in the analysis above, we immediately see that the printed newspaper price is decreasing in own VAT if the value of the advertising market is sufficiently large compared to marginal costs. The results in Section 3 were thus not driven by the assumption that readers are indifferent to the ad level.

With single-homing, as with ad neutrality, simulations show that the digital platform responds by increasing its subscription price if the VAT is reduced (and it will reduce the number of ads it sells when consumers dislike ads). It is also still the case that the printed platform responds by increasing its reader price.

5.3 Positive marginal cost at the digital platform

Another simplification we have made is setting the digital platform's marginal cost to zero. This can easily be relaxed. Let the digital platform face a marginal cost c_D , where

²⁵ At this point we have $\frac{d^2\pi_D}{dp_D^2} = -\frac{2}{t(1+\tau_D)} < 0$, $\frac{d^2\pi_D}{dA_D^2} = -\frac{1}{4t(1+\tau_D)} \left(9b^2(\tau_D + 1)^2 + 4\gamma(v_D - d) \right) < 0$ and $\frac{d^2\pi_D}{dp_D^2} \cdot \frac{d^2\pi_D}{dA_D^2} - \left(\frac{d^2\pi_D}{dp_D dA_D} \right)^2 = \frac{b^2}{2t^2} + \frac{4\gamma(v_D - d)}{2t^2(1+\tau_D)^2} > 0$.

²⁶ We have then implicitly assumed that $A_D < 1$, which amounts to requiring that $\gamma > \frac{b(1+\tau_D)}{2}$. For lower values of γ we can trivially set $A_D = 1$.

²⁷ At this point we have $\frac{d^2\pi_P}{dp_P^2} = -\frac{2}{t(1+\tau_P)} < 0$,

$\frac{d^2\pi_P}{dA_P^2} = -\frac{1}{4t(1+\tau_P)} \left(9b^2(1 + \tau_P)^2 - 4\gamma c_P(1 + \tau_P) + 4\gamma(v_P - d) \right)$ and

$\frac{d^2\pi_P}{dp_P^2} \cdot \frac{d^2\pi_P}{dA_P^2} - \left(\frac{d^2\pi_P}{dp_P dA_P} \right)^2 = \frac{1}{2t^2(1+\tau_P)^2} \left(b^2(1 + \tau_P)^2 - 4\gamma c_P(1 + \tau_P) + 4\gamma(v_P - d) \right)$.

A necessary and sufficient condition for the second-order conditions to hold is therefore $c_P < \frac{v_P - d}{(1 + \tau_P)} + \frac{b^2(1 + \tau_P)}{4\gamma}$.

$c_P > c_D > 0$. The core mechanism underlying Propositions 1 and 5 remains, although now the directional findings are caveated. Mirroring the caveated findings for the printed paper in the main analysis with multi-homing we find that digital prices reduce with higher digital tax rates if the incremental value per consumer of second impressions is higher than the marginal cost of producing an extra copy ($\sigma b > c_D$). If there is a reduction in the tax rate on the digital platform under single-homing, both platforms will increase consumer prices if $b > c_D$ and both will decrease consumer prices if $b < c_D$ (see the Appendix).

6 Conclusion

We have assessed the impact of VAT policy in a two-sided market with asymmetric news platforms. Reducing the tax rate on digital subscriptions increases the profitability of the digital platform's reader market, but has no direct effect on the advertising market. The downward pressure on the digital subscription price exerted by the ad market is reduced and the price increases. Digital consumption decreases.

This intriguing finding does not depend on whether consumers multi-home. With shared consumers, reader prices are strategically independent, while they are strategic complements when consumers single-home. Both situations yield the same inverse relationship between the tax rate on digital subscriptions and the digital price.

Nor do our results hinge on the asymmetry of costs or tax rates. The models we provide can easily assess the impact of tax on two horizontally differentiated digital platforms by using a common tax rate and setting the marginal cost of the printed platform to zero. The inverse relationship between the tax rate and subscription prices holds.

To highlight the underlying mechanisms we assumed consumers were ad neutral in the main analysis. In Section 5 and the Appendix we showed that this assumption was not crucial for our core result.

Our model suggests an interesting relationship between two media policy goals. We have seen that reducing tax rates on digital news can reduce the degree of multi-homing and harm the media policy goal of pluralism. We note that in a more general setting the increased profitability stemming from the reduced tax rates could stimulate entry. Thus, in a two-sided market VAT reductions could still support a media policy aiming to increase media diversity, albeit by harming media pluralism. We leave formal analysis of this trade-

off to future research.

The predictions of the model presented here are strikingly different from the intuition of many policy makers and economists. In Europe, this would matter little if the straight jacket of the VAT Directive was to be maintained. But the European Commission has committed to reforming the restrictive VAT Directive and extending VAT reductions to digital newspapers and ebooks (European Commission, 2016a, 2016b) and the European Council expects agreement in the second half of 2017. Norway has already reduced the VAT rate on digital newspapers to zero assuming that reader prices will decrease. Our results predict the opposite effect.

For printed newspapers, empirical investigation would be particularly useful to identify whether ad market profits exceed variable costs or whether the incremental advertising profit from an extra reader exceeds the marginal cost of reaching that reader. The existence of free newspapers suggests the former holds and when the latter holds our results would question the effectiveness of the existing beneficial tax rates for printed newspapers. Although too late for the European debate, careful empirical investigation of the VAT change for electronic newspapers in Norway could test our theoretical predictions.

Acknowledgements

The authors are grateful to the Editor and two referees who substantially improved the paper. We would also like to thank Paul Belleflamme, Emilio Calvano, Chiara Canta, Tommy Staahl Gabrielsen, Steffen Juraneck, Marko Koethenbueger, Florian Schuett, Lukas Toth, Simen Aardal Ulsaker, Andre Veiga and participants of the CRESSE 2016 conference in Rhodes, Greece for helpful comments. We would also like to thank seminar participants in Bergen, Bonn, Oslo, Mannheim, Munich and Paris for helpful discussions.

7 References

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. and S. Coate (2005): Market provision of broadcasting: A welfare analysis. *Review of Economic Studies*, 72, 947-972.

Anderson, S., Ø. Foros, and H. J. Kind (2017): Product functionality, competition, and multi-purchasing. *International Economic Review*, 58(1) 183-210.

Anderson, S., Ø. Foros, and H. J. Kind (2018): Competition for advertisers and for viewers in media markets. *The Economic Journal*, 128, 34-54.

Anderson, S. and J. J. Gabszewicz (2006): *The media and advertising: A tale of two-sided markets*. In Ginsburgh, V. and D. Throsby (eds.) *Handbook of the Economics of Art and Culture*, North Holland, The Netherlands.

Armstrong, M. (2002): Competition in two-sided-markets. Mimeo, Nuffield College, Oxford.

Armstrong, M. (2006): Competition in two-sided markets. *The RAND Journal of Economics*, 37(3), 668-691.

Athey, S., E. Calvano, and J. Gans (forthcoming): The Impact of Consumer Multi-Homing on Advertising Markets and Media Competition. *Management Science*.

Chandra, A. and U. Kaiser, (2015): *Newspapers and magazines*. In Anderson, S., Stromberg, D. and J. Waldfogel (eds.) *Handbook of Media Economics*, North Holland, The United Kingdom, vol 1A, 397.

EFTA Surveillance Authority (2016): EFTA Surveillance Authority raising no objections to a zero VAT rate for electronic news services (Norway), Case No: 78337, Document No: 776581, Decision No: 023/16/COL.

European Commission (2015): VAT Rates Applied in the Member States of the European Union, Taxud.c.1(2015) - EN.

European Commission (2016a): Public Consultation on reduced and super-reduced VAT rates for electronically supplied publications. Retrieved from

http://ec.europa.eu/taxation_customs/common/consultations/tax/super_reduced_rates_vat_en.htm on 17 August 2016.

European Commission (2016b): Towards a single EU VAT area - Time to decide. COM(2016) 148 final.

European Commission (2016c): Impact Assessment SWD(2016) 392 final. Retrieved from

https://ec.europa.eu/taxation_customs/sites/taxation/files/swd_2016_392.pdf on 27 September 2017.

European Council (2017): Proposal for a Council Directive on reduced VAT rates for

electronic publications – Council general approach, April 2017. Retrieved from <http://data.consilium.europa.eu/data/2017/INIT/en/pdf> on 4 November 2017.

Fan, Y. (2013): Ownership consolidation and product characteristics: A study of the US daily newspaper market. *American Economic Review*, 103(5), 1598-1628.

Gentzkow, M. (2007): Valuing New Goods in a Model with Complementarity: Online Newspapers. *American Economic Review*, 97(3), 713-744.

Gentzkow, M., J. Shapiro, and M. Sinkinson (2011): The Effect of Newspaper Entry and Exit on Electoral Politics. *American Economic Review*, 101(7), 2980-3018.

Gentzkow, M., J. Shapiro, and M. Sinkinson (2014): Competition and Ideological Diversity: Historical Evidence from US Newspapers. *American Economic Review*, 104(10), 3073-3114.

Hotelling, H. (1929): Stability in Competition. *Economic Journal*, 39, 41-57.

Kind, H.J., M. Koethenbueger, and G. Schjelderup (2008): Efficiency-Enhancing Taxation in Two-Sided Markets. *Journal of Public Economics*, 92(5-6) 1531-39.

Kind, H.J., G. Schjelderup, and F. Stahler (2013): Newspaper Differentiation and Investments in Journalism: The Role of Tax Policy. *Economica* 80, 131-148.

Lebedev, E. (2016): Independent closure: Evgeny Lebedev's letter to staff. Retrieved from <http://www.theguardian.com/media/2016/feb/12/independent-closure-evgeny-lebedevs-letter-to-staff> on 14 April 2016.

Ofcom (2015): News consumption in the UK: research report. Retrieved from http://stakeholders.ofcom.org.uk/binaries/research/tv-research/news/2015/News_consumption_in_the_UK on 17 August 2016.

Peitz, M. and M. Reisinger (2015): *The economics of internet media*. In Anderson, S., Stromberg, D. and J. Waldfogel (eds.) *Handbook of Media Economics*, North Holland, The United Kingdom, vol 1A, 445.

Picard, R. G (2004): *The economics of the daily newspaper industry*. In Alexander, A., Owers, J., Carveth, R., Hollifield, C.A., and A.N. Greco (eds) *Media economics: Theory and Practice*, Lawrence Erlbaum Associates, The United States of America, vol 3, 109.

Statministerens Kontor (2015): Notification VAT rating for electronic news services. Retrieved from <https://www.regjeringen.no/contentassets/95cf7028fac3455d859305d0d72399f2/14-4781-notifikasjon-01-12-2015.pdf>. on 17 August 2016.

Tait, A. (1988): Value added tax: International practice and problems. *International*

Monetary Fund.

Williams, K. (2009): *Read All About It!: A History of the British Newspaper*. Routledge, The United Kingdom.

1 Appendix

This Appendix provides more detail on the results we present in the Main Paper.

1.1 Single-homing consumers

Suppose that each consumer buys one and only one of the media products.¹ We assume market sharing, and we find the location of the consumer who is indifferent between buying D or buying P by equalising the utility of purchasing each product and solving for x . Consumers to the left of x will purchase from D , those to the right will purchase from P . Consumer demand for good i thus equals

$$X_i^{SH} = \frac{1}{2} + \frac{v_i - v_j}{2t} - \frac{p_i - p_j}{2t} \quad (1)$$

which resembles the standard Hotelling set-up.

As in the model with shared consumers (Section 3 of the Main Paper) the platforms earn revenue from the subscription fee, p_i , and from advertising fees. Since each pair of eyeballs is available only through one outlet, ad revenue for platform i is simply equal to bX_i^{SH} . We can thus write platform profit as:

$$\pi_D = \frac{p_D}{1 + \tau_D} X_D^{SH} + bX_D^{SH} \quad (2)$$

$$\pi_P = \left(\frac{p_P}{1 + \tau_P} - c_P \right) X_P^{SH} + bX_P^{SH}. \quad (3)$$

The first term on the right-hand side of (2) and (3) is profits from the reader market, and the second term is profits from the advertising market.

Using the equilibrium outcomes from the Main Paper we have

$$\pi_D^{SH*} = \frac{(3t + (v_D - v_P) + (1 + \tau_P)(c_P - b) + (1 + \tau_D)b)^2}{18t(1 + \tau_D)} \quad (4)$$

$$\pi_P^{SH*} = \frac{(3t + (v_P - v_D) - (1 + \tau_P)(c_P - b) - (1 + \tau_D)b)^2}{18t(1 + \tau_P)}. \quad (5)$$

We can check the intuitive result that own profits are decreasing in own tax rate by

¹The market is covered iff $t \leq \frac{1}{3}(v_D + v_P + 2b + b(\tau_D + \tau_P) - c_P(1 + \tau_P))$. The market sharing condition is $t > \max \left[\frac{1}{3}(v_D - v_P + b(\tau_D - \tau_P) + c_P(1 + \tau_P)), \frac{1}{3}(v_P - v_D - b(\tau_D - \tau_P) - c_P(1 + \tau_P)) \right]$.

differentiating (4) and (5) by their respective tax rates:

$$\begin{aligned} \frac{d\pi_D^{SH*}}{d\tau_D} &= -\frac{(b^2(\tau_P - \tau_D)(\tau_P + \tau_D + 2) - 2b(\tau_P + 1)(c_P - v_P + v_D + \tau_P c_P) \\ &\quad + 3t(3t + 2c_P - 2v_P + 2v_D + 2\tau_P c_P) - 6bt(\tau_P + 1) + (c_P - v_P + v_D + \tau_P c_P)^2)}{18t(\tau_D + 1)^2} \\ \frac{d\pi_P^{SH*}}{d\tau_P} &= -\frac{(b^2(\tau_D - \tau_P)(\tau_P + \tau_D + 2) - 6bt(\tau_D + 1) + 2bc_P(\tau_P + 1)^2 - 2b(v_P - v_D)(\tau_D + 1) \\ &\quad + 3t(3t + 2v_P - 2v_D) - (c_P - v_P + v_D + \tau_P c_P)(c_P + v_P - v_D + \tau_P c_P))}{18t(\tau_P + 1)^2}. \end{aligned}$$

Since both numerators are positive under the market sharing condition, it immediately follows that $\frac{d\pi_i^{SH*}}{d\tau_i} < 0$.

1.2 Disutility of ads

With disutility of ads the profit level of platform i is

$$\pi_i = \left(\frac{p_i}{1 + \tau_i} - c_i \right) (X_i^e + X^s) + \beta A_i X_i^e + \sigma \beta A_i X^s \quad (i = D, P), \quad (6)$$

where $X^s = 0$ under single-homing.²

The utility of buying only from platform i for a consumer located at point x is given by $u_i = v_i - t|x - x_i| - p_i - \gamma A_i^2$, where $x_D = 0$, $x_P = 1$, and $\gamma > 0$. The utility of buying both is still $u_{(D+P)} = u_D + u_P - b$ and the incremental utility of multi-homing is specified as $u_{(D+P)} - \max[u_D, u_P]$. The location of the consumer indifferent between multi-homing and buying only the digital newspaper is $x_{DP} = 1 - \frac{v_P - d - p_P - \gamma A_P^2}{t}$ and the location of the consumer indifferent between multi-homing and buying only the printed newspaper is $x_{PD} = \frac{v_D - d - p_D - \gamma A_D^2}{t}$. As in the Main Paper, the digital platform's exclusive consumers can be represented as $X_D^e = x_{DP}$, shared consumers as $X^s = x_{PD} - x_{DP}$, and the print platform's exclusive consumers as $X_P^e = 1 - x_{PD}$.

²We maintain the assumption of homogenous advertisers and incremental pricing in the ad sector but verify that this is indeed an equilibrium also under ad disutility. Let platform i sell A_i ads and price them at their incremental value. Platform j could also price its ads incrementally ($a_j = \beta A_j X_j^e + \sigma \beta A_j X^s$) and then sell its ads to A_j advertisers. All advertisers, whether or not they were served by i , would be willing to purchase the ads. However those $1 - A_i$ advertisers not served by i would be willing to pay the full value for each consumer. If platform j sets the per ad price at $\beta(X_j^e + X^s)$ it could earn $\beta(1 - A_i)(X_j^e + X^s)$ on the ad market (and set its reader price accordingly). In the simulation we check that neither platform has an incentive to deviate in this way. When the value of second impressions is low enough compared to first impressions (i.e σ is low) then there may not be an incremental pricing equilibrium.

Let us start out by considering the multi-homing case ($X^s > 0$), where the first-order conditions for the platforms are

$$\frac{d\pi_i}{dp_i} = \left(\frac{X_i^e + X^s}{1 + \tau_i} + \left(\frac{p_i}{1 + \tau_i} - c_i \right) \frac{dX^s}{dp_i} \right) + \sigma\beta A_i \frac{dX^s}{dp_i} = 0 \text{ and} \quad (7)$$

$$\frac{d\pi_i}{dA_i} = \beta X_i^e + \sigma\beta X^s + \left(\frac{p_i}{1 + \tau_i} - c_i + \sigma\beta A_i \right) \frac{dX^s}{dA_i} = 0. \quad (8)$$

Assume first that $\sigma = 1$. We can then solve equations (7) and (8) and find closed form solutions. For the digital platform we arrive at

$$p_D = \frac{v_D - d}{2} - \frac{3b^2(1 + \tau_D)^2}{8\gamma} \text{ and } A_D = \frac{b(1 + \tau_D)}{2\gamma}.^3$$

The ad level is increasing in the value of the ad market and decreasing in consumer's distaste for ads ($dA_D/d\gamma < 0$).⁴ The reader price is decreasing in the value of the ad market ($dp_D/db < 0$) and increasing in consumer's distaste for ads ($dp_D/d\gamma > 0$). The intuition for the latter is that since higher disutility of ads reduces the ad volume, the willingness to pay for the newspaper increases. More interesting for our point of view, is the fact that we still have that the reader price is decreasing in own VAT rate ($dp_D/d\tau_D < 0$).

For the printed platform we find

$$p_P = \frac{v_P - d}{2} - \frac{3b^2(1 + \tau_P)^2}{8\gamma} + \frac{c_P(1 + \tau_P)}{2} \text{ and } A_P = \frac{b(1 + \tau_P)}{2\gamma}.^5$$

As in the case with ad neutrality (Section 3 of the Main Paper), we immediately see that the printed newspaper price is decreasing in own VAT if the value of the advertising market is sufficiently large compared to marginal costs (more precisely, $dp_P/d\tau_P < 0$ if $b > \sqrt{2c_P\gamma/[3(1 + \tau)]}$). The results in Section 3 of the Main paper were thus not driven by the assumption that readers are indifferent to the ad level.

³At this point we have $\frac{d^2\pi_D}{dp_D^2} = -\frac{2}{t(1+\tau_D)} < 0$, $\frac{d^2\pi_D}{dA_D^2} = -\frac{1}{4t(1+\tau_D)} \left(9b^2(\tau_D + 1)^2 + 4\gamma(v_D - d) \right) < 0 =$ and $\frac{d^2\pi_D}{dp_D^2} \cdot \frac{d^2\pi_D}{dA_D^2} - \left(\frac{d^2\pi_D}{dp_D dA_D} \right)^2 = \frac{b^2}{2t^2} + \frac{4\gamma(v_D - d)}{2t^2(1+\tau_D)^2} > 0$.

⁴We have then implicitly assumed that $A_D < 1$, which amounts to requiring that $\gamma > \frac{b(1+\tau_D)}{2}$. For lower values of γ we can trivially set $A_D = 1$.

⁵At this point we have $\frac{d^2\pi_P}{dp_P^2} = -\frac{2}{t(1+\tau_P)} < 0$,

$\frac{d^2\pi_P}{dA_P^2} = -\frac{1}{4t(1+\tau_P)} \left(9b^2(1 + \tau_P)^2 - 4\gamma c_P(1 + \tau_P) + 4\gamma(v_P - d) \right)$ and

$\frac{d^2\pi_P}{dp_P^2} \cdot \frac{d^2\pi_P}{dA_P^2} - \left(\frac{d^2\pi_P}{dp_P dA_P} \right)^2 = \frac{1}{2t^2(1+\tau_P)^2} \left(b^2(1 + \tau_P)^2 - 4\gamma c_P(1 + \tau_P) + 4\gamma(v_P - d) \right)$.

A necessary and sufficient condition for the second-order conditions to hold is therefore $c_P < \frac{v_P - d}{(1 + \tau_P)} + \frac{b^2(1 + \tau_P)}{4\gamma}$.

The case of $\sigma = 1$ is useful to understand the intuition. Less attractive is the implication that the value of readers on the ad market is then independent of whether the consumers read only one or both newspapers. For the more realistic case of $\sigma < 1$ it is not possible to find closed-form solutions, and we therefore have to rely on numerical analysis. Figure 4 shows how the price of the digital newspaper varies with τ_D for different values of σ (here we have set $\gamma = 0.35$).⁶ The inverse relationship between digital tax rates and digital reader prices is maintained. Following a VAT reduction on digital readers, the digital platform responds to the increased profitability of the reader market by increasing its reader price and reducing the number of ads it carries.

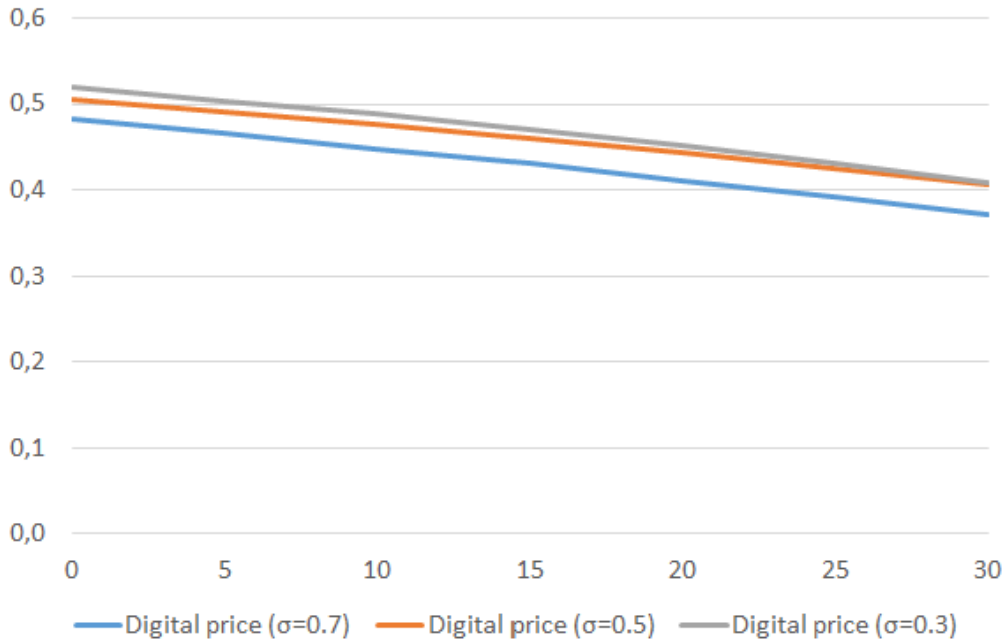


Figure 1: The impact of digital tax rates on digital prices.

Simulations show that more favourable taxation of digital newspapers need not reduce multi-homing if consumers dislike ads. The ambiguity hinges on the two countervailing effects of reducing the VAT rate: consumers dislike the higher prices but like the reduced number of ads. If the value of second impressions is high (σ large), the digital price increase can be large enough that the price increasing effect on consumers dominates the ad reduction effect. Then the number of multi-homing readers decreases, as in the main

⁶When solving (7) and (8) numerically, we have used the following parameter values: $v_D = v_P = 2$, $c_P = 0.2$, $t = 1.2$, $b = 0.45$, $d = 0.7$, $\tau_P = 0$.

analysis. If the value of second impressions is sufficiently low, however, the price increasing effect can be outweighed by the ad reduction effect and the number of multi-homing readers increases.

The ambiguous effects of digital tax reductions on the extent of multi-homing have knock-on implications for the printed platform's profitability. If a digital tax reduction reduces the number of multi-homers, as in the main analysis, the printed newspaper ends up with more exclusive readers (who are more valuable on the ad market), and the printed platform is better off. On the other hand if a digital tax reduction increases the number of multi-homers, then the printed platform ends up with fewer exclusive readers and is worse off. Thus, depending on parameter values, the printed platform's profits can increase or decrease with the digital tax rate.

With single-homing we follow a similar process, and platform i 's demand is $X_i^{SH} = \frac{1}{2} + \frac{v_i - v_j}{2t} - \frac{p_i - p_j}{2t} - \frac{\gamma A_i^2 - \gamma A_j^2}{2t}$. Platform i 's profit is then $\pi_i = \left(\frac{p_i}{1 + \tau_i} - c_i \right) X_i^{SH} + \beta A_i X_i^{SH}$, and we have the first-order conditions

$$\begin{aligned} \frac{d\pi_i}{dp_i} &= \left(\frac{1}{1 + \tau_i} \right) X_i^{SH} + \left(\frac{p_i}{1 + \tau_i} + \beta A_i - c_i \right) \frac{dX_i^{SH}}{dp_i} = 0 \text{ and} \\ \frac{d\pi_i}{dA_i} &= \beta X_i^{SH} + \left(\frac{p_i}{1 + \tau_i} + \beta A_i - c_i \right) \frac{dX_i^{SH}}{dA_i} = 0. \end{aligned}$$

As with single homing under ad neutrality (Section 4 in the Main Paper), simulations show that the digital platform responds by increasing its subscription price if the VAT is reduced (and it will reduce the number of ads it sells when consumers dislike ads). It is also still the case that the printed platform responds by increasing its reader price.

1.3 Positive marginal costs

This sections shows the robustness of our results to positive marginal costs at the digital platform. Let the digital platform face a marginal cost c_D such that, without loss of generality, $c_P > c_D > 0$.

1.3.1 Multi-homing consumers

The profit function of the digital firm is now

$$\pi_D = \left(\frac{p_D}{1 + \tau_D} - c_D + b \right) X_D^e + \left(\frac{p_D}{1 + \tau_D} - c_D + \sigma b \right) X^s$$

where, as in the main paper, the digital platform's exclusive consumers are represented as $X_D^e = x_{DP}$ and shared consumers are represented as $X^s = x_{PD} - x_{DP}$. The profit function of the printed platform is unchanged:

$$\pi_P = \left(\frac{p_P}{1 + \tau_P} - c_P + b\right)X_P^e + \left(\frac{p_P}{1 + \tau_P} - c_P + \sigma b\right)X^s.$$

Using the same process as in the main paper the best responses and equilibrium prices are

$$p_i = \frac{v_i - d}{2} + \frac{1}{2}(c_i - b\sigma)(1 + \tau_i).^7$$

As such we can amend Proposition 1:

Proposition 1*: Multi-homing with digital costs. *The price of the print platform is independent of the tax rate on the digital platform, and vice versa. Suppose that there is a decrease in the tax rate on*

a) *the digital platform. Then own price will increase if the incremental value of second impressions is worth more per consumer than the marginal cost of distributing an extra copy ($\sigma b - c_D > 0$).*

b) *the print platform. Then own price will increase if the incremental value of second impressions is worth more per consumer than the marginal cost of printing an extra copy ($\sigma b - c_P > 0$).*

As in the main paper, price independence does not imply profit independence. A tax reduction at the digital platform will increase profits at the printed platform if and only if the digital price increases. Specifically for $i = D, P$ and $i \neq j$ we now observe $\frac{d\pi_i}{d\tau_j} = -\frac{b(\sigma b - c_j)(1 - \sigma)}{2t}$ which is negative if $b\sigma > c_j$. Thus we can amend Proposition 2:

Proposition 2*: Multi-homing with digital costs. *The print platform's profit decreases in the tax rate of its rival if $b\sigma > c_D$. The digital platform's profit decreases in the rival's tax rate if $b\sigma > c_P$.*

The number of multi-homers is equal to

$$X^s = \frac{(v_D + v_P) + b\sigma(2 + \tau_D + \tau_P) - c_D(1 + \tau_D) - c_P(1 + \tau_P) - 2d}{2t} - 1, \quad (9)$$

⁷The second-order conditions are $\frac{d^2\pi_i}{dp_i^2} = -\frac{2}{t(1 + \tau_i)}$.

from which it immediately follows that the impact of tax changes on multi-homing consumers will also depend on the direction on the price change, $\frac{dX^s}{d\tau_i} = \frac{1}{2t} (b\sigma - c_i)$, and so we can amend Proposition 3:

Proposition 3*: Multi-homing with digital costs. *Reducing the tax rate on platform i decreases the number of multi-homers if the incremental value of multi-homing consumers is larger than that platform's marginal cost ($b\sigma > c_i$).*

It also immediately follows from (9) that the impact of changes in quality on the extent of multi-homing is unaffected by the inclusion of marginal costs for the digital firm. Thus, Proposition 4 is unamended.

Proposition 4*: Multi-homing with digital costs. *Assume that media quality improves (v_D and v_P increase). Then subscription prices and the extent of multi-homing increase.*

1.3.2 Single-homing consumers

The profit function of the digital firm is now

$$\pi_D = \left(\frac{p_D}{1 + \tau_D} - c_D \right) X_D^{SH} + bX_D^{SH}$$

and the profit function of the printed firm is unamended:

$$\pi_P = \left(\frac{p_P}{1 + \tau_P} - c_P \right) X_P^{SH} + bX_P^{SH}.$$

The best responses are now

$$p_i = \frac{t + (v_i - v_j) - (b - c_i)(1 + \tau_i)}{2} + \frac{p_j}{2}$$

and we have equilibrium prices of

$$p_i = t + \frac{(v_i - v_j) - 2(b - c_i)(1 + \tau_i) - (b - c_j)(1 + \tau_j)}{3}.^{8,9}$$

From this we can update Proposition 5:

⁸The second-order conditions are $\frac{d^2\pi_i}{dp_i^2} = -\frac{1}{t(1+\tau_i)}$

⁹Now the market is covered iff $t \leq \frac{1}{3} (v_D + v_P + 2b + b(\tau_D + \tau_P) - c_D(1 + \tau_D) - c_P(1 + \tau_P))$. The market sharing condition is $t > \max \left[\begin{array}{l} \frac{1}{3} (v_D - v_P + b(\tau_D - \tau_P) - (c_D(1 + \tau_D) - c_P(1 + \tau_P))), \\ \frac{1}{3} (v_P - v_D - b(\tau_D - \tau_P) + (c_D(1 + \tau_D) - c_P(1 + \tau_P))) \end{array} \right]$.

Proposition 5*: *Single-homing with digital costs. Suppose that there is a reduction in the tax rate on*

a) the digital platform. Then both platforms will increase consumer prices if $b > c_D$. Both will decrease consumer prices if $b < c_D$.

b) the print platform. Then both platforms will increase consumer prices if $b > c_P$. Both will decrease consumer prices if $b < c_P$.

As in the main paper with single-homing consumers, if platform i responds to a tax reduction by increasing its price under single-homing, platform j will capture a larger number of readers (and charge a higher price, since prices are strategic complements) and thus make higher profits from both reader and advertiser markets. We can thus simplify Proposition 6:

Proposition 6*: *Single-homing with digital costs. Profit is decreasing in own tax rate. The profit of a platform is decreasing in its rival's tax rate if both prices decrease in that tax rate.*

Finally given the location of the indifferent consumer

$$x = \frac{1}{2} + \frac{1}{6t} (v_D - v_P + (1 + \tau_D)(b - c_D) - (1 + \tau_P)(b - c_P))$$

we observe $\frac{dx}{d\tau_D} = \frac{1}{6t} (b - c_D)$. We can therefore amend Proposition 7:

Proposition 7*: *Single-homing with digital costs. Reducing the tax rate on the digital platform (τ_D) will increase sales of printed newspapers and reduce sales of digital newspapers if $b > c_D$.*

CHAPTER III

Price Increasing Tax Reductions for Electronic Newspapers: Implications for State Aid Policy

Price Increasing Tax Reductions for Electronic Newspapers: Implications for State Aid Policy

*Małgorzata Cyndecka and Tim Wyndham **

Abstract

The economic implications of indirect taxation have long been understood. Reducing indirect taxes results in lower prices. This straightforward and intuitive result has been an undisputed input into State aid policy, allowing robust legal analyses of tax exemptions. However, recent advances in the economic theory of two-sided markets yield new predictions for the impact of one type of indirect tax, the value added tax (VAT), on electronic newspapers. The surprising new theoretical prediction is that reader prices increase following VAT reductions. The prediction is driven by low marginal costs and the fact that advertisers would pay more for an ad seen by more readers. We bring this insight into the State aid literature and assess the implications for State aid policy. In particular, we raise questions over the legal assessment of a Norwegian State aid scheme, where a VAT exemption for print newspapers was extended to digital newspapers. If prices go up instead of down the policy will not have the intended effect and, under the given argumentation, the aid should not be viewed as compatible with the functioning of the internal market. The economic and legal arguments we present are important for the wider European debate on whether to reduce VAT on electronic newspapers.

I. Introduction

Economists are notorious for not being able to agree with each other. A rare and consistent area of agreement has been the implications of indirect taxations. Reducing indirect taxes will reduce prices, although the quantum of the price reduction ensures space for disagreement remains. This clarity of message, not always obtainable from economics in the realm of competition policy, has proved a useful input for the assessment of favourable tax treatments in State aid policy.

*Małgorzata Cyndecka, University of Bergen, Faculty of Law. <Malgorzata.Cyndecka@jur.uib.no>; Tim Wyndham, NHH Norwegian School of Economics. <Timothy.Wyndham@nhh.no>. We thank Hans Jarle Kind, Alf Erling Risa, Håvard Sandvik, Iraj Hashi, Niklas Dürr and Jaap Bos for helpful comments.

In the last two decades or so, there has been significant development in the economic understanding of two-sided markets. Two-sided markets have been defined as being markets in which a firm acts as a platform: It sells two different products to two groups of consumers, while recognising that the demand from one group of consumers depends on the demand from the other group and, possibly, vice-versa.¹

In this paper we ask first what economics has to say regarding the implications of one type of indirect tax, value added tax (VAT), in two-sided markets. Second, we investigate the implications for State aid policy. We are not aware of any other paper that assesses the State aid implications of this new economics literature. To give tractability and direct policy relevance we focus on the market for digital news.

State aid law prevents EU Member States from favouring certain undertakings by granting them aid that distorts competition and prevents a ‘level playing field’. Yet, the ban on granting aid is not absolute. Aid measures that address a well-defined market failure may be declared compatible with the internal market. For this purpose, however, the positive effects of such aid must outweigh its negative effects on competition.

An example of aid that may be declared compatible with the internal market is financial support that Member States give to their printed media industries. Indeed, most EU Member States grant printed newspapers lower rates of VAT than their benchmark national rates. In 2017, 26 out of 28 Member States had some form of reduced VAT rate available to printed newspaper subscriptions.² Despite this support to printed newspapers, national policy makers in the EU wishing to extend beneficial VAT rates to their electronically supplied equivalents face an immovable constraint: The 2006 VAT Directive³ that ‘excludes any possibility of a reduced VAT rate being applied to “electronically supplied services”’.⁴ This constraint is about to be

¹ L Filistrucchi, D Geradin, E van Damme and P Affeldt, ‘Market Definition in Two-Sided Markets: Theory and Practice’, 10(2) *Journal of Competition Law and Economics* (2014).

² See European Commission, VAT Rates applied in the Member States of the European Union, <http://ec.europa.eu/taxation_customs/sites/taxation/files/resources/documents/taxation/vat/how_vat_works/rates/vat_rates_en.pdf> (accessed 6 August 2018).

³ Council Directive 2006/112/EC of 28 November 2006 on the common system for value added tax, 2006 OJ L 347/1.

⁴ In 2015 the CJEU found that the VAT Directive ‘excludes any possibility of a reduced VAT rate being applied to “electronically supplied services”’. See Case C-479/13 *Commission v France* ECLI:EU:C:2015:141; and Case C-502/13 *Commission v Luxembourg* ECLI:EU:C:2015:143, as referred to by the Commission’s Impact Assessment on the proposed amendment to the VAT Directive, https://ec.europa.eu/taxation_customs/sites/taxation/files/swd_2016_392.pdf, p 11 (accessed 6 August 2018).

removed. There is widespread support within the European Council to give each Member State the flexibility to give electronic and printed newspapers the same VAT rate.

Norway, as a member of the European Economic Area (EEA) but not the EU, has a competition policy regime that is substantially equivalent to that of the EU.⁵ In other policy areas, including VAT policy, Norway has more freedom. The VAT Directive is not a part of the EEA Agreement.⁶ Thus, when Norway wished to reduce VAT on electronic newspapers the only legal impediment was State aid policy.

In January 2016, the European Free Trade Association Surveillance Authority (ESA), declared a zero VAT rate for electronic news services in Norway to be a State aid that is compatible with the functioning of the EEA Agreement.⁷ A crucial part of the decision was an assumption that following a decrease of the VAT rate on electronic newspapers, prices would decrease. An emerging economics literature on the impacts of taxation in two-sided markets suggests that reader prices might actually increase if the VAT rate falls.⁸ Media firms have been characterized as a prominent or typical example of a two-sided market, selling content to readers and advertising space to advertisers.⁹ The two-sidedness of newspaper markets, as in other media markets, arises because the demand for ads in a newspaper increases with the number of readers.¹⁰

In part II we will introduce and explain the economics literature behind this puzzling prediction of price increasing tax reductions. In part III we proceed to assess whether ESA was

⁵ Supra-national monitoring of compliance with competition rules, including State aid rules, is conducted by the European Free Trade Association Surveillance Authority, ESA, whose competences mirror the competences of the European Commission (the Commission) with regard to the EU Member States, see Protocol 26 to the EEA Agreement and the Agreement between the EFTA States on the establishment of a Surveillance Authority and a Court of Justice (the SCA) and its Protocol 3 sets out the powers of ESA in the field of state aid. The latter is complemented by EFTA Surveillance Authority, Decision of 14 July 2004 on the implementing provisions to referred to under Article 27 in Part II of Protocol 3 to the Agreement between the EFTA States on the establishment of a Surveillance Authority and a Court of Justice, 195/04/COL.

⁶ Agreement on the European Economic Area, 1994 OJ L 1/3, and EFTA States' official gazettes.

⁷ See EFTA Surveillance Authority, Decision of 25 January 2016 raising no objections to a zero VAT rate for electronic news services, 023/16/COL, (the VAT decision).

⁸ See H J Kind, M Koethenbueger and G Schjelderup, 'Efficiency-Enhancing Taxation in Two-Sided Markets', (2008) 92 (5-6) *Journal of Public Economics*, 153, and H J Kind and J Moen, 'Effects of taxes and subsidies on media services' in R G Picard and S S Wildman (eds) *Handbook on the economics of the media* (Edward Elgar Publishing 2015).

⁹ See L Filistrucchi, D Geradin, E van Damme and P Affeldt, 'Market Definition in Two-Sided Markets: Theory and Practice', 10(2) *Journal of Competition Law and Economics* (2014); L Filistrucchi, TJ Klein and T Michielsen, 'Assessing Unilateral Merger Effects in a Two-Sided Market: An Application to the Dutch Daily Newspaper Market', 8(2) *Journal of Competition Law and Economics* (2012).

¹⁰ *Ibid.*

correct to assess that there was aid, and whether it was correct to assess that the aid was compatible with the functioning of the internal market. Despite the measure equalizing VAT rates for printed and electronic newspapers, the ESA was correct to find that the measure is a distortion of competition. The new insights from economic theory however directly question whether the aid was compatible with the internal market. ESA's approval was based on the assumption that VAT reductions would lead to falling prices and increasing demand. Part IV concludes.

II. The Economics of Taxation in Electronic Newspaper Markets

How indirect taxes affect prices in imperfectly competitive markets has been an important question in economics for nearly 200 years.¹¹ In perfectly competitive markets the firm is a price taker and it does not matter whether a firm faces a specific tax per unit (the tax payable depends on quantity) or an *ad valorem* tax (the tax payable depends on price and quantity).¹² In markets where a firm has some control over its price (either directly through setting its price or indirectly through setting its quantity) the impact of the different forms of taxation on price, profits and tax revenue can be starkly different.¹³ Although the size of effects vary between different types of taxes and in different competitive conditions the *direction* of price changes is consistently predicted.¹⁴ An increase in the tax rate would be expected to increase the price.

New results from economic theory now suggest that this positive relationship cannot be taken for granted in two-sided markets.¹⁵ Imagine an electronic newspaper facing an exceptionally high VAT rate on reader subscriptions. The newspaper can just set the reader price

¹¹ See M Keen, 'The Balance between Specific and Ad Valorem Taxation', (1998) 19(1) Fiscal Studies 1, for a summary. A A Cournot and I Fisher, *Researches into the Mathematical Principles of the Theory of Wealth*, (Macmillan Co., 1897 - original version 1838) is recognized by Keen as the first to identify differential effects of specific and ad valorem taxes in the most extreme case of imperfect competition: monopoly.

¹² Ibid, Keen.

¹³ Ibid. Keen specifies an identity for the producer price, P_n , in relation to the *ad valorem* tax rate, v , the consumer price, P , and the specific tax rate; $P_n \equiv (1-v)P - s$. Thus specific taxes act to shift the demand curve down whereas *ad valorem* taxes rotate the demand curve down.

¹⁴ One notable exception is the case of vertical differentiation. See H Cremer and J-F Thisse. 'Commodity taxation in a differentiated oligopoly', (1994) International Economic Review 613. We restrict ourselves from further discussing the wider tax literature.

¹⁵ See, H J Kind, M Koethenbueger and G Schjelderup, 'Efficiency-Enhancing Taxation in Two-Sided Markets', (2008) 92 (5-6) Journal of Public Economics, 1531, and HJ Kind and J Moen, 'Effects of taxes and subsidies on media services' in R G Picard and S S Wildman (eds) *Handbook on the economics of the media* (Edward Elgar Publishing 2015).

to zero, make no profit on the reader market and instead make all its profit from advertisers. In this simple example, presented by Kind and Moen,¹⁶ we observe high taxes on the reader price and a low reader price. This is in stark contrast to the standard intuition that high taxes will lead to high prices.

1. The Economic Literature on Taxation in Two-Side Markets

There have now been a number of papers that have demonstrated this novel result. The first paper to highlight this was Kind et al which studied the impact of tax changes on a monopolist and in perfect competition and maintained the finding when one side exerts negative externalities on the other side (i.e. ads reduce readers utility).¹⁷ Another paper demonstrated that the finding was maintained when duopolists chose their editorial stance and invested in journalistic quality.¹⁸ These papers assumed symmetric firms and that consumers only bought one newspaper. A recent contribution by Foros et al extended these results by including asymmetric tax rates, asymmetric costs and allowing for multi-homing readers.¹⁹

In all these papers, a crucial part of the mechanism is the fact that advertisers want their advert to be seen by more readers. Thus, the same ad space yields higher ad revenues for a platform if it has more readers. That is to say that readers yield positive indirect network effects.

2. The Economic Intuition

To help us explain the intuition we imagine an electronic newspaper that is prohibited from selling adverts. Naturally, the profit maximizing firm will set its price such that its marginal revenue equals its marginal cost. At this point a small decrease in price would not be profitable since the extra revenue gained from new sales would not outweigh the lost profit from the

¹⁶ Ibid.

¹⁷ See H J Kind, M Koethenbueger and G Schjelderup, 'Efficiency-Enhancing Taxation in Two-Sided Markets', (2008) 92 (5-6) *Journal of Public Economics*, 1531.

¹⁸ See H J Kind, G Schjelderup and F Stähler, 'Newspaper differentiation and investments in journalism: The role of tax policy' (2013) 80(317) *Economica*, 131.

¹⁹ See Ø Foros, H J Kind and T Wyndham, 'Tax-free digital news?', *International Journal of Industrial Organisation* (forthcoming).

reduced margin on pre-existing customers, and the cost of producing and distributing the extra copies to reach these new customers.²⁰

What would happen if the prohibition on adverts were lifted? To give an intuitive description we will assume that the utility of a reader depends only on the level of content in a newspaper and not on the volume of ads.²¹ We can re-evaluate the impact of a small price decrease at the point before ads were allowed (where reader market marginal revenue equalled marginal costs). The three effects we identified for the one-sided firm remain exactly the same. But now there is an additional revenue gain from reducing price. The lower reader price will lead to more readers and since advertisers, *ceteris paribus*, will pay more for an ad viewed by more readers, advertising revenues will increase. The profit maximizing two-sided firm will therefore set a lower price than the one-sided firm.

One way to visualize this is to think about the extra ad revenue each individual reader generates as acting to lower the marginal cost curve of producing and distributing newspapers.²² The net costs of reaching an extra consumer will be lower in the presence of advertising than without (and if the advertising market is valuable enough, the net marginal cost may be negative). In this setting it is entirely intuitive that the firm with lower (net) costs has lower prices.

More generally, the two-sided firm performs a balancing act. It seeks to maximise its total profits, but there is an inherent tension between its two sources of profit: the reader market and the advertising market. We have illustrated that the two-sided reader price is lower than the one-sided reader price. But given that the one-sided price was maximising reader market profits, we know that reducing the reader market price will reduce the reader market profit. It is optimal for the two-sided firm to reduce its price because the additional profit from the advertising market exceeds the reduced reader market profit. The two-sided firm will therefore reduce prices until the profit increases in the advertising market are exactly offset by the profit decreases in the reader market. This balancing of profits between the two markets is important for the second part of the mechanism.

²⁰ For electronic newspapers it is reasonable to assume that production and distribution costs of one accessing one extra reader are close to zero.

²¹ This assumption makes the intuition easier to describe. It does not drive the result.

²² We are grateful to Bruno Jullien who made this comment in response to a separate paper.

The second important part of the mechanism is the *ad valorem* nature of the tax. We now consider the introduction of VAT on reader subscriptions. There is no tax on the advertising revenues so the benefits of reducing prices, in terms of increased advertising market profits, are unchanged relative to before the VAT introduction.²³ In terms of foregone reader market profits, the cost of lowering the price has now decreased. It remains true that reducing the price will reduce the (gross) margin on infra-marginal readers. However there is a second effect that partly mitigates this loss, specifically, reducing the price also reduces the amount of the tax. This second effect implies that the opportunity cost (in terms of foregone reader market profits) of lowering prices is now lower than the benefit of reducing prices and so the firm decreases its price following the introduction of the VAT.

Our assumption that the newspaper is electronic has helped us to explain this counterintuitive effect. An electronic newspaper has a marginal cost of zero, or very near to zero. Since in one-sided markets ‘If marginal costs are negligible ... then the VAT rate has almost no effect on price and output’²⁴ we can ignore any countervailing impact from the standard mechanisms that lead to prices increasing in taxes. The higher the marginal cost, the more valuable the advertising market is required to maintain this result, for given curvatures of demand on both sides of the market.

It is important to note that this intuition does not apply to specific taxes or subsidies. The reason is that following the introduction of a specific tax, which does not depend on price, the firm does not benefit from this tax reduction effect from reducing prices. Formally, VAT acts to rotate the firm’s effective demand curve while the specific tax shifts it down.

We end this section by noting that tax reductions simply unwind the effect of introducing a tax. When the electronic newspaper faced the introduction of VAT we showed why it would reduce its price. The tax incentivized the firm to increase its advertising profit by reducing its reader market profit. Reducing or removing the VAT reduces or removes this incentive, and prices increase.

²³ We only need to assume that the tax rate on adverts does not change, which is an accurate reflection of the Norwegian and European policy changes.

²⁴ H J Kind, M Koethenbueger and G Schjelderup, ‘Efficiency-Enhancing Taxation in Two-Sided Markets’, (2008) 92 (5-6) *Journal of Public Economics*, 1531, and HJ Kind and J Moen, ‘Effects of taxes and subsidies on media services’ in R G Picard and S S Wildman (eds) *Handbook on the economics of the media* (Edward Elgar Publishing 2015).

III. Legal Analysis of the ESA Decision

In this part we analyse the ESA decision. We first present the decision before commenting on the ESA's finding on the presence of aid. In particular, we focus on two of five cumulative conditions that must be met under Articles 107(1) TFEU and 61(1) EEA in order to qualify a given measure as aid: selectivity, and distortion or threat of distortion of competition. This is because we do contest that the VAT reduction implies a transfer of public resources that amounts to an economic advantage affecting trade between the Member States (Contracting Parties in the EEA Agreement). The conditions of selectivity and distortion of competition are, however, worth analysing in more detail because of the context in which the measure at stake was introduced. This analysis is provided in sections 2a and 2b. Then, given the finding that the VAT reduction was State aid we assess its compatibility with the internal market. At its simplest, the Commission or ESA may declare a given aid to be compatible with the internal market if the competition distortions it necessarily causes are outweighed by its positive effects. Following the implementation of the State aid Modernisation reform, the Commission has introduced seven Common Assessment Principles that a given aid measure must meet in order to pass the compatibility test. Those are provided in section 3. We focus on the criterion most susceptible to criticism from the new economic theory described above, namely the presence of an incentive effect.

1. The Decision and the Measure

In January 2016, ESA declared a zero VAT rate for electronic news service in Norway to be aid that is compatible with the EEA Agreement.²⁵ As the notified zero VAT rate was a type of aid that was not covered by any existing State aid guidelines, the compatibility assessment was made directly under Article 61(3)(c) EEA.²⁶ ESA's approval is valid until 1 March 2022.²⁷

²⁵ The VAT decision (n 7).

²⁶ As Article 61(3)(c) EEA provides, 'The following may be considered to be compatible with the functioning of this Agreement: aid to facilitate the development of certain economic activities or of certain economic areas, where such aid does not adversely affect trading conditions to an extent contrary to the common interest.' In the EU, aid to the press may be declared compatible under Article 107(3)(c) TFEU whose counterpart is Article 61(3)(c) EEA, or under Article 107(3)(d) TFEU. The latter allows the Commission to approve 'aid to promote culture and heritage conservation'. Although Article 107(3)(d) TFEU does not have its counterpart in the EEA Agreement, ESA considers that cultural aid may be approved on the basis of Article 61(3)(c) EEA. In its assessment, ESA will apply

The zero VAT rate for printed news has existed since 1970 when Norway introduced VAT.²⁸ A printed publication is considered a newspaper that is entitled to VAT zero rating if it: (1) is published at least once a week; (2) informs the public about news and current affairs home or abroad; (3) has a responsible editor; and (4) charges payment for the paper.²⁹

The ‘electronic news services’ are subject to certain eligibility criteria as well. According to the notified VAT Regulation, only those electronic news services that: (1) mainly include news and current affairs content from different areas of society; (2) have the general public as target audience; (3) have an editor-in-chief; and (4) are published at least once a week,³⁰ are not ‘disadvantaged’ by the VAT rate of 25%.³¹

2. Legal Assessment of the Existence of Aid

The notion of aid is based on the EU Courts’ interpretation of Article 107(1) TFEU. The substance of this Article and Article 61(1) EEA, which is of relevance here, is identical. The differences in the wording are only technical.³² A state measure qualifies as aid if the following five cumulative conditions are met: (1) transfer of public resources, (2) grant of an economic

the same criteria as those applied by the Commission for the purposes of Article 107(3)(d) TFEU. See ESA’s 2014 Film and Audiovisual Guidelines, which correspond to Communication from the Commission on State aid for films and other audiovisual works, 2013 OJ C 332/1. More on the Commission’s practice as regards aid to the press, see E Psychogiopoulou, ‘State Aids to the Press: The EU’s Perspective’ (2012) 11(1) European State Aid Law Quarterly 57.

²⁷ The VAT decision, (n 7) Article 2.

²⁸ See Notification: VAT zero rating for electronic news services, 1 December 2015, p 3 (the Notification).

²⁹ See Lov. om merverdiavgift (merverdiavgiftsloven), LOV-2009-06-19-58, § 6-1, (the VAT Act, section 6-1), <https://lovdata.no/dokument/NL/lov/2009-06-19-58/KAPITTEL_6#KAPITTEL_6> (accessed 6 August 2018).

³⁰ See the VAT Act, (n 29), [section 6-2], and Forskrift til merverdiavgiftsloven (merverdiavgiftsforskriften), FOR-2009-12-15-1540, § 6-2-1, (the VAT Regulation, section 6-2-1), <https://lovdata.no/dokument/SF/forskrift/2009-12-15-1540#KAPITTEL_6> (accessed 6 August 2018).

³¹ According to the President of ESA, ‘[t]he new zero VAT rate makes it possible for news media, including the large number of local and regional newspapers in Norway, to publish and sell their content electronically without being disadvantaged by the VAT system ... and will promote the consumption of news and current affairs media published in electronic form, which is of increasing importance for customers in Norway’, see ESA Press Release, (21 January 2016), <<http://www.eftasurv.int/press--publications/press-releases/state-aid/electronic-news-services>> (accessed 6 August 2018).

³² While Article 107(1) TFEU stipulates that ‘Save as otherwise provided in the Treaties, any aid granted by a Member State or through State resources in any form whatsoever which distorts or threatens to distort competition by favouring certain undertakings or the production of certain goods shall, in so far as it affects trade between Member States, be incompatible with the internal market’, Article 61(1) EEA provides that ‘Save as otherwise provided in this Agreement, any aid granted by EC Member States, EFTA States or through State resources in any form whatsoever which distorts or threatens to distort competition by favouring certain undertakings or the production of certain goods shall, in so far as it affects trade between Contracting Parties, be incompatible with the functioning of this Agreement.’

advantage, (3) selectivity, (4) distortion, or a threat of distortion of competition, and (5) an effect on trade between the Member States (Contracting Parties in the EEA Agreement). Based on those conditions ESA found the zero VAT rate for electronic news services to be aid.³³

At first sight, it seems that the notified scheme seemed to have equalized the VAT rate for printed and electronic news and therefore removed any competitive distortions arising from Norway's taxation policy between these two sectors, at least compared to the counterfactual of no revisions to the Norwegian tax code. It is thus reasonable to ask why the proposed measure was considered distortive (or indeed selective) and how the effect on competition should be taken into account for the purposes of declaring the measure as aid. The answers to those questions require providing an explanation of the conditions of selectivity and distortion or threat of distortion of competition under State aid rules.

a. The Condition of Selectivity

Under Article 61(1) EEA, a state measure is selective if it favours 'certain undertakings or the production of certain goods'. Consequently, measures of general application, ie those that apply across all sectors and to all economic operators within a given Contracting Party, do not fall within the scope of aid.³⁴ The fact that the given measure is not aimed at one or more specific recipients defined in advance, but that it is subject to a series of objective criteria according to which it may be granted, does not exclude the selective nature of that measure.³⁵

In the present case, in line with the well-established practice concerning selectivity of tax measures, the ESA applied a three-step test.³⁶ The first step amounts to identifying the system of reference. The very existence of an advantage may be established only when compared with the 'common or normal taxation'.³⁷ In the present case, the reference system was the general VAT

³³ See also Commission Notice on the notion of State aid as referred to in Article 107(1) of the Treaty on the Functioning of the European Union (NoA), 2016 OJ C 262/1, (NoA).

³⁴ See Joined Cases E-17/10 and E-6/11 *Liechtenstein v ESA* [2012] EFTA Ct. Rep. 114 [53], and the case law cited.

³⁵ See Case T-55/99 *Confederación Española de Transporte de Mercancías v Commission*, ECLI:EU:T:2000:223, [40].

³⁶ See NoA (n 33) [127 et seq].

³⁷ See Case C-88/03 *Portugal v Commission* ECLI:EU:C:2006:511, [56].

rule in Norway according to which the supply of goods and services is subject to VAT at a rate of 25%.³⁸

The second step determines whether a given measure is a derogation from that system insofar as it differentiates between economic operators who, in light of the objectives intrinsic to the system, are in a comparable factual and legal situation.³⁹ This allows to conclude whether the measure in question is *prima facie* selective. If it does not constitute a derogation from the reference system in question, it is not selective.⁴⁰ *Prima facie* selectivity, however, is verified and may be ruled out in the third step.

In the present case, one could identify two types of potential aid beneficiaries. First, the ESA excluded selectivity with regard to the direct beneficiaries, ie the consumers of electronic news services, who were mostly private individuals who would benefit from the anticipated fall in prices. They could not receive aid within the meaning of Article 61(1) EEA as only entities that qualify as undertakings may qualify as aid beneficiaries under State aid law.⁴¹ As regards undertakings that purchase electronic news services, any undertaking in Norway that purchases electronic news services would be able to benefit from this effect. This excluded the condition of selectivity in relation to such potential direct beneficiaries as well.⁴² Indeed, all such undertakings could and would benefit from the introduced measure.⁴³ Second, with regard to media companies selling electronic news, the zero VAT rate was *prima facie* selective. This is because it would only apply to publishers of electronic news services that fulfil the above-mentioned eligibility criteria in light of the objective intrinsic to the system, ie generating income for the state. Such publishers were in the same factual and legal situation as those that did not benefit from the zero VAT rate.

³⁸As provided in NoA, ‘The reference system is composed of a consistent set of rules that generally apply — on the basis of objective criteria — to all undertakings falling within its scope as defined by its objective. Typically, those rules define not only the scope of the system, but also the conditions under which the system applies, the rights and obligations of undertakings subject to it and the technicalities of the functioning of the system. In the case of taxes, the reference system is based on such elements as the tax base, the taxable persons, the taxable event and the tax rates,’ NoA (n 33) [133-134].

³⁹In this respect, see the leading Case C-143/99 *Adria-Wien Pipeline*, ECLI:EU:C:2001, [41].

⁴⁰See NoA, (n 33) [128].

⁴¹Under competition law, which includes State aid rules, the notion of undertaking is defined as an entity engaged in an ‘economic activity’, regardless of its legal status and the way in which it is financed. ‘Economic activity’ means ‘any activity consisting in offering goods and services in a given market’, see Case C-41/90 *Höfner and Elser* ECLI:EU:C:1991:161, [21]; Joined Cases C-159/91 and C-160/91 *Poucet and Pistre* ECLI:EU:C:1993:63, [17]; Joined Cases C-180/98 to C-184/98 *Pavlov and Others* ECLI:EU:C:2000:428, [74].

⁴²The VAT decision, (n 7) [52].

⁴³As ESA clarified, ‘Every undertaking can benefit from the measure if it purchases electronic news services. This is an objective condition not subject to any discretion by the tax administration’, the VAT decision, (n 7)[54].

Importantly, a Member State may legitimately invoke a few objectives when arguing that the beneficiary undertakings were not in a comparable factual and legal situation to those that did not benefit from a given measure. Amongst such objectives are the protection of public health, environment or culture.⁴⁴ In the present case, however, none of them was of relevance.

Having established that a given tax measure is *prima facie* selective, one proceeds to the third step. In essence, it amounts to verifying whether the derogation is justified by the logic or general nature of the reference system.⁴⁵ If this is the case, a *prima facie* selective measure will not be considered selective within the meaning of State aid law. This is the case where a measure derives directly from the intrinsic basic or guiding principles of the reference system or where it is the result of inherent mechanisms necessary for the functioning and effectiveness of the system.⁴⁶ In this respect, one may mention the need to fight fraud or tax evasion, administrative manageability, the principle of tax neutrality or the need to take into account specific accounting requirements. Yet, any external (extrinsic) policy objectives that are not inherent to the system at stake may not be taken into account.⁴⁷

In the present case, the Norwegian authorities argued that, in general, reduced and zero VAT rates for certain goods and services formed an integral part of the Norwegian VAT system and, consequently, were not selective.⁴⁸ Yet, as the ESA stressed, one of the principles of the Norwegian tax system was that the consumption of goods or services should be charged with VAT at a standard rate of 25% in order to generate revenues for the state. Exceptions from that principle could be justified by an objective of common interest, but they were not part of the logic and general nature of the consumption tax system.⁴⁹ Indeed, the objective of media pluralism and media diversity was not amongst the inherent objectives of the Norwegian VAT system. The zero VAT rate for electronic news services was thus selective.

⁴⁴ More on this issue, see M Honoré, 'Selectivity' in P Werner and V Verouden (eds) *EU State aid control law and economics* (Wolters Kluwer 2017) 119, 129.

⁴⁵ See, for example, Case C-279/08 P *Commission v Netherlands* ECLI:EU:C:2011:551, [62]; Joined Cases E-17/10 and E-6/11 *Liechtenstein v ESA* [2012] EFTA Ct. Rep. 114, [74].

⁴⁶ See, for example, Joined Cases C-78/08 to C-80/08 *Paint Graphos and others* ECLI:EU:C:2011:550, [69].

⁴⁷ See, for example, Case C-88/03 *Portugal v Commission* ECLI:EU:C:2006:511, [81]; Joined Cases C-78/08 to C-80/08 *Paint Graphos and others* ECLI:EU:C:2011:550, [69].

⁴⁸ The Norwegian authorities used the same argument when denying selectivity of the zero VAT rate for electric vehicles, see EFTA Surveillance Authority, Decision of 21 April 2015 on State aid measures in favour of electric vehicles, 150/15/COL, [101].

⁴⁹ See the VAT decision, (n 7)[59]. See, likewise, EFTA Surveillance Authority, Decision of 8 May 2014 on certain amendments to Act 50/1988 on Value Added Tax applicable to customers of Icelandic data centers, 193/14/COL, [58].

This selectivity finding is consistent with the ESA's decisional practice. For example, a zero VAT rate for electronic vehicles in Norway was deemed selective since environmental protection was not one of the inherent objectives of the Norwegian VAT system.⁵⁰ Similarly, in the case of non-imposition of VAT on supply of mixed services to customers of data centres and for import of servers, the objective of bringing the Icelandic data centre industry to a comparable level with the data centre industry in the EU, and thus attracting a mobile and tax sensitive service sector to Iceland, were not amongst the inherent objectives of the Icelandic VAT system.⁵¹ Iceland pursued a political and economic objective when it introduced amendments to its VAT system.⁵²

It is understandable that the ESA wishes to hold the line that VAT reductions are selective, by their nature. We agree that this should indeed be the presumption. Were this not the case individual countries would, in effect, be able to argue that any VAT reductions are outside of the scope of State aid rules.

We agree with the ESA's findings of selectivity in this case, albeit with two caveats relating to argumentation. Our position is that qualifying media companies should be classified as direct, not indirect, beneficiaries. The first order effects of the tax change is that, for a given price, a company pays less tax. In many cases a company will also lower its prices, but the present case is a counter-example. Moreover, a presumption that reductions of indirect tax burdens will primarily benefit the consumer contradicts the Commission's practice concerning indirect taxes.⁵³ Unless proven otherwise, it is thus the taxable person who is to be regarded as the beneficiary of aid.⁵⁴

⁵⁰ EFTA Surveillance Authority, Decision of 21 April 2015 on State aid measures in favour of electric vehicles, 150/15/COL, [103].

⁵¹ EFTA Surveillance Authority, Decision of 8 May 2014 on certain amendments to Act 50/1988 on Value Added Tax applicable to customers of Icelandic data centers, 193/14/COL, [60].

⁵² See, likewise, European Commission, Decision of 17 February 2003 State aid implemented by the Netherlands for international financing activities, 2003/515/EC, [95]. Therein, the Dutch scheme's express aim was to encourage large multinationals to transfer their financing activities back to the Netherlands. This was an economic aim and it was not inherent in a taxation system.

⁵³ See J Englisch, 'State Aid and Indirect Taxation' in: A Rust and C Micheau (eds), *State aid and tax law* (Kluwer Law International 2013) 69, 76. In this respect, Englisch refers to European Commission, Decision of 15 December 2009 The Dutch exemption from environmental taxes for ceramic producers, COM(2009)9972, 31; European Commission, Decision of 3 September 2010 - The Latvian tax exemptions for primary producers, COM(2010)5769 21; European Commission, Decision of 8 April - The British exemption from the climate change levy for electrified rail, 2011 COM(2011)2613, 15.

⁵⁴ See Englisch (n 53). In this respect, Englisch relied on the Commission practice as well, see, European Commission, Decision of 30 October 2009 - The Finnish excise duty exemption for biofuel, COM(2009)8497, 18; European Commission, Decision of 9 April 2010 Bulgaria - Tax reduction for biofuels, C(2010)2219, 29; European Commission, Decision of 19 April 2010 - The German tax rebates for biofuels, COM(2010)2557, 4.

b. The Condition of Distortion of Competition

As regards the condition of distortion, or a threat of distortion of competition, given the pre-existing VAT reduction for printed newspapers one could argue that the zero VAT rate on electronic newspapers removes rather than contributes towards a distortion of competition. Indeed, according to the Norwegian authorities, the scale of the pre-existing distortion problem has increased since newspaper circulation has fallen and the digitization of news media services has increased, so now a larger share of the news consumption falls outside of the scope of the zero rated VAT.⁵⁵ If the measure does not distort competition, it does not qualify as State aid.

However, the ESA did not agree. In its view, the proposed measure benefits undertakings selling electronic news services by increasing demand for their services. These undertakings are active in the publishing and/or broadcasting sectors, which are open to competition and trade within the EEA. Moreover, these undertakings are active in other markets, for example for advertisement space, which are also subject to trade and competition. In addition, the proposed zero VAT rate only benefits those undertakings that sell electronic news services fulfilling the eligibility conditions. This strengthens their competitive position in comparison to other publishers and broadcasters that do not benefit from the measure.⁵⁶

Whilst we agree that the proposed measure benefits undertakings selling electronic news services we do not agree that demand for their services will increase. Such a prediction is based on the assumption that companies will respond to the VAT reduction by reducing prices. However, the economic theory highlighted in this article predicts that they will respond to the lower tax rate on reader prices by increasing reader prices. Non-recipients of the VAT reduction would expect to be (weakly) better off. If they compete with recipients who subsequently increase their reader price they would expect to benefit from this, if not they would be unaffected. In any case, the competitors of aid beneficiaries are not affected in a negative way.

Under State aid law, in order to establish a competition distortion, it is not necessary to define the market or to carry out a thorough investigation as regards the impact that the given

⁵⁵ See the Notification (n 28) p 7.

⁵⁶ The VAT decision (n 7) [61-64].

state measure may have on the undertakings concerned.⁵⁷ No actual assessment of that condition is required.⁵⁸ While the Commission or ESA do not have to establish that competition would be affected, it is sufficient that the measure in question is capable of having such an effect.⁵⁹ In this regard, the mere fact of having obtained an economic advantage that a given undertaking would not have obtained in normal market conditions is sufficient to conclude that the measure in question is liable to distort competition.⁶⁰ In the case at hand, the fact that the new zero VAT rate for electronic newspapers undoubtedly reduces the differences in the level of state support between printed newspapers and electronic newspapers is not sufficient to exclude the presence of aid. Given the arguments provided by the ESA, we agree that one cannot exclude that competition is distorted even if to a smaller extent than prior to introducing the aid measure at stake.

Yet, the measure's 'positive' effect on competition is duly taken into account in the course of compatibility assessment. This next step of evaluating aid is dealt with below.

3. Legal Assessment of the Compatibility of the Aid

As mentioned above, there are seven Common Assessment Principles on the basis of which the Commission or ESA may declare aid to be compatible.⁶¹ These principles are: (1) contribution to a well-defined objective of common interest; (2) need for state intervention; (3) appropriateness of State aid as a policy instrument; (4) existence of an incentive effect; (5) proportionality of the aid amount (aid limited to minimum necessary); (6) avoidance of undue negative effects on competition and trade; and (7) transparency. The ESA decision hinged on the existence of an incentive effect. We now explain the legal background to the incentive effect before assessing

⁵⁷ See Case C-211/05 *Italy v Commission* ECLI:EU:T:2009:304, [157-160].

⁵⁸ See Joined cases T-204/97 and T-270/07 *EPAC v Commission* ECLI:EU:T:2000:148, [85].

⁵⁹ See Case T-288/97 *Regione Friuli Venezia Giulia v Commission* ECLI:EU:T:2001:115, [133].

⁶⁰ The VAT decision, (n 7) [61-64]. See also Joined cases T-298/97, T-312/97, T-313/97, T-315/97, T-600/97 to 607/97, T-1/98, T-3/98 to T-6/98 and T-23/98 *Alzetta Mauro and others v Commission* ECLI:EU:T:2000:151, [141-147]; Case C-280/00 *Altmark Trans* ECLI:EU:C:2003:415.

⁶¹ The common assessment principles were introduced following the State Aid Modernisation reform (SAM) that was launched in 2012. See Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions, State Aid Modernisation, COM/2012/0209 final, and <http://ec.europa.eu/competition/state_aid/modernisation/index_en.html> (accessed 6 August 2018).

this component of the ESA decision. Lastly, some comments on reducing the distortion of competition will be provided.

a. Legal Background to the Incentive Effect

The Commission has gradually put more emphasis on the existence of an incentive effect when assessing the compatibility of aid under Article 107(3)(c) TFEU⁶² on which Article 61(3)(c) EEA is based. This is illustrated by formalizing the requirement of the incentive effect in successive State aid reforms.⁶³ In the State Aid Action Plan (SAAP) of 2005,⁶⁴ the incentive effect was included in the so-called balancing test that verified whether the benefits of aid could prevail over the distortions of competition resulting from granting that aid.⁶⁵ As put by Nicolaidis, an incentive effect exists if a given aid measure can induce the aid beneficiary to take measures that it would not normally take under conditions of free competition.⁶⁶ Following the 2012 State aid Modernisation reform (SAM), the incentive effect became one of the seven Common Assessment Principles. Thus, the Commission will consider a State aid measure compatible with the internal market only if the aid changes the behaviour of the undertaking(s)

⁶² V Verouden, 'EU State Aid Control: The Quest for Effectiveness', (2015) 4 European State Aid Law Quarterly 459, 460.

⁶³ See, in particular, Communication from the Commission — Criteria for the analysis of the compatibility of State aid for training subject to individual notification, 2009 OJ C 188/1, section 2.3; European Commission, Regulation No 651/2014 of 17 June 2014 declaring certain categories of aid compatible with the internal market in application of Articles 107 and 108 of the Treaty, 2014 OJ L 187/1, Article 6; European Commission, Guidelines on regional State aid for 2014-2020, 2013 OJ C 231/1, sections 3.5-3.6; European Commission, Framework for State aid for research and development and innovation, 2014 OJ C 198/1, section 4.4 ; European Commission, Guidelines on State aid for environmental protection and energy 2014-2020, 2014 OJ C 200/1, sections 3.2.4-3.2.5.

⁶⁴ State Aid Action Plan - Less and better targeted state aid: a roadmap for state aid reform 2005 – 2009, Consultation document adopted on 7 June 2005.

⁶⁵ The balancing test consisted of the following questions: (1) Is the aid measure aimed at a well-defined objective of common interest? (2) Is the aid well designed to deliver the objective of common interest ie does the proposed aid address the market failure or other objectives? (i) Is the aid an appropriate policy instrument to address the policy objective concerned? (ii) Is there an incentive effect, ie does the aid change the behaviour of the aid recipient? (iii) Is the aid measure proportionate to the problem tackled, ie could the same change in behaviour not be obtained with less aid? (3) Are the distortions of competition and effect on trade limited, so that the overall balance is positive? See Common principles for an economic assessment of the compatibility of State aid under Article 87(3) EC [now 107(3) TFEU] 1, 9, <http://ec.europa.eu/competition/state_aid/reform/economic_assessment_en.pdf> (accessed 6 August 2018).

⁶⁶ P Nicolaidis, 'Incentive Effect: is State Aid Necessary when Investment Is Unnecessary?' (2008) 2 European State Aid Quarterly 230, 230.

concerned in such a way that it engages in additional activity which it would not carry out without the aid or which it would carry out in a restricted or different manner.⁶⁷

The requirement of an incentive effect must be seen in the context of safeguarding the effectiveness of given aid in the pursuit of a well-defined objective of common interest. Contributing to such an objective is the first Common Assessment Principle.⁶⁸ Consequently, the Commission has the right to deny compatibility of aid that is not likely to induce the beneficiary to contribute to one of the objectives provided in Article 107(3) TFEU.⁶⁹

The ESA acceptance of the zero VAT rating for the supply, import and leasing of electric vehicles in Norway is an example of the incentive effect.⁷⁰ Those measures aimed to decrease the price of electronic vehicles thus encouraging demand for these vehicles and leading to environmental benefits resulting from using electronic vehicles instead of fossil-fuelled ones.

b. The Incentive Effect in the VAT Decision

The Norwegian authorities argued that the objective of the notified zero VAT rate for electronic newspapers was the promotion of media pluralism and media diversity.⁷¹ In explaining the need for state intervention, the authorities referred to a major shift from newspapers to electronic media over the past years, which made the existing aid scheme that concerned only printed newspapers obsolete. The rapid digitalization of news media raised important equity concerns related to considerable demographic differences in media consumption. The Norwegian authorities had grounds to fear that the existing zero VAT rate for printed newspapers, which aimed to lower prices for the consumers and thus increase the consumption of news, did not meet its objective with regard to the younger generation. The younger generation simply uses different publication platforms.

The ESA considered the support of the consumption of news media (to increase the demand for news services) to be the objective pursued by the notified scheme.⁷² Lower prices of

⁶⁷ Instead of issuing a separate document of those principles, the Commission decided to include them in its updated guidelines.

⁶⁸ More on that perspective in: Verouden, (n 62) 459.

⁶⁹ Case T-126/99 *Graphischer Maschinenbau v Commission* ECLI:EU:T:2002:116, [34].

⁷⁰ EFTA Surveillance Authority, Decision of 21 April 2015 on State aid measures in favour of electric vehicles. This decision also covered the supply and import of batteries for such vehicles, 150/15/COL.

⁷¹ The VAT decision, (n 7) [71].

⁷² *Ibid* [89].

electronic news would lead to higher consumption. In its decision, the ESA relied on the arguments and reasoning as they were presented by the Norwegian authorities, in particular, a study provided in Copenhagen Economics (2007) Taxation Papers.⁷³ Consequently, the ESA assumed that if one permanently lowered the VAT rate on a given good or service, sooner or later this would result in lowering the price of the good or service. The reduction on price would be more or less equal to the monetary equivalent of the lower VAT rate.⁷⁴ The cited study, however, admitted that the effect of the lower VAT rate depended on the consumer response to reduced prices of the particular good or service and the level of competition within the given sector. While the ESA recognized that the study cited provided limited evidence regarding the effects of lowering VAT rates on prices and consumption in the market for electronic news services, it agreed with the Norwegian authorities that an increase in VAT rate changed consumer prices, and thus demand.⁷⁵ The ESA found the Norwegian news media sector to be highly competitive since many news media companies competed for readers. It considered it likely that a zero VAT rate on electronic news services would lead to lower prices for the consumer compared to the current situation, ie charging VAT at 25%.⁷⁶ As a result, the ESA concluded that the notified zero VAT rate did have an incentive effect for consumers by bolstering their demand for electronic news services.⁷⁷

The economic theory highlighted in this article questions whether VAT reductions would bolster demand for electronic news services. Indeed, the emerging predictions from the economics literature of two-sided markets directly question the presumption that prices will decrease. It is important to note that this literature is not, per se, at odds with the empirical evidence provided by Norway and discussed by the ESA.⁷⁸ The literature predicts that the direction of price changes depends on the value of an extra reader on the ad market relative to the newspapers marginal costs. In the electronic newspaper market, the zero marginal cost yields unambiguous predictions for an inverse relationship between VAT and prices. But in the printed

⁷³ Copenhagen Economics (2007) Taxation Papers, Study on reduced VAT applied to goods and services in the Member State of the European Union, Working Paper NO 13 2007, European Commission, available at: <http://ec.europa.eu/taxation_customs/resources/documents/taxation/gen_info/economic_analysis/tax_papers/taxation_paper_13_en.pdf> (accessed 6 August 2018).

⁷⁴ The VAT decision, (n 7) [91].

⁷⁵ Ibid [92-94].

⁷⁶ Ibid [95].

⁷⁷ Ibid.

⁷⁸ Ibid [93].

market there are positive marginal costs and the theoretical predictions depend upon the relative value of the advertising market.

c. Avoidance of Undue Negative Effects on Competition and Reducing the Distortion of Competition

For State aid to be compatible with the functioning of the EEA Agreement, the negative effects of aid in terms of distortions of competition must be limited and outweighed by the positive effects in terms of contribution to the objective of common interest.⁷⁹ As the ESA pointed out, the proposed measure would apply to all eligible general news and current affairs media, irrespective of their distribution channel and the type of media. In particular, the measure is not limited to Norwegian electronic news services, but services of all origins are eligible for the zero VAT rate (provided that they are sold in Norway). Therefore, the effect on competition between undertakings active in the provision of electronic news services in Norway, including in relation to related markets, eg for advertisement space, is limited.

Moreover, the zero VAT rate on electronic news reduces an existing distortion of competition. As the ESA rightly noted, printed newspapers are also subject to a zero VAT rate based on an existing aid scheme. The proposed zero VAT rate for electronic news services will remove any difference in VAT treatment between distribution platforms for news and current affairs media. As a result, media companies will receive equal tax treatment irrespective of their choice of distribution channels for their news and current affairs products or services.⁸⁰

As the zero VAT rate on electronic news puts printed and electronic news on the same level in terms of VAT treatment, we agree that this state measure reduces distortion of competition compared to the situation prior to its implementation.

4. Summary of the Legal Implications from the Economic Analysis

ESA and the Norwegian government expected the proposed VAT reduction to lower the price of electronic news, increase consumption and promote the plurality and diversity of the media. The

⁷⁹ Ibid [102].

⁸⁰ Ibid [105].

new economic prediction from part II questions this expectation and suggests the opposite. Under this view, the VAT reduction should not be viewed as compatible with the functioning of the internal market.

IV. Conclusion

This paper has shown that recent developments in our economic understanding of two-sided markets question long-held intuitions about the impact of VAT changes on prices. The two-sidedness of the market, the cost structure of electronic newspapers and the *ad valorem* nature of VAT together suggest that prices can increase following VAT reductions.

These new predictions from economic theory question the compatibility of VAT reductions with State aid policy. The ESA approved a Norwegian VAT reduction on electronic newspapers as an aid measure that was compatible with the functioning of the internal market. The decision was predicated on prices falling and demand for electronically supplied newspapers increasing. Economic theory now questions this assumption and by extension this legal decision.

The economic and legal analysis in this article is relevant for a wider European debate. Policy makers in the EU are seeking to allow lower VAT rates on electronic publications. To date, the VAT Directive has restricted EU Member States from introducing lowered or zero VAT rates for electronic news services.⁸¹ There is widespread support within the European Council to give each Member State the flexibility to give electronic and printed newspapers the same VAT rate. The Council expected agreement to be achieved in the second half of 2017. This, however, has turned out to be a very optimistic scenario. As provided on the Council's website, on 13 July 2018, the Economic and Financial Affairs Council discussed the proposal on reduced VAT rates for electronic publications. Ministers were unable to reach an agreement and agreed to discuss it again at the next Council in October 2018.⁸²

Moreover, the analysis appears to have been neglected. For example the Commission's impact assessment recognizes the presence of the advertisement sector in electronic newspaper

⁸¹ Despite this prohibition, France has applied a super-reduced rate of 2.1% to the digital press since February 2014, Belgium applied a super-reduced rate between 2012 and 2015. Since 2016 Italy has applied a super-reduced rate of 4% to digital newspapers. The Commission discuss observed VAT rates in their Impact Assessment on the proposed amendment to the VAT Directive, (n 4) pp 10-11.

⁸² See <<http://www.consilium.europa.eu/en/policies/reduced-vat-e-publications/>> (accessed 6 August 2018).

and periodical markets,⁸³ but then assess that prices for e-publications (which comprise electronically supplied books, newspapers and periodicals) would all fall by the same proportion (half) of any VAT reduction.⁸⁴ We have not seen any mention of State aid policy in the policy process, and indeed the Commission's proposal for a Directive did not mention State aid policy.⁸⁵

Whether the theoretical predictions presented in this article translate into actual price rises is an empirical question. This appears to be a fruitful area for future research, in particular given the possible staggered roll out of any changes in the wider European Union.

⁸³ The VAT Directive (n 4) p 14.

⁸⁴ Ibid [25].

⁸⁵ See Proposal for a Council Directive amending Directive 2006/112/EC, as regards rates of value added tax applied to books, newspapers and periodicals available at <http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM:2016:0758:FIN> (accessed 6 August 2018).

CHAPTER IV

Outlet proximity, alcohol sales and sick leave:

Evidence from Norway

Outlet proximity, alcohol sales and sick leave: Evidence from Norway*

Oddmund Berg[†]

Tim Wyndham[‡]

Abstract

We present new evidence on the relationship between alcohol consumption and sick leave. The rapid expansion of a State-owned monopolist of high strength alcohol provides a novel opportunity to cleanly identify the impact of increased proximity to outlets on sales. We exploit this expansion as a plausibly exogenous increase in the regional availability of alcohol, or a decrease in the generalized price, to estimate the causal effect of alcohol consumption on sick leave. We find that an increase of alcohol sales of 1 percent in a quarter leads to 0.16 percent more men taking sick leave in that quarter, at the mean.

Keywords: Alcohol consumption; travel costs; sick leave

JEL codes: D12 ; I18; L12; L66

*We gratefully acknowledge comments and suggestions from Gernot Doppelhofer, Andreas Fagereng, Martin Holm, Bjørn Olav Johansen, Hans Jarle Kind, Øivind Anti Nilsen, Anja Roth, Morten Sæthre and Andreas Tveito. We would also like to thank seminar participants at the Norwegian School of Economics, Statistics Norway and the 2018 NHH-UoB Phd Workshop.

[†]Norwegian School of Economics

[‡]Norwegian School of Economics

1 Introduction

Alcohol is a commonly used drug in most countries. As well as private costs and benefits, consumption of alcohol often yields externalities. One potential wider cost to society is alcohol related sick leave. If alcohol consumption leads to employees being absent, there will be costs to the employer, to colleagues and potentially to the wider economy.

This paper aims to establish the causal impact of alcohol consumption on observed sick absence. Previous studies have assessed the association between alcohol and sick leave, but reverse causality and self-selection have prevented the attribution of causality (Norström (2006), Norström and Moan (2009) and Schou and Moan (2016)). We use data from an expansion of the Norwegian State-owned monopolist of high strength alcohol, Vinmonopolet, to show that there is a positive relationship between alcohol availability and sales (our first stage). In the second stage, we exploit this plausibly exogenous variation in availability over time and between regions, to study the effect of alcohol consumption on sick absence.

Our first stage results suggest that if the average driving distance to the nearest Vinmonpolet in a region decreases by 1km in a quarter then quarterly per capita expenditure on alcohol in that region increases by 1.45 percent. This translates to an implied travel cost per kilometer of 40 cents, which is largely in line with previous findings in the literature on proximity and demand. We mitigate potential concerns that our first stage is picking up the effects of omitted variables, by performing a synthetic control analysis of our store openings. This analysis provides additional evidence that the increased alcohol demand is driven by increased proximity arising from new store openings.

In the second stage, we find that an increase in alcohol consumption of 1 percent leads to an increase of sick leave in men of around 0.3 per 10,000 men, an increase of around 0.16 percent, at the mean. Our finding, using official sick leave data, is robust across a range of specifications. Apart from Pidd, Berry, Roche, and Harrison (2006), who use survey data, we are not aware of any papers that have estimated this causal relationship. Our results for women, and when we aggregate across genders, are of a similar magnitude but are less statistically robust.

For our first stage analysis, we consider there to be three main underlying mechanisms driving the observed results. To a greater or lesser degree, they all rely on a proportion of consumers having an element of time inconsistency or constraints in the storage or transportation of alcohol. Hinnosaar (2016) found that of 16 percent of consumers who bought beer regularly displayed time inconsistency in their purchases. The first mechanism relates to consumers who make trips to specifically buy alcohol. This mechanism relies on the seminal work of Hotelling (1929) where consumers make purchase decisions based on generalized price that incorporates not only the cost of the goods to be bought, but also the cost of getting to the store. A new store in a region will reduce the travel time for some consumers, and for those consumers the generalized

price of visiting a Vinmonopolet store reduces, and they may make more trips to buy alcohol. The second mechanism relates to consumers who can now plan alcohol purchases as part of their shopping routine, as opposed to having to change their routine to purchase alcohol. Since Vinmonopolet stores are usually based in shopping centers that also include supermarkets and other shops, this seems plausible. These consumers will also face a reduction in travel costs to purchase alcohol and may make more trips to the store. The final mechanism arises from this co-location effect. Some consumers may spontaneously enter Vinmonopolet whilst in a shopping centre, even though they had no intention of making a purchase when they initially planned their wider shopping trip. To the extent such purchases are not substitutes for previous purchases (that is they do not anticipate such spontaneity, nor adjust in later visits) an increase in proximity will increase consumption.

Although there is an extensive literature investigating the association between alcohol consumption and sick leave, there is limited empirical evidence of alcohol consumption causing sick leave. [Norström and Moan \(2009\)](#) used time series data from Norway between 1957 and 2001 to assess the relationship between sickness absence for manual employees and per capita alcohol sales. Using annual data, they found that a 1 liter increase in alcohol consumption was associated with a 13 percent increase in sick leave amongst men. This result was similar to that found previously in Sweden by [Norström \(2006\)](#). Both studies only claim an association. [Schou and Moan \(2016\)](#) reviewed the association literature and found consistent relationships between alcohol consumption and short term sick leave. They found that relationships between alcohol consumption and long term sick leave were less consistently found, although high quality studies, measured by the variables used and sample characteristics, always found a statistically significant link. [Johansson, Böckerman, and Uutela \(2008\)](#) used survey data to establish the association between alcohol consumption and sick leave in Finland. One survey has suggested a causal relationship: [Pidd, Berry, Roche, and Harrison \(2006\)](#) use Australian data where 3.5 percent of those who were in work and were current drinkers reported having missed at least one day in the previous 3 months due to alcohol consumption, suggesting that alcohol related sick days could represent about 6 percent of total sick days. We are not aware of any papers that use recorded sick leave data to establish a causal relationship between alcohol consumption and sick leave.

Our first stage evidence is consistent with the two prominent strands of literature assessing the impact of proximity on demand. One strand studies the impact of proximity on purchase decisions for consumers facing similar regulatory conditions and estimates the monetary value of proximity (the marginal cost of travel). [Seim and Waldfogel \(2013\)](#) find a travel cost of between 39 and 157 cents per kilometre depending on the proportion of households who have access to a car. Analysis of American movie markets by [Davis \(2006\)](#) suggested that the marginal cost of travel starts at 31 cents initially and then falls by about 8 cents per mile. In an analysis of the market for speciality coffee at the University of Virginia, [McManus \(2007\)](#) suggests that consumers would pay 40 cents to avoid traveling a tenth of a mile. The analysis of commuting

paths in Quebec City by Houde (2012) suggests that the median consumer's value of a minute of shopping is 90 cents, though concedes that this is likely an over-estimate.

Another strand of literature has quantified the role of distance in cross-border shopping. The further consumers are from a border the less likely they are to travel to benefit from lower taxes and duties. In terms of alcohol, two papers have taken advantage of price differentials in Scandinavia, where Norwegian prices exceed Swedish prices, which in turn exceed Danish prices. Asplund, Friberg, and Wilander (2007) found that the cross-price elasticity of regional alcohol demand in Sweden with respect to Danish alcohol prices was about 0.3 at the border. 150 (400) kilometers away from the border this reduced to 0.2 (0.1). Beatty, Larsen, and Sommervoll (2009) find that store-level revenues in Norway increase with distance from the Swedish border in an economically significant manner, up to about two and a half hours travel time.

Taking these two strands of literature together, there is a clearly demonstrated role of proximity on demand for a number of goods. When stores are nearer, travel costs (and therefore the generalized price faced by the consumer) is lower and demand increases. Our empirical environment has the benefit that there is no price competition. Vinmonopolet outlets have identical prices nationwide, and are not subject to outside competition for the vast majority of their products.¹ In any case, Vinmonopolet sets prices according to a transparent mark-up rule and does not seek to maximize profits. Thus, we can cleanly identify the role of distance without confounding competitive effects.

To further ensure confidence in our first stage results, and since the functional form of distance reduction on demand is the key ingredient in our two-stage estimation, we also treat the expansion as a series of natural experiments. We use the latest techniques from the econometrics of case studies to do a non parametric investigation. Specifically, we adopt the synthetic control approach of Abadie, Diamond, and Hainmueller (2010), ensuring that we find control groups that best match the 166 treatment groups. This methodology has been applied to such diverse economic questions as the impact of economic liberalisation (Billmeier and Nannicini, 2013) or natural disasters (Cavallo, Galiani, Noy, and Pantano, 2013), and hospital pricing (Garmon, 2017). The results we generate from treating each opening as an individual policy experiment enhance our confidence in our reduced form results, namely that increased proximity increases customer demand. We therefore consider that our first stage evidence of the role of proximity on demand is robust, economically and statistically significant and in line with previous literature.

Our clean identification of the role of proximity in alcohol consumption allows us to make causal statements about the impact of alcohol consumption on sick leave in a highly transparent manner. The critical assumption for our analysis is that the distance reduction from a store opening is only related to sick leave through increased alcohol consumption. The nature of the rollout of the new stores, which lead to the

¹Supermarkets can sell beer up to 4.75 percent alcohol. Wine and spirits are by far the largest revenue sources for Vinmonopolet. In 2016, wine constituted 63 percent of revenues, spirits a further 33. Authors calculations from Note 2 of Vinmonopolet's 2017 Annual Report.

changes in proximity, gives us confidence in this assumption.²

We proceed by describing our data, before discussing the expansion of Vinmonopolet. We then discuss our identification strategy before we present the detailed results from our first and second stages. Finally, we conclude.

2 Data

Our sick leave data are quarterly and cover the period from 2000 to 2016 at the municipality level. Sick leave is defined as the number of people registered as having taken at least one day of sickness absence. This data is publicly available from Statistics Norway, where we also collect municipality level data on socio-economic variables. Specifically, we collect information on employment shares, age composition and median financial characteristics such as income and bank deposits. For each municipality we calculate the share of the population that are of working age.

Our alcohol data comes from Vinmonopolet. We have monthly store level data on revenue and liquid volume from 2000 to 2016. The volume data can be further broken down on the five categories of product sold: Beer, wine, strong wine, liquor and non-alcoholic, although we focus on overall sales. We use alcohol purchases at Vinmonopolet as a proxy for alcohol consumption. A concern of this approach is that we might observe an increase in expenditure, and mistake this for an increase in consumption, if consumers simply substitute from other sources of alcohol from other sources. Another concern is that alcohol can be stored, which may cause discrepancies between purchase and consumption dates.

We have two arguments against this. First, [Hinnosaar \(2016\)](#) shows that sixteen percent of regular purchasers of beer display time inconsistent preferences. Thus, when it is easier to buy beer, we can reasonably expect that consumption will go up and not just be substituted. To the extent that consumers of other alcohol types also display time inconsistency, and are also likely to substitute to either low strength beer, or consumption at bars and restaurants we can be confident that the increase in sales at Vinmonopolet reflects an increase in overall consumption. Second, the available substitutes for Vinmonopolet are poor. Bars and restaurants offer the same products as Vinmonopolet, but consumption of those products is legally restricted to the time and place of purchase. Most beers sold in supermarkets are imperfect substitutes for the beers sold in Vinmonopolet, due to the restriction on alcoholic strength.³ Consumers can produce some types of alcohol in Norway, although it is hard to replicate the quality of commercial alternatives

²Previous studies have used rollouts as their identification strategy. For example, to investigate the role of the internet on sex crimes, [Bhuller, Havnes, Leuven, and Mogstad \(2013\)](#) used the rollout of broadband in Norway and [Dinkelman \(2011\)](#) used the rollout of electricity access to look at the impact of electrification on employment in South Africa.

³Most aisle space in supermarkets is allocated to half liter cans of medium strength beer which are not frequently sold at Vinmonopolet.

and requires planning. Taking these arguments together we might expect to see that beer sales are more responsive to distance reductions than the other categories, as releasing the spatial constraint allows them to buy their preferred beer rather than an imperfect substitute, indeed this might be considered a validity check.⁴ In essence, our argument is that releasing these externally imposed spatial constraints allows consumers to purchase their preferred type of alcohol rather than imperfect substitutes and, since this increases the availability of their preferred bundle of alcoholic goods their consumption of alcoholic goods increases. Indeed, this is part of the argument for the continuing spatial restrictions.

All data are aggregated to quarterly Labor Market Region (LMR) level observations for our analysis. The essence of a LMR is that if you live within a given LMR, you will also work in that LMR. We connect sales to the population by assuming sales within an LMR represents the alcohol consumption by the respective population living and working there. We believe that this is a reasonable approximation, since the definition of LMRs is based on residential and commuting patterns.⁵ The precise definition of LMRs we apply comes from [Bhuller \(2009\)](#).⁶

2.1 Proximity data

We follow [Seim and Waldfogel \(2013\)](#) and define clusters where people live to calculate the population weighted distance from the center of the cluster to the nearest store. We use driving distance calculated using GPS coordinates and the Georoute software from [Weber and Péclat \(2017\)](#). We label the clusters as “population centers”. In our data, each LMR consists of between 3 and 94 such population centers.

In our setting, a population center is a place that either already has, or will eventually receive, a store. In municipalities where there are no stores present at any time, we use the administrative center as the population center. We implicitly assume that residences are evenly distributed around these centers, such that the traveling distance associated with buying alcohol for the population living around center m is given by the distance from center m to the nearest population center with a Vinmonopolet store. If there is a store in the population center, the traveling distance is set to zero.

To find the population associated with each population center, we use municipality level data and divide the population on the number of centers if the municipality has more than one. The average traveling distance per capita in a LMR is calculated by summing the population weighted distances for all population centers m and dividing by the total population.

⁴Although we do not present the results here, if we breakdown our first stage analysis by type of product, we observe that beer is more responsive than other categories of alcohol.

⁵Individuals will also buy alcohol abroad or from duty-free stores. Furthermore, the proximity and attractiveness of bars and restaurants may differ across LMRs. We control for this with regional fixed effects.

⁶We have redone the first stage using the official LMR-classification of 90 LMRs. Our conclusions do not depend upon which coding we use. The official classification imposes a LMR to belong to one and only one county. Norway has 14 counties. We find the [Bhuller \(2009\)](#) classification the most reasonable for our purposes.

$$AD_{rt}^m = \frac{\sum_{m=1}^{n^r} d_{mst} * P_{mt}}{\sum_{m=1}^{n^r} P_{rt}}$$

To assess the precision of our approach, we calculate a more precise per capita distance using data from around 55,000 population grids. This is only available on a yearly basis from 2008-2016. We use the same approach as above, but use these 1 squared kilometer areas instead. Since the population data at municipality level is available quarterly and for a longer time range, we use that data for our main analysis.

2.2 Summary statistics

Table 1 displays summary statistics of our data. The proportion of individuals taking leave during a quarter has fluctuated during the period of our analysis. Women and men display similar trends and movements, with women at a higher absolute level. Per capita alcohol sales have increased over the period of analysis, growing during the first decade, and since partially falling back. In terms of volume of liquid per capita sold, wine has increased steadily over the period with a less pronounced tailing off at the end. The increase and subsequent reduction of liquor sold, more closely track the sales figures. The remaining three categories are much less important. Strong wine has declined from a low base, while non-alcoholic beverages have increased but from an even lower base. Beer declined initially, but has increased in the last decade, most likely reflecting the growing market for specialty or "craft" beers. Beer sales remain low, however. These data suggest that for every bottle of beer sold in the beginning of 2015 twelve bottles of wine were sold.

3 The expansion of Vinmonopolet

3.1 The market for high strength alcohol in Norway

The sale of alcohol for consumption "off-premises" in Norway is strictly controlled: Vinmonopolet is the only legal vendor of beverages with more than 4.75% alcohol, aside from the usual tax-free stores for travelers. In 1997, the Government decided to partially relax the restrictions on the number of outlets. For our period of analysis, between 2000 and 2016, the number of outlets increased from 129 to 324.⁷ A cap on the number of stores, set by the Ministry of Health remains in place.⁸ The number and distribution of stores remains subject to plans set out by the Ministry of Health. The new outlets were relatively evenly spread across Norway, as can be seen from Figure 1 and across time as can be seen by Figure 2. Furthermore, Figure

⁷Norway is the same area as the United Kingdom and is twice the size of Florida. In 1997, Norway had 114 outlets.

⁸Vinmonopolet itself predicts a further ten to fifteen fold increase in the number of stores under privatisation. Taken from Today's Vinmonopolet - a modern chain with a social responsibility, accessed 13 June 2018 <https://www.vinmonopolet.no/social-responsibility>.

Table 1: Descriptive statistics

	Overall	Q1.2000	Q1.2003	Q1.2006	Q1.2009	Q1.2012	Q1.2015
Sick absence per 10 000 inhabitants in the LMR:							
Sick absence	250.82	231.96	262.62	243.88	271.69	246.3	283.93
	.73	3.20	2.91	2.71	2.69	2.58	2.80
Sick absence (men)	195.95	190.71	216.36	190.91	216.11	185.5	207.49
	.6	2.98	2.70	2.37	2.39	2.02	2.11
Sick absence (women)	306.79	273.86	308.99	296.83	328.7	309.24	362.99
	.96	3.82	3.49	3.28	3.26	3.39	3.80
Alcohol sales (NOK/cap)	743.98	574.17	637.4	757.28	832.54	808.89	787.89
	4.27	17.25	14.80	16.35	17.20	15.53	17.41
Wine (ml/cap)	2640.81	1916.94	2148.28	2534.95	2910.48	3037.51	3035.79
	16.58	62.85	53.79	57.86	63.24	60.55	68.09
Strong wine (ml/cap)	40.66	65.87	54.01	45.6	37.99	29.07	23.88
	.39	2.49	1.63	1.35	1.08	.80	.73
Liquor (ml/cap)	677.53	529.07	656.29	700.89	760.5	713.61	664.06
	3.95	15.91	13.80	15.88	16.51	15.22	15.22
Non-alcoholic (ml/cap)	7.68	4.78	4.13	4.72	5.27	9.92	15.03
	.11	.15	.11	.14	.14	.38	.51
Beer (ml/cap)	55.19	42.97	38.05	35.85	39.46	59.8	109.13
	.90	2.91	2.49	2.39	2.36	3.27	4.71
Average distance (km)	10.86	20.96	13.5	11.51	8.98	7.69	6.67
	.17	1.13	.70	.59	.55	.54	.53
Average grid distance (km)	8.34				9.13	8.38	7.85
	.05				.29	.27	.27
Employment share	.51	.50	.49	.51	.52	.51	.50
	0	.003	.002	.003	.002	.002	.002
Share in working age	.51	.50	.49	.51	.52	.51	.50
	0	.003	.002	.003	.002	.002	.002
Share in working age(men)	.59	.59	.59	.59	.59	.59	.59
	0	.001	.001	.001	.001	.001	.001
Share in working age(women)	.56	.56	.56	.56	.56	.56	.56
	0	.002	.001	.001	.001	.001	.001
Median income(NOK)	151252.14	85582.39	112179.8	131521.48	170526.25	185418.35	194055.47
	882.21	1931.18	1730.61	1959.90	2389.77	2628.62	3040.13
Median bank deposits(NOK)	48984.51	26423.85	33149.57	41481.29	50722.45	60002.31	72836.52
	353.59	630.59	670.15	745.57	887.69	981.39	1201.11

Notes: Mean in first row, standard deviation below. Working age is defined as 20-65. Grid data are only available from 2008-2016. Alcohol expenditure deflated to 2015-values.

3 shows that in 2000 half of the regions had 1 or fewer stores. However most of these regions increased their number of stores. In fact only one region did not receive a new store.

3.2 The store opening process

There is a designated process for the opening of new stores. The new store process can be initiated by either the local municipality or by Vinmonopolet, but a new store can only open with mutual consent. Upon receipt of an application, a municipality is placed upon the decision list. Every autumn the Vinmonopolet board decide upon their new openings. In making their decisions they assess local purchasing power, population data, proximity to the nearest store, whether the proposed location is already a population center and a range of ad-hoc factors, such as seasonal tourism or abstinence cultures.

Every year Vinmonopolet select between five and fifteen municipalities to receive new stores. For example, in 2017 they approved seven new municipalities, from a list of more than a hundred active applicants. For successful locations, Vinmonopolet then formally applies to the municipality for permission to open a location, and if successful, advertises for a place to open. Subsequently Vinmonopolet chooses between bidding locations taking into account characteristics such as proximity to other stores, and the availability of parking and delivery spaces. Vinmonopolet then draws up and applies for approval of the interior plans. Around six to twelve months after the initial Vinmonopolet decision the new store is open.

3.3 Vinmonopolet's objectives

As a State-owned monopolist Vinmonopolet does not seek to maximise profits. Instead it seeks "to secure responsible social control of sales".⁹ In 2016, annual revenues were around 13 billion NOK, and an operating profit of 150 million NOK was split between a dividend and equity. These profits derive from a simple mark up rule. The mark up has a per-liter component and a 22 percent mark up on the pre-tax wholesale price, subject to a cap of 110 NOK per item.¹⁰

3.4 The impact of the expansion on driving distances

As is evident in Figure 4, the expansion of Vinmonopolet reduced average per capita travel distance to the nearest store. The lower line represents average driving distances calculated using municipalities as the basis for our population centers, whereas the upper line uses the grid data. The expansion reduces the

⁹Taken from Today's Vinmonopolet - a modern chain with a social responsibility, accessed 13 June 2018 <https://www.vinmonopolet.no/social-responsibility>.

¹⁰Vinmonopolets purchasing procedures and product range, accessed 20 June 2018 <https://www.vinmonopolet.no/purchasing-and-product-range>. This cap has the effect that, for very expensive wines, consumers may pay less than in other jurisdiction.

Figure 1: Outlet locations in 2000 and 2015




Figure 2: Number of stores over time





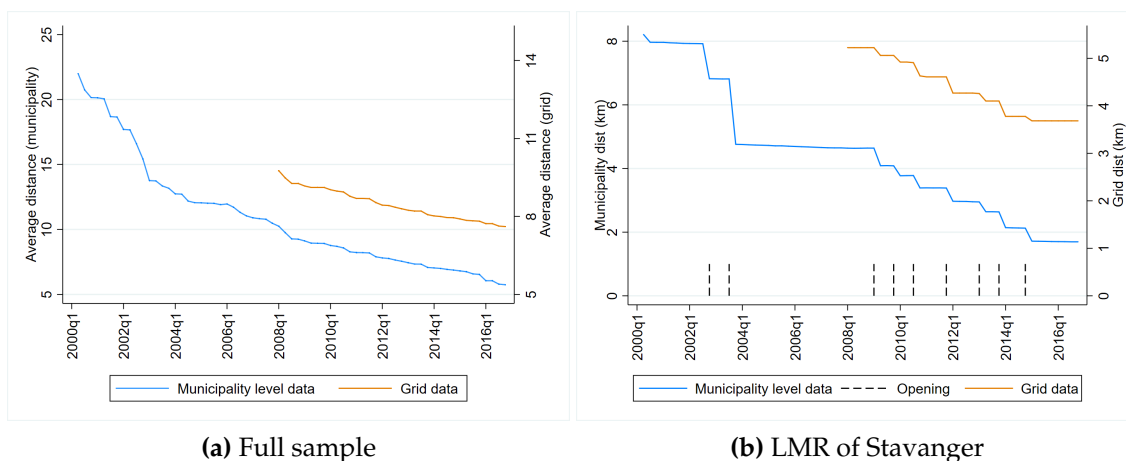
Figure 3: Frequency count of outlets in a LMR



Note: Figure 1 shows the GPS-coordinates of outlets in 2000 and 2015. Figure 2 plots the total number of outlets on a monthly basis over the same time period. Figure 3 counts the number of LMRs by their total number of outlets. In 2000, 22 LMRs had only one store. In 2016, only one LMR had one store.

average traveling distance per capita quite smoothly over time. When we zoom in on a specific region, we see that the reductions are stepwise and almost purely driven by openings. Furthermore, using the municipality level population data captures the same pattern as using the finer grid data, suggesting that population movement is not a key factor. The grid data reports a larger mean since it has non-zero distance for everyone located more than one kilometer away from the store, but we consider that the relative changes in the grid and population data are comparable.

Figure 4: Average distance per capita, population and grid data.



Note: Panel (a) shows how our measures for average distance per capita change over time across all LMRs. Panel (b) shows the impact of store openings in the LMR of Stavanger.

4 Identification strategy

Our goal is to understand how alcohol consumption, measured as quarterly sales per capita in a LMR, affects sick leave. Since alcohol consumption can be affected by sick leave, or be correlated with the error-term in (1), we adopt a two stage IV-approach where alcohol is instrumented by proximity in (2),

$$S_{rt} = \alpha_r + \gamma A_{rt} + \beta X_{rt} + \varepsilon_{rt} \quad (1)$$

$$A_{rt} = \pi_r + \pi_1 AD_{rt} + \delta X_{rt} + u_{rt} \quad (2)$$

where S_{rt} is the number of persons on sick leave per 10 000 inhabitants, A_{rt} , is alcohol consumption, proxied by log transformed revenue from alcohol sales measured in NOK per capita, deflated by the alcohol CPI.¹¹

¹¹The results are almost identical using volume.

AD_{rt} , the per capita average driving distance, is our instrument for alcohol consumption. X_{rt} is our vector of controls.

Assuming that the openings only affect sick leave through changes in alcohol sales, the effect we estimate is causal. While this assumption cannot be tested, there are some notes to be made. Many of the confounding factors are taken care of through inclusion of time fixed effects, and the selection of time and place for new openings contains elements of randomness. While the first stage effect may be affected by selection of places, this should only affect the strength of the relationship between distance reductions and demand, under the assumption that sick leave is only affected by proximity to an outlet through alcohol consumption. This assumption fails if, for example, new stores are located in growing regions that simultaneously experience increased accessibility to doctors, increasing reported sick leave per capita. We are not too concerned by this effect. Due to the restricted nature of the number of stores and Vinmonopolet's decision process, it is likely that size rather than growth determines new store location.¹² This also chimes with the evidence presented above on the store opening process. Furthermore, short term sick leave can often be self-reported, without the need for sign off by a doctor.

We include a LMR fixed effect, π_r , to pick up the fact that these regions differ substantially in their geographical features and a range of other attributes that are likely to affect both sick leave and alcohol demand. Within our control variables, X_{rt} , we always include a vector of quarterly Q and yearly fixed effects y_t . To control for time varying features that might affect sick leave or alcohol consumption we also include controls for the LMR's age composition, employment share and median financial characteristics.

Since previous studies looking at sick leave and alcohol consumption at the aggregate level time differentiate their series to achieve stationarity, we have performed panel unit root tests on our data. These indicate that our panel is stationary, and we control for seasonality and time fixed effects. Still, we include robustness to capture persistence of sick leave over time, and concerns related to underlying trends, by including lags of the dependent variable.

5 Results

As explained in the previous section, we investigate the relationship between alcohol demand and sick leave in the period 2000-2016. Since we use changes in proximity as our instrument for alcohol consumption, it is important to clearly establish the relationship between distance reductions and demand for alcohol before proceeding to the main results. We therefore devote the next section to our first step, to assess the effect of distance on demand.

¹²Recall that the Ministry of Health continues to restrict the number of outlets and Vinmonopolet consider that privatisation would lead to a ten-fold increase in the number of outlets. VM also consider local purchasing power in their decision process.

Table 2: First stage results of distance on log per capita revenue

	(1)	(2)	(3)	(4)	(5)
Linear:					
Average distance	-0.0158*** (0.00301)	-0.0150*** (0.00291)	-0.0149*** (0.00281)	-0.0147*** (0.00285)	-0.0145*** (0.00278)
Quadratic:					
Average distance	-0.0242*** (0.00496)	-0.0240*** (0.00466)	-0.0238*** (0.00477)	-0.0229*** (0.00432)	-0.0224*** (0.00449)
Average distance squared	0.000181 (0.000131)	0.000194 (0.000126)	0.000194 (0.000129)	0.000178 (0.000120)	0.000170 (0.000127)
Linear with grid data :					
Average grid distance	-0.0218*** (0.00508)	-0.0218*** (0.00509)	-0.0224*** (0.00460)	-0.0233*** (0.00487)	-0.0225*** (0.00462)
Employment	No	Yes	Yes	Yes	Yes
Age composition	No	No	Yes	No	Yes
Financial controls	No	No	No	Yes	Yes
Observations	2948	2948	2948	2948	2948

All specifications include LMR fixed effect, year and quarter dummies.

Standard errors are heteroskedasticity robust and clustered at LMR-level.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

5.1 The effect of outlet proximity on alcohol demand

Results from estimating equation (1), with and without the additional controls, are shown in Table 2. Driving time has a negative impact on demand in all specifications. The reduced form linear estimates suggest that if average driving distance per capita is reduced by 1 km, demand increases by approximately 1.45%. Per kilometre, this corresponds to a price of 3,62 NOK. In 2015 this was equivalent to 40 cents/km.¹³ The magnitude is largely comparable to previous findings by [Seim and Waldfogel \(2013\)](#) (39 to 157 cents/km) and [Davis \(2006\)](#) (31 cents/mile).

The inclusion of age composition and economic condition reduces the effect somewhat, indicating that there might be factors driving demand that also affect opening decisions. Controlling for such factors, a large and economically meaningful effect remains.

Since functional form is key in an IV-setting like this, we also allow for a diminishing or increasing effect of distance. The quadratic rows in Table 2 indicate that, within our data, the effect on demand may

¹³NOK and USD in 1. January, 2015

be diminishing in distance. However, the squared distance term is not statistically significant, suggesting that the linear model is reasonable.

The final row in Table 2 shows that using distance measures from the grid data yields comparable results. The point estimates are slightly higher which can be explained by the fact that in the grid data the same store openings lead to smaller decreases in distance.

These results indicate that there is a clear relationship between distance and demand. If the effect on demand is not causally driven by distance changes exogenous to the consumers, but is merely a consequence of locational choices, one could worry that this could also drive the second stage effects. We refer to Figure 4, showing that the distance reductions mainly comes from openings of new stores, and not by more people living in central places. After presenting the main results, we devote a section to showing that the relationship between distance changes and demand is directly related to the openings of new stores.

5.2 The effect of alcohol on sick leave

From the previous section, we know that the distance reduction associated with each opening leads to increased sales. Table 3 shows the results from estimating (1) and (2) on the proportion of the population on sick leave, and broken down by gender. As we move to the right of the table, more controls are included. We also report the coefficients and F-values of the instrument from the first stages. These indicate that the instrument is strong, even when using only the simple linear specification.

The first column suggests that an increase in alcohol of 1 percent increases the number of people on sick leave per 10 000 inhabitants by 0.29. At the mean, this is a 0.12 percent increase in the proportion of people taking sick leave.¹⁴ In columns 2 to 5, we also control for age composition, employment share, and median financial characteristics. Our parameter estimates remain stable, although we lose some precision. Columns 6 and 7 show that when we include the full set of controls, and add up to two lags of the dependent variable, our parameter estimates decrease by approximately half but maintain significance.

Disaggregating our analysis to assess potentially different results across genders is revealing. The point estimates show that when alcohol consumption in a region increases by 1 percent, sick leave amongst men increases by 0.16 percent, at the mean. In our baseline specifications we find no significant effects for women. Furthermore, at the mean, the point estimates suggest an elasticity about half that of men.

In their time series study of the association between alcohol consumption and sick leave in Norway, [Norström and Moan \(2009\)](#) found a significant association for men, but not for women. They found that a 1 percent increase in alcohol consumption was associated with a 0.62 percent increase in sick leave for men. Given the differences in measures, they used the proportion of sickness absence days of all working days,

¹⁴The overall mean of sick absence per 10 000 is 250.82.

Table 3: IV-effect of alcohol sales on sick leave

	Baseline specification					With lags of dep var		
	(1)	(2)	(3)	(4)	(5)	1 lag (6)	2 lags (7)	3 lags (8)
All	29.07* (13.06)	32.89* (14.43)	30.38* (15.26)	34.08* (15.19)	30.92 (15.98)	16.70* (7.982)	15.58* (7.773)	15.33 (8.033)
L1.Sickness absence						0.500*** (0.0202)	0.467*** (0.0190)	0.465*** (0.0193)
L2.Sickness absence							0.0769*** (0.0204)	0.0762*** (0.0204)
L3.Sickness absence								-0.00939 (0.0232)
Men	33.43* (13.41)	34.47* (14.11)	31.74* (15.04)	35.96* (15.04)	33.01* (15.82)	14.41* (7.121)	13.91 (7.220)	14.03 (7.287)
Women	25.69 (14.83)	27.62 (16.06)	28.71 (17.00)	31.27 (16.62)	29.87 (17.87)	20.67* (10.53)	19.54 (10.15)	18.29 (10.45)
First stage:								
Average driving distance	-0.0158*** (0.0029)	-0.0150*** (0.0029)	-0.0151*** (0.0029)	-0.0149*** (0.0029)	-0.0150*** (0.0029)	-0.0141*** (0.0027)	-0.0141*** (0.0027)	-0.0138*** (0.0027)
F-value of instrument	27.75	28.83	27.96	27.83	27.29	27.53	26.95	26.80
Age composition	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Employment share	No	No	Yes	No	Yes	Yes	Yes	Yes
Financial controls	No	No	No	Yes	Yes	Yes	Yes	Yes
Observations	2948	2948	2948	2948	2948	2904	2860	2816

All specifications include LMR fixed effect, year and quarter dummies.

Standard errors are heteroskedasticity robust and clustered at LMR-level.

* p < 0.05, ** p < 0.01, *** p < 0.001

and the potential positive feedback loops between sick leave and alcohol consumption our estimates are in accordance with their findings.

Interestingly, when we add a lag of the dependent variable, the point estimates suggest similar elasticities for both men and women, at the mean. This could suggest that there is some persistence in sick leave that vary across genders. Controlling for such persistence, reduces the first-stage estimates which could imply that sick leave has a positive feedback on alcohol consumption. In eliminating such effects, we also remove the longer term impacts of alcohol consumption on sick leave, which could explain the lower elasticity. Nonetheless, we still find economically and statistically significant results suggesting that alcohol consumption may have both short and long term impacts.

6 Robustness

In this section we explain the two main reasons behind our confidence in the results. First, we perform a synthetic control analysis to underline the validity of our first stage. Second, based on those results we perform additional robustness checks on how we implement our estimation strategy.

6.1 Non-parametric approach

By viewing each opening as a natural experiment, we can perform a comparative case study on each of them. Compared to our reduced form first stage in the previous section, this will give us a cleaner and more direct identification of the effect from an opening, showing that the distance changes drive the effect on demand. Since this is the case, the only remaining threat to identification is that opening a Vinmonopolet impacts sick leave beyond changed alcohol consumption.

We use the synthetic control method of [Abadie, Diamond, and Hainmueller \(2010\)](#), matching LMRs that receive a new outlet with a linear combination of the regions that do not receive a new outlet either one year before or one year after the opening. The post period allows us to estimate the one-year effect on demand from each new opening, which can be used as a validation of the reduced form set up, and further increase our understanding of the first stage effects.

For inference, we follow [Abadie, Diamond, and Hainmueller \(2010\)](#) and produce distributions based on a placebo exercises. For each available control group in each experiment, we produce a synthetic control group and estimate the placebo effect. We then compare the estimated effect of an actual opening to the distribution of placebo effects.

To start the analysis, for each opening, we run the algorithm of [Abadie, Diamond, and Hainmueller \(2010\)](#). We use log per capita revenues as our dependent variable¹⁵ and match on pre-treatment observations

¹⁵We have also tested various residualisation-procedures removing seasonality before matching, yielding similar

of the dependent variable,¹⁶ population, average distance, age composition and income and assets.

From the list of 198 new openings, we are left with 166 events after removing those with an insufficient amount of observations at the beginning and end of the sample. The openings have an average donor pool of 20 LMRs, the smallest being 13 and largest 30. The placebo exercise creates 3 364 additional experiments from which we generate our control distribution.

Given the volume of cases, we cannot display figures of each treatment akin to [Abadie, Diamond, and Hainmueller \(2010\)](#), so we first show the results from estimating the treatment effect on the pooled sample. We simply estimate,

$$A_{rt} = \alpha post_t + \eta T_{rt} + \omega post_t * T_{rt} + \varepsilon_{rt} \quad (3)$$

where $post_t$ indicates the post period, η quantifies the level difference between the control and treatment group in the pre period, and ω is the treatment effect. Column 1 of Table 4 shows the treatment effect of the LMRs who were actually treated, while column 2 shows the placebo effects. There is a clear positive effect of 4.8 percent from an opening, while we see a precisely estimated zero effect in the placebo regions. On average, an opening results in a decrease of 3.6 kilometer, indicating that these results are broadly in line with our reduced form estimates. A linear fit between the effects and the distance reductions have a slope of 0.0133. This implies a marginally smaller linear relationship between distance and demand compared to our baseline specification where we include all control variables (1.50 percent).

We proceed by estimating the effect for each opening (real or placebo) against its synthetic control group, separately. We then evaluate each of the real effects against its distribution of placebo-effects. We do this by computing each opening's rank in the distribution of its placebo effects. If there are 19 placebos and 1 real effect estimated, and the real effect is the largest one, its absolute rank is 20. We then compute the relative rank, meaning we normalize by the size of the control group,

$$\text{relative rank} = \frac{\text{absolute rank}}{N_p + 1}$$

where N_p is number of placebos for that treatment (the size of the control group). If the actual opening has the largest effect relative to its placebos it will have a relative rank of 1, if it's the median it will have a rank of relative rank of around 0.5, and if it's the smallest it will be $\frac{1}{N_p+1}$. This normalization enables us to compare ranks across experiments with different sized placebo groups. Figure 5 shows histograms with 5 percent bins of the relative rankings. 111 of the estimated SCM-effects are above .90 in the treated group, results.

¹⁶Since [Kaul, Klößner, Pfeifer, and Schieler \(2016\)](#) warns against matching on all pre-intervention outcomes, we exclude every other month, as is done in [Garmon \(2017\)](#).

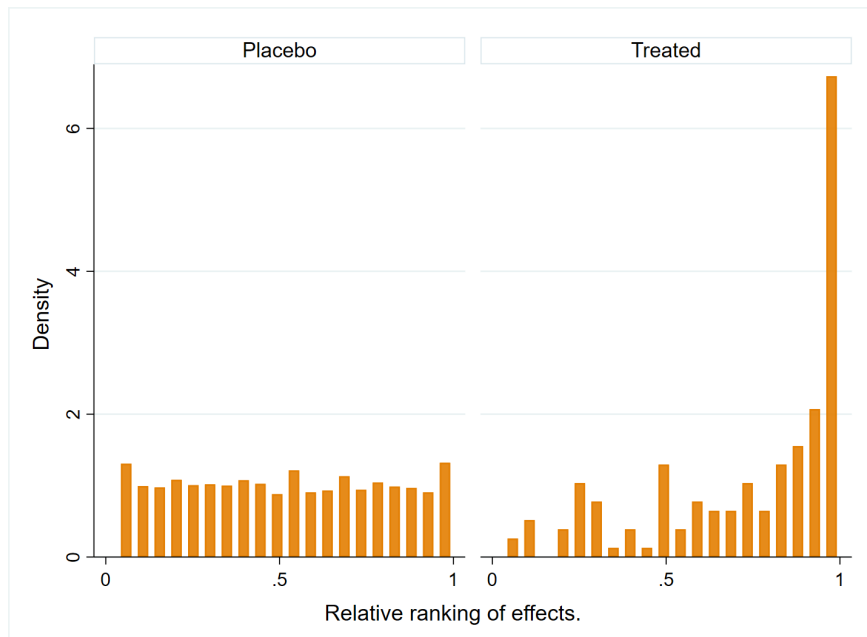
Table 4: Aggregate synthetic control results

	Treated	Placebo treatment
α_{post_t}	-0.0003 (0.002)	0.0013 (0.00297)
ηT_{rt}	0.0159 (0.009)	0.0021 (0.002)
$\omega_{post_t} * T_{rt}$	0.0479 *** (0.0140)	-0.0011 (0.0042)
Observations	62 172	62 172
F	13.52	0.27
R^2	0.0007	0.0000

Note: Results from estimating the average treatment effect on log revenue per capita one year after opening on placebo openings versus real openings. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

while the distribution of the placebo effects is uniform, as expected.¹⁷

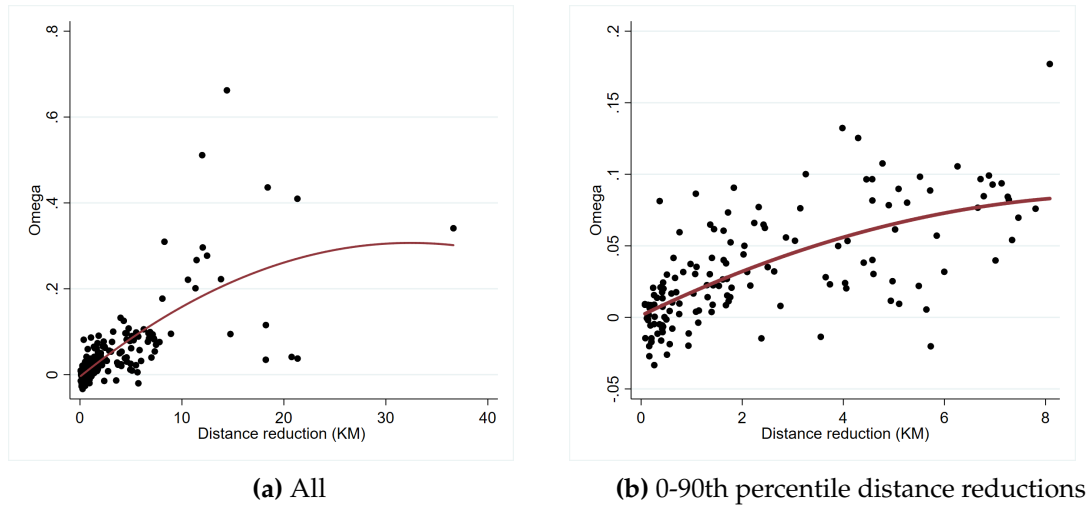
Figure 5: Relative rankings of effects



Note: Histogram showing the distribution of the relative rank, by treatment status. The left panel shows the relative rank of the effect estimated for the placebo openings. The right panel shows the relative rank of the estimated effect from the actual openings.

¹⁷Another way of doing this is presented in [Garmon \(2017\)](#), who plots the effects relative to the placebo effects and look whether its distribution is different from zero.

Figure 6: SCM-effects against distance reduction



Note: Panel (a) plots the average treatment effect (ω from equation 3) of an opening against its associated distance reduction for all openings included in the SCM-analysis with a quadratic line fit on top. Panel (b) plots the same effects as in (a), excluding the largest distance reductions.

Since the functional form the first stage is essential in an IV-estimation, we plot the individual effects of real openings against their respective distance reductions in Figure 6 to assess the shape of the relationship between distance and demand. There is a positive relationship between the two, but the functional form is less clear cut. The implied shape of the quadratic fit is concave, but seems influenced by the largest distance reduction. Therefore, we also include a similar plot including only the distance reductions smaller than the 90th percentile. Also in these less extreme cases we see a weakly concave pattern, but not to the extent that specifying a linear first stage seems inappropriate.

6.2 Estimation variations

Our synthetic control analysis revealed that there might be a non-linear relationship between distance and demand, so we re-run our analysis allowing for a quadratic term. Furthermore, Figure 6 suggests that outliers might be influencing the size of the estimated relationship of distance on alcohol demand. We therefore re-run our analysis based on openings that lead to reductions of distance that are less than 10km per capita.

In light of the weakly concave pattern seen in Figure 6, incorporating a quadratic term relationship between distance and alcohol consumption in the first stage does not lead to significantly different parameter estimates in the second stage. The parameter estimates for men and women are largely similar to the baseline

Table 5: IV-effect of alcohol sales on sick leave, robustness

	Baseline specification				
	(1)	(2)	(3)	(4)	(5)
Quadratic first stage					
All	28.91* (12.81)	32.57* (13.88)	30.28* (14.76)	32.86* (14.85)	29.70 (15.81)
Men	31.49* (13.52)	32.10* (13.97)	29.16 (15.08)	32.89* (15.02)	29.59 (16.10)
Women	27.36 (14.46)	34.60* (15.75)	32.47 (16.72)	34.77* (16.84)	31.09 (17.90)
Observations	2948	2948	2948	2948	2948
Excluding large distance reductions					
All	31.93** (11.80)	36.62** (12.84)	35.30* (13.93)	33.35* (14.04)	31.49* (15.13)
Men	35.32** (11.48)	35.35** (11.34)	33.50** (12.55)	32.95* (13.03)	31.15* (14.08)
Women	29.34* (14.77)	39.36* (17.24)	37.78* (18.62)	35.73 (18.43)	33.07 (19.80)
Observations	2792	2792	2792	2792	2792
Age composition	No	Yes	Yes	Yes	Yes
Employment share	No	No	Yes	No	Yes
Financial controls	No	No	No	Yes	Yes

All specifications include LMR fixed effect, year and quarter dummies. Standard errors are heteroskedasticity robust and clustered at LMR-level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

specifications presented in Section 5.

The second panel of Table 5 shows that when we exclude the largest distance reductions, the precision of the estimates increases. It is reassuring that the relationship we uncover between sick absence and alcohol consumption is not driven by the large distance reductions we exclude, rather they come from openings involving reasonable distance changes.

7 Summary and conclusion

We use the expansion of Norway's State-owned monopolist of high strength alcohol to provide causal evidence on the relationship between alcohol consumption and sick leave. Under the assumption that changes in proximity to a Vinmonopolet store only affects sick leave through alcohol consumption, we show that an increase in alcohol consumption of 1 percent increases the proportion of men who take sick leave in that quarter by 0.16 percent, at the mean. This finding is robust across specifications. In our baseline specifications we find no significant effects for women. Furthermore, at the mean, the point estimates suggest an elasticity about half that of men.

Our analysis has focused on the LMR (Labor Market Region) level, allowing us to fully exploit the regional changes in consumption that arise from new store openings. A drawback is that we do not have sick leave (or employment) data by industry type at the regional level. Given the variability of sick leave (and possibly drinking patterns) by industry, such data would have improved the scope and gains of our analysis. Another interesting avenue for future research would arise if we could increase the precision of our measure of alcohol consumption, for example by having access to a wider range of alcohol purchases. Finally, it would be useful to have individual level residential and working addresses to further improve the precision of our first stage, by allowing us to construct commuting paths.

References

- ABADIE, A., A. DIAMOND, AND J. HAINMUELLER (2010): "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program," *Journal of the American statistical Association*, 105(490), 493–505.
- ASPLUND, M., R. FRIBERG, AND F. WILANDER (2007): "Demand and distance: evidence on cross-border shopping," *Journal of public Economics*, 91(1-2), 141–157.
- BEATTY, T. K., E. R. LARSEN, AND D. E. SOMMERVOLL (2009): "Driven to drink: Sin taxes near a border," *Journal of Health Economics*, 28(6), 1175–1184.
- BHULLER, M. (2009): "Classification of Regional Labour Markets in Norway," *Statistics Norway*.
- BHULLER, M., T. HAVNES, E. LEUVEN, AND M. MOGSTAD (2013): "Broadband internet: An information super-highway to sex crime?," *Review of Economic Studies*, 80(4), 1237–1266.
- BILLMEIER, A., AND T. NANNICINI (2013): "Assessing economic liberalization episodes: A synthetic control approach," *Review of Economics and Statistics*, 95(3), 983–1001.
- CAVALLO, E., S. GALIANI, I. NOY, AND J. PANTANO (2013): "Catastrophic natural disasters and economic growth," *Review of Economics and Statistics*, 95(5), 1549–1561.
- DAVIS, P. (2006): "Spatial competition in retail markets: movie theaters," *The RAND Journal of Economics*, 37(4), 964–982.
- DINKELMAN, T. (2011): "The effects of rural electrification on employment: New evidence from South Africa," *American Economic Review*, 101(7), 3078–3108.
- GARMON, C. (2017): "The accuracy of hospital merger screening methods," *The RAND Journal of Economics*, 48(4), 1068–1102.
- HINNOSAAR, M. (2016): "Time inconsistency and alcohol sales restrictions," *European Economic Review*, 87, 108–131.
- HOTELLING, H. (1929): "Stability in Competition," *The Economic Journal*, 39(153), 41–57.
- HOUDE, J.-F. (2012): "Spatial differentiation and vertical mergers in retail markets for gasoline," *American Economic Review*, 102(5), 2147–82.
- JOHANSSON, E., P. BÖCKERMAN, AND A. UUTELA (2008): "Alcohol consumption and sickness absence: evidence from microdata," *The European Journal of Public Health*, 19(1), 19–22.
- KAUL, A., S. KLÖSSNER, G. PFEIFER, AND M. SCHIELER (2016): "Synthetic control methods: Never use all pre-intervention outcomes as economic predictors," *University of Hohenheim, Department of Economics*.

- McMANUS, B. (2007): "Nonlinear pricing in an oligopoly market: The case of specialty coffee," *The RAND Journal of Economics*, 38(2), 512–532.
- NORSTRÖM, T. (2006): "Per capita alcohol consumption and sickness absence," *Addiction*, 101(10), 1421–1427.
- NORSTRÖM, T., AND I. S. MOAN (2009): "Per capita alcohol consumption and sickness absence in Norway," *The European Journal of Public Health*, 19(4), 383–388.
- PIDD, K. J., J. G. BERRY, A. M. ROCHE, AND J. E. HARRISON (2006): "Estimating the cost of alcohol-related absenteeism in the Australian workforce: the importance of consumption patterns," *Med J Aust*, 185(11-12), 637–41.
- SCHOU, L., AND I. S. MOAN (2016): "Alcohol use–sickness absence association and the moderating role of gender and socioeconomic status: A literature review," *Drug and alcohol review*, 35(2), 158–169.
- SEIM, K., AND J. WALDFOGEL (2013): "Public monopoly and economic efficiency: evidence from the Pennsylvania Liquor Control Board's entry decisions," *American Economic Review*, 103(2), 831–62.
- WEBER, S., AND M. PÉCLAT (2017): "GEOROUTE: Stata module to calculate travel distance and travel time between two addresses or two geographical points," .

CHAPTER V

Paywalls and the demand for online news

Paywalls and the demand for online news*

Frode Skjeret

SNF Centre for Applied Research at NHH

Frode Steen

NHH Norwegian School of Economics and CEPR

and

Timothy G.A. Wyndham

NHH Norwegian School of Economics

This version: March 2019

Abstract

The digitisation of society has posed a challenge to news outlets. Seeking advertising revenues and facing competition for the attention of their readers, many news outlets entered the digital era with unrestricted access to their online content. More recently, news outlets have sought to restrict the amount of content available for free. We quantify the impact of introducing a paywall on the demand for news in Norway. The short-run average impact of a paywall is negative and between 3 and 4%, in the long run the effect increases to between 9 and 11%. We find heterogeneity in the response to paywalls. The largest news outlet within its market experiences larger effects than the other news outlets. After introducing a paywall, the largest news outlets face a long-run reduction in demand between 13 and 15%, as compared to the others who experience a decrease of between 8 and 11%. The timing of introducing a paywall does not seem to affect the demand response very much.

JEL codes: L20, L82, D40

Keywords: Online news, paywalls, business models, two-sided markets

* Frode Skjeret, SNF Centre for Applied Research at NHH, E-mail: Frode.Skjeret@snf.no Frode Steen, Norwegian School of Economics and CEPR, E-mail: Frode.Steen@nhh.no Timothy G.A. Wyndham, Norwegian School of Economics E-mail: Timothy.Wyndham@nhh.no. We are grateful to all the Norwegian news outlets and media houses included in our study for providing access to data on paywall introductions, and, in particular, to TNS Gallup and NRK who provided us with data on online media consumption. We thank Hans Jarle Kind and seminar participants at NHH for comments and discussions. The usual caveats apply.

Introduction

The digital transformation of society has profoundly impacted news producing organisations. Digitisation has lowered the cost of getting content to readers, but has also increased the range of substitutes available to readers and advertisers. News outlets, as two-sided platforms, have always had the option to not charge their readers for news.¹ Most print newspapers have adopted a model where readers pay for news, the free newspapers often available in large cities being the exception.² Online news outlets, on the other hand, have frequently decided not to charge their readers.

In many countries the proportion of news going behind a paywall has increased. We adopt the paywall definition used in Chiou & Tucker (2013): a “digital mechanism that separates free content from paid content on a website”. In 2011, according to Høst (2016), only five out of 194 Norwegian news outlets with an online presence had implemented any sort of paywall.³ By 2015, he found that business models had evolved, and nearly two thirds of news outlets in Norway with an online presence had some sort of paywall.

News outlets have two sources of profits and revenues: readers and advertisers. Changing a business model by introducing a paywall is likely to have opposite effects on these sources. Whilst it will lead to new reader revenues, the number of readers and/or pages viewed is likely to fall, decreasing advertising revenues. If, as is most common, the advertising price is based on number of views, the news outlet will have fewer views and lower revenues.⁴

The business models on how to integrate and price online vs printed news are still not concluded. Though we have seen a trend towards introduction of paywalls across many markets, we also see reversals. In the US, some news outlets have reverted to a non-paywall model (Kim et al. 2018), thus suggesting that knowledge on the effects of paywalls is still developing and the optimal choice of business model is not clear-cut. This is something we also observe in the Norwegian market, where some news providers stick to a business model with no paywall. There is also an example of a paper introducing a paywall early, that later removed the wall.

¹ Casadesus-Masanell and Zhu (2010) provide a more general framework of business models when advertising sponsoring is possible.

² Two examples are the Metro and the Evening Standard in the United Kingdom who offer printed news to readers in urban areas for free (in the morning and the afternoon respectively).

³ Most of our news providers are traditional newspapers that have also developed digital platforms. However, since some of the news providers are either fully digital web based outlets (e.g., Nettavisen), and others are television channels that also operate web based news pages (e.g. NRK), we refer to the group of news providers in our sample as ‘news outlets’.

⁴ Where advertising fees are per ad, as in printed newspapers, fewer expected readers would reduce the attractiveness of ads at that paper, and the ad demand curve would shift down.

The impact on reader demand is clearly a crucial determinant of the profitability of a paywall and therefore the choice of business model. We address two major questions. First, what has been the general quantitative impact of introducing paywalls for the average news outlet in the Norwegian market for online news? Second, is the impact different for news outlets that introduce paywalls after their rivals (timing heterogeneity), and does the effect differ with the outlet's local market position (news outlet heterogeneity)? More particularly, there are reasons to believe that the larger news outlets have a different readership base in their respective markets than the smaller outlets. This might lead to a heterogeneous response from readers faced with paywalls at different sized news outlets. We disentangle the response to the introduction of paywalls and estimate separate effects for the larger news outlets and for the others. We also estimate dynamic models, which allows us to quantify to which degree the effects differ in the short and the long run. To this end, we use longer, more frequent data and analyse a larger number of news outlets that are changing their business model than has previously been done. Distinguishing between short- and long run effects within the same analysis is also new to this literature.

To do this, we utilize weekly data from 122 news producers from January 2012 to December 2015 on the usage of electronic news outlets before and after the introduction of paywalls. Of these, 69 introduced paywalls during our sample period. Furthermore, we contrast these data to the consumption of news from a number of news providers that offered open access to their online content throughout the sample period, including the national public service broadcaster (NRK). NRK has produced online news from dedicated newsrooms tailor-made for all the regions throughout our data period.⁵ Whereas regional news providers without payment walls are often smaller and without a strong presence online, NRK is typically among the largest regional news providers in its respective markets, and is as such a particularly well suited control group to the news providers that introduced paywalls over the period of study.⁶ Given the nature of the Norwegian media topography, we are able to allocate all our news outlets to 13 well-defined markets. These comprise twelve regional markets and one national market for the outlets with a much wider spread.

Analysing a relatively large number of markets enables us to study heterogeneity both with regard to how differences in type of news outlet (relative size in their local markets) affect the introduction of paywalls, and how heterogeneity in the timing of the introduction affects the

⁵ We differentiate between national and regional news providers, applying the regions used by NRK when defining the geographical scope of their district offices. See Figure 1 and Table A1 for more details.

⁶ Note that also some of the news providers that do not impose paywalls are large online providers, examples are e.g., Nettavisen and TV2 that are ranked as number 4 and 5 in the national market (See Table A1 in the Appendix).

demand responses. We apply a difference-in-difference approach, where we look at how the introduction of paywalls affects the number of page hits, unique sessions and unique visitors.

Chiou & Tucker (2013) were the first to empirically investigate the quantitative impact of paywalls on digital news outlet readership. They used state level data from the USA for two periods; before and after the simultaneous introduction of paywalls at three local newspapers.⁷ Employing a difference-in-difference strategy and a control group of 76 similar newspaper owned by the same newspaper group, they estimate that the introduction of the paywall led to a short-term decrease of readership of 51%. Pattabhiramaiah et al (2018) assess the impact of a paywall at the New York Times (NYT). They used five national newspapers to create a synthetic control group. In the thirteen months after the NYT introduced the paywall, the number of unique visitors compared to that of the synthetic control newspaper, fell by 16.8%. They also estimated a reduction in other engagement metrics such as number of visits, pages viewed per visit and duration of visit.

The closest paper to ours is perhaps Kim et al. (2018). They study the rollout of paywalls in 42 newspapers in the US between 2010 and 2015. They analyse the newspapers as a pooled group without being able to attribute them to different regional markets, but rather utilize a rich dataset on individual newspaper characteristics to control for marginal effects of paywall introductions. As opposed to us, they have no control groups that offer free online news. They find that most newspapers' paywalls have long-term negative effects (though they do not estimate short-run effects), but that the amount of the loss varies by reader demographics, newspaper characteristics and when the paywall is introduced. Of their 42 newspapers they find significant decreases in online demand for 36 newspapers. The estimates vary from -54% to -10% with an average decrease of 28.3%. For the remaining six newspapers estimates are positive, but non-significant (calculated from Kim et al. (2018), Table 4). Their dataset has more information on newspaper and readership characteristics than ours, allowing the estimation of a richer set of marginal effects in terms of heterogeneity in consumer responses to paywalls. Our analysis complements theirs in the sense that we can utilise the combination of a number of well-defined markets and the existence of a public, online, freely accessible news provider that has dedicated newsrooms for each regional market. This provides us with better controls throughout the data period and better information on the size and ranking of news outlets within their local markets.

A number of studies have analysed the optimal type of paywall. Lambrecht & Misra (2016) assess the dynamic question of what share of content to place behind the paywall. They find empirical evidence

⁷ Each period contained visit data for four weeks.

showing that news outlets make more news available for free in periods of high demand. Aral & Dhillon (2017) use a series of natural experiments to evaluate the impact of changes to the amount and breadth of news behind a pay wall on cross channel demand and subscription rates.

Our paper is linked to previous qualitative work on the Norwegian roll-out of paywalls. Sjøvaag (2016) found that papers which provide some content free whilst keeping other content behind a paywall are more likely to place content of local relevance behind the paywall and leave more widely relevant news, such as 'wire copy' news and syndicated content open to all. This is in line with Kim et al. (2018) who quantify this and find that newspapers with more unique content tend to perform better after the roll-out than newspapers with more common content. Hognaland and Saebi (2015) investigate qualitative drivers of business model choices (full paywall, partial paywall, freemium etc.) in Norwegian newspapers. An interesting finding for this paper is that experimentation played a central role for news outlets' choice of business model. Hence, knowledge about the effects of introducing paywalls was most likely scarce at the time of introduction.

In our most basic average effect models, we find that the short-run average impact of a paywall is negative and reduces demand by between 3 and 4%, which is smaller than in previous studies. However, the effect is found to be much larger in the long run and when we control for news outlet heterogeneity. Our results suggest that readers' habits take some time to change. The average news outlet experiences between 9 and 11% long-run reduction in demand after a paywall introduction, suggesting that the longer the paywall exists, the stronger the impact from its reader-base will be.

Turning to the relative ranking and size of the news outlets, we find some evidence of heterogeneity in the responses to paywalls. The largest news outlet experiences larger demand effects than the other news outlets within their regional (or national) market. The largest news outlets face a long-run reduction in demand of between 13 and 15% after paywall introductions, as compared to between 8 and 11% decrease in demand for the other outlets.

When estimating the effect of the introduction of a paywall, we control for the share of hits behind other competing news outlets' paywalls. As competing news outlets install paywalls, the share of freely available online news is reduced. If this share is increased by 10% (implying a reduction in freely available news), our models predict a general and significant increase in online consumption for the remaining free news outlets of between 3.6 to 4.4%.

We also analyse to what extent the timing of introducing a paywall affects the demand response. By allowing the effect to differ according to the share of the market that is behind the competitors' paywalls, we find that timing is not very important. However, for the largest news outlets, we find

some indication of an increased negative demand effect as more and more competing news outlets introduce payment walls.

We show that our results are robust both when focusing only on changes around the paywall introductions (event-study) and when we allow for alternative market compositions.

In the next section, we describe our data. We then present our descriptive analysis before describing our empirical strategy. We present our results and robustness analysis before we discuss our results and conclude.

Data and market definitions

We combine data on the usage of electronic news outlets before and after the introduction of paywalls with data on the consumption of news from the regional and national pages of the public service broadcaster, NRK.

Our data on the usage of electronic news outlets is from Kantar TNS. We have removed sites that cannot be regarded as news-media sites, e.g. weekly magazines, special interest group sites and various news aggregators. Our dataset contains only news outlets that actually produce news in-house. We also have data on the usage of NRK's internet sites, both for the nationwide site, and its regional news outlets.⁸ Although our data on private news outlets dates back to 2009, we restrict our data to span from week 1, 2012 to week 52, 2015, in order to maintain consistency with the NRK data. In addition, many of the smaller private news outlets were included in the data from 2013 and onwards.

The dataset from Kantar TNS consists of three measures of weekly internet media consumption:

- Hits: this is the total number of hits (all articles and front-page) from all visitors.
- Unique sessions: Unique sessions are measured as the number of sequences of hits by all unique visitors from first visit to site until leaving the site.
- Unique visitors: All visitors to a news outlet site are uniquely identified and counted.

The three measures are highly correlated; we typically find a correlation coefficient ranging from 0.98-0.99, all measuring online news outlet activity.⁹


⁸ We only have aggregate NRK data for Buskerud, Telemark and Vestfold before week 42, 2013, thus we aggregate these three counties from week 43, 2013 by summing them. This way we have a consistent market for this area (Østafjells) throughout our sample period. We have also checked that there is no apparent shift in the long-run trend for other outlets in the three areas before and after the aggregation takes place. From week 10, 2015 we have separate NRK data for Finnmark and Troms, we thus aggregate these two counties for the last period (week 11, 2015 - week 52, 2015) to obtain a consistent regional market control group.

⁹ The correlation between Hits and Visitors is 0.978, between Hits and Sessions 0.987, and, finally, between Sessions and Visitors 0.994. All are significant at the 1% level.

Our final dataset includes 122 news-producing media-sites. We extend this news outlet level data by adding the dates of introduction of paywalls in the Norwegian market by 69 news outlets. We obtained this data by contacting the individual news outlets, or their owners.

Data on the geographical coverage of the news outlets in the sample is provided by Medietilsynet (the Norwegian Media Authority). We use the NRK district offices as our definition of geographical markets. When the location of a news outlet in our data from Medietilsynet is within the regional boundaries of an NRK district office, the news outlet is assumed to belong to that regional market. Thus, we use the following 12 regional markets: *Hedmark and Oppland, Østafjells, Østlandssendingen, Østfold, Sørlandet, Rogaland, Hordaland, Sogn and Fjordane, Møre and Romsdal, Trøndelag, Nordland, and Troms and Finnmark.*

Figure 1: The regional markets used for our analysis



In addition, we define a national market where all larger, national news outlets are included. Since news providers that are attributed to the national market are larger by orders of magnitude than the regional market providers, allocating these national players to regional markets might lead to biases in our models. However, as a robustness check, we estimate models in which we leave out the national

market, i.e. only estimating models for the 12 regional markets, and our predictions are stable across these models. In Figure 1 we show a map illustrating our 12 regional markets.

Descriptive analysis

While the previous chapter introduced the data used, this chapter describes the dataset in greater detail. Table 1 provides an overview of our markets, number of news outlets and paywalls introduced, the average weekly hits and visitors and Herfindahl-Hirschman concentration Index (HHI) based on hits across markets. A more detailed list providing the names and relevant figures of all outlets is found in Table A1 in Appendix A.

Table 1: Markets, news providers, and usage of data (Week 1, 2012- Week 52, 2015).

Market	Number of news providers	Paywalls introduced	Average weekly hits	Average weekly visitors	HHI
National Market	12	6	32 092 240	4 663 477	2 007
Hedmark and Oppland	6	3	492 595	163 736	3 563
Østfold	7	6	693 226	173 507	2 711
Østlandssendingen	11	8	718 570	233 670	2 584
Østafjells	22	20	1 106 149	344 066	1 614
Sørlandet	7	3	673 722	180 121	3 306
Rogaland	3	2	938 422	238 308	4 173
Hordaland	11	5	1 760 959	419 341	3 722
Sogn and Fjordane	5	4	839 009	144 623	7 798
Møre and Romsdal	10	3	725 251	184 844	2 754
Trøndelag	12	3	1 442 963	337 078	4 195
Nordland	8	5	782 982	242 228	2 769
Troms and Finnmark	8	1	847 941	266 123	3 389
Total	122	69			
Average local markets			918 482	243 970	3 548

Note: Hit and visitor data in this table are aggregates of the weekly average figures in Table A1 in the Appendix. At the market level, the visitor data will be an overestimate since some readers will visit more than one site.

We see that the different markets display different sizes and compositions. News outlets in the national market have substantially more hits and visitors than news outlets in the regional markets. News outlets in the regional markets are reasonably homogenous in terms of weekly hits, but less so in terms of the number of news outlets within each region. We also observe that the number and proportion of paywall introductions differ quite a bit across regional markets. For half of the markets the percentage of news outlets having introduced paywalls is between 43 and 67%, whereas for some markets nearly all news outlets have done so. For Troms and Finnmark, only one of the news outlets

has introduced a paywall. Market concentration is high; measured by HHI using hits per news outlet, the average across the regional markets is as high as 3 548, which is well above competition authorities' 'worrying' threshold. Even the national market has a high concentration index.

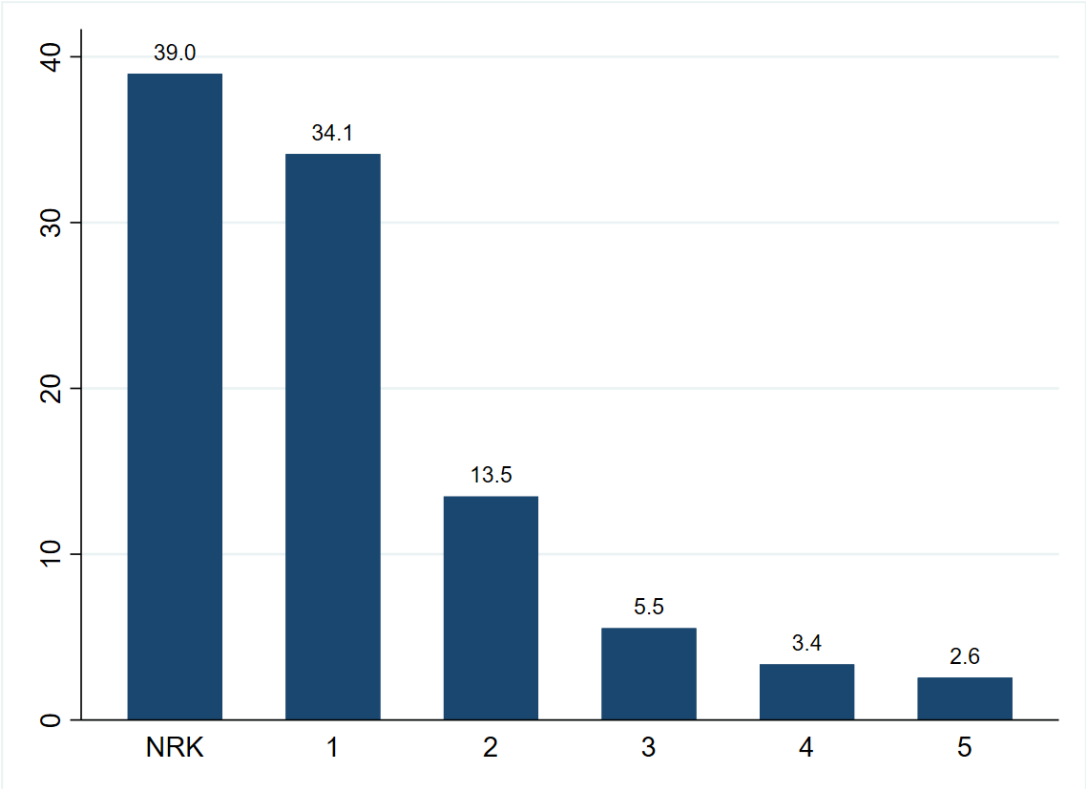
Looking at the data for individual outlets' figures from Table A1 in the Appendix, we observe that in half of the regions, the NRK district page has the highest number of hits. The remaining half consists of the national market, and the regional markets Hordaland, Trøndelag, Rogaland, Sørlandet, Østfold and Østlandssendingen. Using 'unique visitors', it should be noted that only in the national market is a private news outlet (VG.no) larger than NRK. Indeed, the NRK district pages account for 25% of visitors but only 10% of hits. When excluding the national market from the analysis, the NRK regional pages comprise 56 % of visitors on average, but 34 % of hits. In terms of hits per visitor, the NRK district pages receive 2.2 hits per visitor on average, as compared to 3.5 hits per visitor for the regional news outlets. The national market and the NRK national page show a higher ratio, perhaps reflecting a greater breadth of coverage, e.g. all regional news and national common interest topics. In this market, we observe 7.0 hits per visitor per week, on average.

As is evident from Tables 1 and A1, there is great variation in our dataset. First, in the market Østafjells, there are 22 news outlets, while in the market Rogaland, there are only three. The markets also differ when it comes to the relative size of the news outlets. While the number of hits for the smallest news outlet in Rogaland (Haugesund Avis) is slightly less than one third relative to that of the largest outlet (Aftenbladet), the situation is different in the market Østafjells. While NRK Østafjells – the largest outlet – has more than 340 000 hits per week, the 13 smallest outlets each has less than 20 000 hits per week. We also see that the dates of introduction of paywalls vary greatly. While more than half of the news outlets introduced a paywall in 2015, many of the major ones introduced paywalls already in 2012 and 2013 (Fædrelandsvennen, VG, Dagbladet, Aftenbladet, Bergens Tidende, Aftenposten and Agderposten). Hence, while many news outlets introduced a paywall late, a high fraction of the reader-base in our sample had experienced that one or more of their main news outlets introducing a paywall early.

Figure 2 illustrates the average size distribution of media providers. We see that NRK is the most important news source, and their news outlets are still freely accessible. We also see a clear pattern where the largest regional news outlets are substantially larger and have a market share that is two to three times as high as the number two outlets.¹⁰

¹⁰ The numbers are much the same if we exclude the national market: 40.2% (NRK), 34.4% (no.1), 13.2% (no.2), 5.2% (no.3), 3.0% (no.4) and 2.2% (no.5). The major difference in the national market is that here there is somewhat less dispersion between the news outlets. The largest news outlet is a newspaper (*Verdens Gang*)


Figure 2: Market share of NRK and the five largest news outlets, average of market shares in 12 regional markets and the national market (based on number of hits).



This picture suggests that the number one news outlets have a different market position in most of the markets. Later we discuss their content profile as compared to that of the other news providers and analyse whether their consumers respond differently to the introduction of paywalls as compared to those of the others. There is also great variation in the weekly development of hits. In addition, the weekly patterns of hits for news outlets in similar markets display similar movements, see Figure 3 below.

that has 30.9% of the market, whereas NRK is number 2 with a market share of 24.2%. Still the number two newspaper that comes in as number 3 has half the market share of the largest newspaper (16.7%).

Figure 3: Development for selected papers of weekly hits in the national market from 2012 to the end of 2015, in thousands.




Note: The news outlets in the figure are Verdens Gang, NRK, Dagbladet, TV 2, Aftenposten, and Dagens Næringsliv.

There is no apparent trend in the data, although there does appear to be some seasonality. From Figure 3, we can clearly identify a reduction in demand for news during week 52 (Christmas) and during the Easter holidays. There is also a significant drop in demand for news during the summer holidays, in particular during the month of July. In addition to the seasonal variation, there are instances where the national news outlets are strongly correlated, while there are times when this strong correlation breaks down. Hence, visual inspection of Figure 3 indicates that the national news outlets belong to a common market. In particular, we see that the largest news outlets in the national market – Verdens Gang, NRK and Dagbladet – experience similar increases and decreases of demand for news during the period we analyse. However, there is also news-outlet specific variation in the time-series. This is expected, since the news-producing media will create outlet-specific demand. Finally, there does not seem to be a long-term trend when it comes to hits in our dataset, neither increase nor decrease in the variables *hits*, *sessions* and *visitors*.

Turning to the regional markets, we observe the same pattern as in the national market. The seasonal cycles are clearly present, and there are periods when regional markets correlate strongly, particularly

when there are major events taking place in the regions. We give two examples of this below (Figures 4 and 5). We also observe the difference in magnitudes in terms of viewings. Whereas the national news outlet *Verdens Gang* (Figure 3) has around 10 million weekly hits, the largest news outlet in Trøndelag, *Adresseavisen* (Figure 4) has 1 million hits. (See also Table 1A).


Figure 4: Weekly hits for selected papers in the regional market *Trøndelag* from 2012 to the end of 2015, hits in thousands.



We now focus on the regional market Trøndelag in Figure 4. In early 2014, there were a number of wildfires in the western parts of Norway, notably the fires in Lærdal, Flatanger and Frøya. The latter of these, the fire on the island of Frøya, started at midday on January 29th, and due to extremely dry weather and windy conditions, the fire quickly spread across large areas of the island. There was widespread evacuation of people living in areas affected by the fire. Thus, the fire created a strong demand for news, and, as observed in Figure 4, hits at the news outlets in the region of Trøndelag increased strongly. From Figure 4, we also see that there are other periods where the demand for news from NRK Trøndelag and Adresseavisen correlates in the same manner. This is an indication that these news outlets belong to the same market. There are also periods where the smaller news outlets in the region are clearly strongly correlated; however, this is due to differences in scale and therefore not easily observed in Figure 4. The correlation coefficient between NRK Trøndelag and Adresseavisen is about 0.36.

A similar situation occurred in the same month in the region of Sørlandet. In January 2018, the southern parts of Norway experienced a massive snowfall during a short period of time, and the region of Sørlandet was heavily hit. For several days, roads were kept closed due to lack of snow-removal personnel. In addition to the snowfall, the region was hit by a storm, making it difficult to move around. The extreme weather period increased the demand for news about weather forecasts, traffic-operations and general news about the region. This is also observed by visual inspection of Figure 5. Most of the news outlets in the region witnessed all-time high observations of number of weekly hits and number of unique visitors. This was particularly so for *Fædrelandsvennen* – the largest news outlet of the region - and for NRK Sørlandet.

Figure 5: Weekly hits in the Sørlandet-region for selected papers, in thousands. Data period from the start of 2013 to end of 2015.



As seen from the discussion above, there are clear links between news providers' weekly hits both in the national market and in the regional markets. The other important observation is that the online activity across news outlets and markets correlates well with that of the NRK sites. The NRK sites have always been freely available, and have a significant readership and, as such, NRK is a well suited control group when analysing private news outlets introducing paywalls.


This correlation in online activity within markets, together with the fact that there are no apparent differences in trends across markets, supports a common trend assumption in our difference-in-difference modelling. Since we later define the treatment group as all news outlets introducing

paywalls sequentially over time we cannot perform traditional common trend tests here. In the empirical models, we also include a full set of weekly time dummies to account for seasonality due to holidays etc.

To get a first impression of how the introduction of paywalls affects consumption across all our introductions we scrutinize an event-window of ten weeks around all introductions. In Figure 6 we pool all the news outlets in our sample that introduced a paywall in the period of study. We normalize the week of the introduction to zero and look at the average hits before and after the introduction.

Given that introducing paywalls makes it harder to access news from a site, we expect that news consumption will decrease as a result. Figure 6 does indeed suggest that we see fewer hits after the introduction of a paywall.

Figure 6: Average weekly hits 5 weeks before and after the introduction of the paywall, hits measured in logs.



In the next section we present our empirical strategy and models that will be used to analyse our research questions and hypotheses.

Econometric framework and models estimated

Our main research question asks whether introducing a paywall affects the pattern of consumption of news at news-producing media sites. To test for this we apply a difference-in-difference framework.

Main models - average effects across all news outlets

In order to empirically test the relationship between paywalls and news consumption, we estimate the following generic fixed effect model for the logarithm of our time-series of data on hits, unique sessions and visitors, $\ln(x_{i,j,t})$:

$$(1) \ln(x_{i,j,t}) = \alpha_0 + \alpha_1 \ln(x_{i,j,t-1}) + \alpha_2 \ln(x_{i,j,t-2}) + \beta_1 Post_{i,j,t} + \beta_2 Share_{-i,j,t} + \gamma_t + \delta_i + \epsilon_{i,j,t},$$

where subscript i refers to online news outlet, j to regional (or national) market, and t to time period (week). To account for serial correlation in our time series we specify an autoregressive process of second order AR(2), allowing two lags of the dependent variable on the right hand side. The parameter δ_i is the news outlet-specific fixed effect parameter. A fixed effect for all 52 weeks in our dataset, γ_t , accounts for common demand shocks across all news outlets such as holidays. The $\epsilon_{i,t}$, is the standard error-term, anticipated to have the standard properties and being *iid*.

All models will be estimated with robust standard errors where we allow the error term to be clustered for each news provider.

We account for the general effect of the number of news providers in a market that are behind a paywall. The number of news outlets implementing paywalls increases over time. To take account of this, we explicitly consider the paywall share of our regional news markets. The variable ' $Share_{-i,j,t}$ ' measures the share of readers in a market behind competing papers' paywalls. That is, when paper i introduces a paywall, the 'Share'-variable increases with the market share of paper i for all *other* papers. When paper i introduces a paywall, this variable is unaffected for paper i .

The variable $Post_{i,j,t}$ is an indicator-variable taking the value 1 in all periods t after paper i introduces a paywall, and is equal to 0 otherwise. We expect the parameter measuring this effect, β_1 , to be negative, since the paywall introduces restrictions on the news readership. This parameter also serves as the difference-in-difference parameter in our model. Since we apply a fixed effect regression, the treated term used in ordinary regression difference-in-difference models is omitted from the model. All firm-specific differences will be picked up by the δ_i -parameters, and will be omitted due to collinearity if included in a fixed effect regression. In our dataset, treatments do not take place simultaneously. Rather, the introduction of paywalls takes place throughout our sample period. In this respect, we have two types of control groups. First, papers that introduce paywalls late in the sample

period, will act as control group for papers that introduce paywalls early. The second type of control group outlets includes papers that never introduce a paywall, and the NRK internet sites (national and regional). NRK is a significant player both in the regional markets as well as in the national market, with dedicated newsrooms with tailor-made news. NRK is financed via a mandatory licence fee for all households that own a television, and is not allowed to charge for their online news services. These sites will act as control group for all news outlets that introduce a paywall.

In line with the results found by Kim et al. (2018) we expect that online demand response will differ according to several factors, such as reader demographics, news outlet characteristics and the time of introduction of the paywall. Thus, we will extend our model to account for heterogeneity on both timing and news outlet size.


Heterogeneity due to timing of introduction of paywall

The average effect Model (1) accounts for the general effect on consumption of an increasing number of news providers implementing paywalls. However, we do not allow for any changes in the difference-in-difference effect due to this development. We might expect that the first news outlet that introduces a paywall in a particular market will potentially experience a different negative impact than a news outlet that introduces a paywall after a large fraction of news outlets have already done so. The pool of freely available substitutes is decreasing in the share of the market behind a wall. Studies suggest that consumers will be less price sensitive when there is a smaller number of substitutes (*e.g.*, Gumus, Kaminsky, Mathur 2016). Another argument put forward by Kim et al. (2018) relates to the reference price research, which seems to suggest that the more accustomed consumers are to paying, the less price sensitive they become. Or, as in the case of paywalls, where the product was originally provided for free, with paywalls slowly evolving into a new 'normal', where most providers charge for access. This development will potentially affect consumers' product choice (Mazumdar, Raj, and Sinha 2005). News providers that adopt a paywall late are likely to experience a smaller reduction in demand. Both because late movers have a smaller number of substitutes to compete against and because consumers' references have changed. On the other hand, when a news outlet enters late, consumers may already have purchased access to news outlets that have introduced paywalls early, and are thus potentially less likely to purchase additional subscriptions. This effect will work in the opposite direction of those put forward by Kim et al. (2018).

When we look at the data, we see heterogeneity in paywall introductions. For instance, first-moving papers typically have a relatively strong position within their regional market. The largest news outlets (in our dataset) were typically the first outlets to introduce a paywall in eight of twelve regional

markets.¹¹ We see a clear pattern where the largest news outlets introduce paywalls early. This is illustrated in Figure 7. The horizontal axis measures the order of paywall introduction in a market, (1 means first mover), whereas the vertical axis is the average size (in rank form) of that order. Thus, the average rank of the first mover was 3, and for the second mover it was 5. More generally, the graph shows a clear pattern where the larger news outlets move earlier than the smaller ones.

Figure 7: Average rank of regional news outlets vs their timing of introducing a paywall.



Note: Horizontal axis: The order of papers according to when they introduced a paywall. Vertical axis: The average size (measured by rank) of the ordered introductions.

In sum, there might be a selection regarding the timing of introducing paywalls. We thus introduce an additional model where we allow for an interaction between the ‘Share’ and ‘Post’ variables. This variable ($Post_{i,j,t} \cdot Share_{-i,j,t}$) measures whether there is a marginal difference in the paywall effect depending on whether the outlet is an early or a late mover. Thus, we expand (1) and also estimate:

$$(2) \ln(x_{i,j,t}) = \alpha_0 + \alpha_1 \ln(x_{i,j,t-1}) + \alpha_2 \ln(x_{i,j,t-2}) + \beta_1 Post_{i,j,t} + \beta_2 Share_{-i,j,t} + \beta_3 Post_{i,j,t} \cdot Share_{-i,j,t} + \gamma_t + \delta_i + \epsilon_{i,j,t},$$

¹¹ The eight regional markets and the news providers are as follows: Hedmark and Oppland (Oppland Arbeiderblad), Hordaland (Bergens Tidende), Møre and Romsdal (Sunnmørsposten), Sogn and Fjordane (Firda and one other news outlet in the same week), Sørlandet (Fædrelandsvennen), Trøndelag (Adresseavisen) and Østafjells (Drammenstidende) (see Table A1 for details on paywall introductions and news outlets’ size).

where β_3 measures the additional effect of news outlet i introducing a paywall early or late: Early, the share of competing firms already behind a paywall will be low, later it will be higher. For instance, the average effect of introducing a paywall will be given by β_1 in Model (1) and by $(\beta_1 + \beta_3 \cdot \overline{Share_{-i,j,t}})$ in Model (2), where $\overline{Share_{-i,j,t}}$ refers to the mean of $Share_{-i,j,t}$. Since the $Share_{-i,j,t}$ variable differs according to local and national markets, this formulation is more precise in measuring potential time heterogeneity in local demand responses, than simply introducing a time trend across all wall-introductions and markets (which is e.g., what is done in Kim et al. (2018)).

Heterogeneity within news outlets due to readership and news outlet size

The ability to analyse a set of markets also allows us to analyse the effect of news providers' rank and relative size within their relevant markets. In particular, the definition of markets allows us to precisely define which news provider is number one in its relevant market. As we saw above, (see Figure 2) the largest news outlet stands out by being more than double the size of its number two competitor across most markets. This might make the largest news outlets different in terms of what effects to anticipate after the introduction of a paywall.

There seems to be a positive correlation between news providers' size and breadth of content coverage. It has long been common to see news providers differentiate news content to their readers' preferences (Litman and Bridges 1986). Larger news providers thus aim towards covering broader and more general news in order to attract as many consumers as possible. These outlets typically have less unique content and are attractive for their breadth rather than for special news from local areas or local interest groups. This is the case for most of the larger regional news outlets, as well as for the major national news outlets covering the national market. In this respect, one could argue that the larger news outlets have a different readership group than more specialized news outlets.

A number of (often smaller) news providers specialize towards (smaller) more narrowly defined reader groups, either through political slant (e.g., towards more extreme political groups) or through adapting their content to certain groups (e.g., religious news outlets), or simply by having much more local coverage for a smaller population group/area. For instance, Gentzkow and Shapiro (2010) have shown that news outlets respond to their readership's political opinions and tailor their political orientation to differentiate themselves from other papers. Some readers will have a particular interest in local news from e.g., their home municipality, where dedicated and narrower news outlets cover the local market for more in-depth news than regional or national papers are able to. This connection between local coverage and circulation is shown by Lacy and Sohn (1990). In a more recent study, Mitchell, Holcomb, and Page (2015) find evidence that a large majority of readers follow local news very, or somewhat, closely. All this suggests that consumers are less

price-sensitive when a brand, in our case a news outlet, has a more unique positioning, which is also shown by Nagle, Hogan, and Zale (2016).

In general, we therefore anticipate that the degree of content uniqueness might influence the effect of introducing a paywall. This is also in line with Kim et al. (2018) who quantify whether the response to introducing paywalls differs according to content uniqueness. Indeed, they find empirical evidence supporting the idea that more unique content providers tend to perform better after paywall rollout than more general content providers. Previous qualitative work on the Norwegian rollout of paywalls also seems to support this finding. Sjøvaag (2016) found that papers which provide some content for free whilst keeping other content behind a paywall, are more likely to place content of local relevance behind the paywall and leave more widely relevant news, such as syndicated content open to all.

Since the largest news providers within our regional markets and the national market typically have a broader scope for their news coverage, we would anticipate that they observe a larger reduction in hits following the introduction of a paywall than the news providers with a larger proportion of unique content.

We thus expand Model (1) to include an estimate for separate effects for the largest and the other news outlets:

$$(3) \quad \ln(x_{i,j,t}) = \alpha_0 + \alpha_1 \ln(x_{i,j,t-1}) + \alpha_2 \ln(x_{i,j,t-2}) + \beta_1^L Post_{i,j,t}^L + \beta_1^O Post_{i,j,t}^O \\ + \beta_3 Share_{-i,j,t} + \gamma_t + \delta_i + \epsilon_{i,j,t},$$

where superscript 'L' refers to the largest news outlet in market j introducing a paywall, and superscript 'O' refers to all other news outlets introducing paywalls in market j , hence parameters β_1^L and β_1^O quantify the effect of introducing a paywall for the largest news outlets (L) and for the other outlets (O) respectively.

Empirical results

In this section, we present our empirical results. We start out by estimating and discussing our two main models. Then we extend these to allow for potential heterogeneity between the largest and the other news outlets.

Main models - average effects across all news outlets

Model (1) is presented in Table 2. The model perform well, we explain between 34 and 48% of the variation in the data and all parameters come in significant. The autoregressive components, $x_{i,t-1}$ and $x_{i,t-2}$ are highly significant.¹²

Table 2: Difference-in-difference results, logarithm of hits, unique sessions and unique visitors, Model (1), clustered standard errors, all markets.

	Log of hits b/se	Log of sessions b/se	Log of visitors b/se
$\ln(x_{i,j,t-1})$	0.501*** (0.026)	0.404*** (0.025)	0.372*** (0.025)
$\ln(x_{i,j,t-2})$	0.198*** (0.012)	0.213*** (0.016)	0.207*** (0.016)
$Post_{i,j,t}$	-0.035*** (0.010)	-0.040*** (0.010)	-0.040*** (0.010)
$Share_{-i,j,t}$	0.036* (0.019)	0.039** (0.020)	0.044** (0.020)
Long-term effect	-0.117*** (0.030)	-0.104*** (0.022)	-0.094*** (0.021)
Constant	3.402*** (0.312)	3.895*** (0.327)	4.092*** (0.332)
r2	0.475	0.383	0.343
N	18 451	18 451	18 451
Fixed effect; news outlet	Yes	Yes	Yes
Fixed effect; time	Yes	Yes	Yes

* p<0.10, ** p<0.05, *** p<0.01

Across all volume measures we find an adjustment speed of around 30 and 42%, suggesting that when

¹² To determine the lag length we used STATA's varsoc routine for all news providers independent time series on the right side variable (note that this was only doable for 89 news outlets due to sample lengths) and calculated optimal number of lags for the autoregressive process based on the distribution of the AKAIKE information criterion (Akaike, 1974). For the majority of the news outlets the criterion suggested at most two lags for the autoregressive process. See histogram in Figure A1 in Appendix A.

we see changes in volumes in the short run it takes some time before we return to the long run trend.¹³ As expected, the difference-in-difference ‘Post’-variable is highly significant and negative, thus, introducing a paywall entails a reduction in hits for the papers as well as in the number of sessions and unique visitors falls. The average short-term effect of introducing a paywall is relatively small, reducing volumes by 3.5 to 4.0%. This result is considerably lower than previous studies have found, but it is a short-run effect and clearly in line with the expectation that volumes are reduced when a news outlet introduces a paywall.

In addition, the $Share_{-i,j,t}$ variable is positive and significant, suggesting that as more competing news outlets move behind paywalls, volumes generally increase for those that do not introduce paywalls. One implication is thus, that there is a significant increase in consumption of news on free internet media sites after a competing paper introduces a paywall.

We use the log-log framework in our model, hence the parameter estimates are elasticities, and changes could be interpreted as percentage changes, suggesting that a 10% increase in the Share variable increases volumes by between 3.6 to 4.4%.¹⁴

In row 5 of Table 2 we have calculated the long-run elasticity for introducing a paywall.¹⁵ In line with the low adjustment speed the estimated long-term effects are substantially larger than the short-term effects. All long-run effects are significant, and suggest a reduction in traffic between 9 to 12%. The difference between short- and long-run effects is highly significant for all measures.¹⁶ The highest figures are found for number of hits. These figures are more in line with previous findings but remain smaller. Chiou & Tucker (2013) found an effect as high as 51% and Pattabhiramaiah et al. (2018) had an estimate of 17%. However, these two studies include very few news outlets in the treatment group (one and three, respectively) and they have either fewer periods before and after the introduction of paywalls, or less relevant control groups. We look at 69 paywall introductions, and have a larger control group operating in the same markets, and use high frequency data for a longer period. Comparing

¹³ The adjustment speed can be written as $(1-\alpha_1-\alpha_2)$, which means that for e.g., hits we find the adjustment speed to be $1-0.501-0.198=0.301$.

¹⁴ Note that since $Share_{-i,j,t}$ is a percentage variable, a one percentage change is not the same as a one percentage-point increase, e.g., a change in the Share from 0.50 to 0.51 represents a 2% increase when we refer to the elasticity.

¹⁵ In the estimated models the long-run solution is the steady-state solution, implied from equalizing all periods and thus treating variables as equal regardless of their time lag. In our case this implies that we assume that $x_{i,j,t} = x_{i,j,t-1} = x_{i,j,t-2} = x_{i,j}$. Hence, the long-run effect in Model (1) is given as $\beta_1^* = \beta_1 / (1 - \alpha_1 - \alpha_2)$. Since the long-run parameters are non-linear combinations of the short-run parameters we use the “delta method” to approximate standard errors.

¹⁶ The differences between short- and long run are all significant; 0.082 (0.021), 0.064 (0.014) and 0.055 (0.012) for hits, session and visitors respectively. Delta calculated standard errors in parentheses.

these numbers to the findings of Kim et al. (2018), we are in the lower end of their distribution of long run individual news outlet estimates in the range of -54 to -10%.

We now scrutinize the timing heterogeneity by allowing the difference-in-difference effects to differ with the $Share_{-i,j,t}$ variable.

Heterogeneity due to timing of entry of paywall

Model (2) allows for timing heterogeneity in the difference-in-difference effects. The results are tabulated in Table 3. Our results parallel the results discussed above. Most of the joint parameters and explanatory power stay the same. The 'Post' variable is still significant and in the same range as in Model (1). The interactions between 'Post' and 'Share' are not significant however.

Table 3: Difference-in-difference results, logarithm of hits, unique sessions and unique visitors, Model (2), clustered standard errors, all markets.

	Log of hits b/se	Log of sessions b/se	Log of visitors b/se
$\ln(x_{i,j,t-1})$	0.501*** (0.026)	0.404*** (0.025)	0.372*** (0.025)
$\ln(x_{i,j,t-2})$	0.198*** (0.012)	0.213*** (0.016)	0.207*** (0.016)
$Post_{i,j,t}$	-0.037*** (0.013)	-0.039*** (0.010)	-0.041*** (0.012)
$Post_{i,j,t} \cdot Share_{-i,j,t}$	0.003 (0.035)	-0.003 (0.037)	0.004 (0.040)
$Share_{-i,j,t}$	0.035* (0.019)	0.040* (0.020)	0.043** (0.021)
Short-term effect	-0.036*** (0.009)	-0.039*** (0.007)	-0.040*** (0.008)
Long-term effect	-0.119*** (0.028)	-0.103*** (0.018)	-0.096*** (0.018)
Constant	3.403*** (0.311)	3.895*** (0.326)	4.093*** (0.331)
r2	0.475	0.383	0.343
N	18451	18451	18451
Fixed effect; news outlet	Yes	Yes	Yes
Fixed effect; time	Yes	Yes	Yes

* p<0.10, ** p<0.05, *** p<0.01

To measure the effect of introducing the paywall we still need to consider both the Post and the interaction parameter between Post and Share.¹⁷ Since the interaction variable is very small, the

¹⁷ We estimate the effect (elasticity) of introducing the paywall for the average Share, i.e.,

elasticities for introducing a paywall for the competitors' average 'Share' behind the wall are much the same as what we found for Model (1), ranging from 3.6 to 4%, and the difference is not significant.

In row 7 of Table 3, the long-run effects for the difference-in-difference results for Model (2) are tabulated. They are of similar size to those found in Model (1), suggesting only marginally higher long-run elasticities in the order of 10 to 12%.

Hence, when allowing also the difference-in-difference effects to vary with the number of competing news providers behind a paywall, we still find a relatively low short-term negative effect of 4% from introducing a paywall for the average news outlet in our sample. The long-term effect increases somewhat, but is still lower than what has been found in earlier studies.

The elasticities from Model (2) are calculated using the average *Share* of competitors behind a paywall across the sample (=0.210). Since the interaction parameter is very small, and insignificant, the difference between adopting a payment early or late is negligible, and not significant.

We have also estimated a model where we introduce a time trend ('*Trend*') and an interaction between '*Trend*' and our difference-in-difference variable '*Post*'. This specification is closer to the model estimated by Kim et al. (2018). The interaction ($Post_{i,j,t} \cdot Trend_t$) measures whether there is a potential overall linear trend in the effects from introducing paywalls, but now across all paywall introductions across all markets.¹⁸ The results are presented in the Appendix in Table A2. In this trend model we find similar but lower and less precise elasticities for the average news outlet, in both the short- and the long run, suggesting a short-term decrease in demand between 2 to 3% and a long run decrease between 4 to 10%. However, now the whole difference-in-difference effect is picked up by the interaction parameter, which is only significant (at a 5%-level) for the two models. The positive effect from the share of competitors behind the wall found in Model (2) seems to be picked up by the trend-parameter (these variables have a correlation coefficient of 0.596). Since the significant interaction terms come out negative, this suggest that, if anything, the negative effect of introducing paywalls increases linearly over time. Note however, that in our setting with 13 different markets the trend model is biased due to the fact that the sequence of entering payment walls in local markets is

$$\mu_{paywall} = \beta_1 + \beta_3 \cdot \overline{Share}_{-i,j,t}$$

¹⁸ We expand Model (1) to include the $Trend_t$ and the $Post_{i,j,t} \cdot Trend_t$ interaction:

$$\ln(x_{i,j,t}) = \alpha_0 + \alpha_1 \ln(x_{i,j,t-1}) + \alpha_2 \ln(x_{i,j,t-2}) + \beta_1 Post_{i,j,t} + \beta_2 Trend_t + \beta_3 Share_{-i,j,t} + \beta_4 Post_{i,j,t} \cdot Trend_t + \gamma_t + \delta_i + \epsilon_{i,j,t}.$$

The $Trend_t$ takes the value 1 in the first week of observation for all news outlets. The variable increases by (1/52) for each week, taking the value 2 after one year of observations, 3 after three years etc. The parameter β_4 measures the interaction effect; how does the difference-in-difference effect change over time across markets.

not captured through a 'national' trend variable. The trend formulation picks up a more general trend across the whole country, rather than the development in each regional market. Thus, we are reluctant to place too much weight on this formulation of our models.

In sum, timing heterogeneity in the sense of when to introduce a paywall does not seem to be very important in the Norwegian market. As we argue above, Model (2), where we allow for timing heterogeneity and measure the effect through the local market based share of competing news outlets behind a wall, measures timing heterogeneity more precise than a general linear trend across all markets. For instance, the trend model assumes that the introduction of paywalls in the Sørlandet-region in the South of Norway is as important for the local market in the North as if regional news outlets in the North (more than 3000 km away) had introduced paywalls. Thus, we need to be careful when interpreting the trend-model, and conclude that in terms of timing heterogeneity, we believe this effect to be very modest in the Norwegian data.

[Heterogeneity within news outlets due to readership and news outlet size](#)

Generally, Kim et al. (2018) show that the paywall effect differs according to news outlet demographics and readership. We are able to expand on their analysis, using more detailed information on the relative market position of the news outlets in terms of size distribution in their relevant markets. In this subsection, we thus open up for size heterogeneity among the news outlets.

In Models (1) and (2) we have looked at average effects for any news outlet introducing a paywall, but have allowed the difference-in-difference parameter to change as more competing news outlets go behind a paywall. As we discussed above, aside from NRK, the largest news outlets are significantly larger than their competitors within their respective markets. In order to serve a large fraction of consumers within the market, the largest news outlets typically have a more general news content, suggesting heterogeneous demand responses to paywalls for the largest providers as compared to the others. In Model (3), we will allow for this heterogeneity by estimating separate effects for the largest news outlets and the others. The results are shown in Table 4.

In essence, the results are similar to those of Model (1). This comes as no surprise since the only difference from Model (1) is the way in which we divide the *Post*-variable into two separate variables. However, we still find significant difference-in-difference effects for the introduction of the paywall across all models and parameters, but we now observe some signs of heterogeneous responses across groups. For the largest news outlets, the short run effect is a decrease in demand of around 5%, whereas it is around 3% for the others, suggesting that the largest news outlets take a bigger relative

hit when introducing walls. The differences between groups are between 1 and 2 percentage points across the models, but they are not significant.¹⁹

Table 4: Difference-in-difference results for largest and other news providers, Model (3), logarithm of hits, unique sessions and unique visitors, clustered standard errors, all markets.

	Log of hits b/se	Log of sessions b/se	Log of visitors b/se
$\ln(x_{i,j,t-1})$	0.500*** (0.026)	0.404*** (0.025)	0.372*** (0.025)
$\ln(x_{i,j,t-2})$	0.198*** (0.012)	0.213*** (0.016)	0.207*** (0.016)
$Post_{i,j,t}^L$	-0.045*** (0.013)	-0.052*** (0.009)	-0.055*** (0.011)
$Post_{i,j,t}^O$	-0.032*** (0.012)	-0.036*** (0.012)	-0.035*** (0.012)
$Share_{-i,j,t}$	0.034* (0.019)	0.037* (0.020)	0.041** (0.020)
Long-term: Largest	-0.150*** (0.042)	-0.137*** (0.021)	-0.131*** (0.026)
Long-term: Other	-0.107*** (0.037)	-0.094*** (0.028)	-0.083*** (0.026)
Constant	3.405*** (0.311)	3.899*** (0.326)	4.097*** (0.331)
r2	0.475	0.383	0.343
N	18 451	18 451	18 451
Fixed effect; news outlet	Yes	Yes	Yes
Fixed effect; time	Yes	Yes	Yes

* p<0.10, ** p<0.05, *** p<0.01

Turning now to the long-run effects from Model (3), these are tabulated in rows 6 and 7 of Table 4. The long-run negative effects for the largest news outlets are now larger, suggesting a reduction in online activity between 13 and 15%. For the other news providers the figures vary between 8-11%. The differences are thus much larger and vary between 4 to 5%, but are still not statistically significant.²⁰ Hence, when accounting for heterogeneity in size and market rank, we generally see larger negative responses from the demand side following introductions of paywalls.

¹⁹ For hits, sessions and visitors, the differences are -0.013 (0.016), -0.017 (0.012) and -0.020 (0.014) respectively, delta calculated standard errors in parentheses.

²⁰ For hits, sessions and visitors, the differences are -0.044 (0.054), -0.043 (0.033) and -0.048 (0.035) respectively, delta calculated standard errors in parentheses.

To explore both types of heterogeneity simultaneously we also estimate a fourth model where we introduce interaction effects between the difference-in-difference and the share variables for the two groups of news outlets. Thus, this model allows us to also differentiate the timing (share) effect for the largest news outlets and the others:²¹

The results are tabulated together with long-run elasticities in Table A3 in the Appendix. Generally, we find much the same when it comes to the “other” group, and somewhat larger long-run effects for the largest news outlets’ elasticities, that now range from 15 to 17%, as compared to the results for Model 3, which range from 13 to 15%. As in Model (2), the interaction terms are mostly insignificant. We find significance in two cases, for larger news outlets for sessions and visitors, suggesting an increased negative effect as more competing news outlets are behind a wall. As we discussed above, the largest news outlets (in our dataset) in eight of twelve regional markets were typically the first outlets to introduce a paywall (see also footnote 11). Hence, we ascribe this larger and significant effect from the interaction term to be driven by a few news outlets, and we are therefore careful in the interpretation of the results in this combined extended model. Generally, as argued, we typically find only weak evidence for heterogeneity in timing. To the extent that we do find some heterogeneity, it seems that the largest news outlets benefit from being early adopters of payment walls.

Summing up, we find only weak evidence of timing heterogeneity. When looking at heterogeneity within news outlets we find stronger evidence. From Model (3) we find that larger news outlets that typically have a broader news content and a broader readership, will take a bigger hit from introducing a paywall, this is true even when controlling for timing heterogeneity. Our long-run results vary between 8 and 17 percentage points, which is more in line with average estimates found earlier in the literature.

Our dynamic autoregressive distributed lag models will account for likely serial correlation in the left hand side variables. However, we have also performed the Box-Pierce test on the residuals in Model (3) to test for potential autocorrelation in the error terms. Since we have a panel, we undertake tests for our individual news outlets. We can keep the null-hypothesis of no autocorrelation for 87% of the

²¹ The combined model accounting for both types of heterogeneity is thus:

$$\ln(x_{i,j,t}) = \alpha_0 + \alpha_1 \ln(x_{i,j,t-1}) + \alpha_2 \ln(x_{i,j,t-2}) + \beta_1^L Post_{i,j,t}^L + \beta_1^O Post_{i,j,t}^O + \beta_2 Share_{-i,j,t} + \beta_3^L Post_{i,j,t}^L \cdot Share_{-i,j,t} + \beta_3^O Post_{i,j,t}^O \cdot Share_{-i,j,t} + \gamma_t + \delta_i + \epsilon_{i,j,t}$$

The first part of this combined model matches Model (3). However, the last two parameters, β_3^L and β_3^O measure the interaction effect outlined in Model (2), but where we now differentiate between large and small news outlets. For instance, the average effect of introducing a paywall for a large news outlet is β_1^L in Model (3) and by $(\beta_1^L + \beta_3^L \cdot \overline{Share}_{-i,j,t})$ and in this combined model.

news outlets on a 99% significance level.²² In the next Section, we will use Model (3) when exploring the robustness of our results.

Robustness analysis

Here we focus on two issues. First, we look at the time dimension in the sense that we would like to know whether the longevity of the sample is driving our results. That is, one might think that there is a strong underlying time trend that influences our difference-in-difference results since we consider relatively long before and after periods for quite a few of the news outlets, and the variation in these periods could be influenced by underlying long-run trends. Obviously, the inclusion of an autoregressive process as well as a full set of week dummies will account for such effects. However, to make sure that a significant change in demand takes place around the introduction of the paywall, we also undertake an event study looking at only a limited period before and after each news outlets' paywall entry. Second, we look more closely at our market definition. We have so far included the national market as a thirteenth market, thereby treating the national market similarly to the regional markets. Of our 18 451 week-news-provider observations, the national market amounts to 2 049 observations (11%). To investigate the robustness of this assumption, we estimate Model (3) where we exclude the national market. We will also estimate a model where we reduce our control group and take out all NRK sites.

Event analysis

In this part, we define the week in which a news outlet introduced its paywall to be week 0. Then we include all five-week-periods prior to introducing a paywall (the estimation window), and all five-week-periods after the introduction (the event window). We use the average effect Model (1), and apply an identical difference-in-difference structure, but we now drop the competitors' share behind the wall variable and the weekly-indicator variable, and do not include a lag structure.

The results are shown in Table 5. We replicate our main findings from Model (1) in Table 2 even when we look at the ten-week window. We find a significant reduction in demand (*'Post'*) across all three online demand measures. The difference-in-difference effect is now measuring what happens within a five-week span, and, as such, is neither directly comparable to our short-run or our long-run estimates. The estimates in the event-model range between 8 and 12%, which is somewhere between our short- and long-run estimates. Hence, even with a five-week event window we confirm our findings in the much longer main sample.²³

²² For 105 news outlets we keep the null of no autocorrelation of first order on a 99%-level. For 91 we keep the null on a 95% level.

²³ This is the case if we also include the competitors' share behind a payment wall (not reported).

Table 5: Event analysis five weeks before/after introduction of paywall, Difference-in-difference results, logarithm of hits, unique sessions and unique visitors, reduced Model (1) without competitors' share behind payment walls, clustered standard errors, all markets.

	Log of hits b/se	Log of sessions b/se	Log of visitors b/se
$Post_{i,j,t}$	-0.119*** (0.029)	-0.088*** (0.018)	-0.083*** (0.019)
Constant	10.929*** (0.016)	9.581*** (0.010)	9.207*** (0.010)
r2	0.078	0.066	0.053
N	748	748	748

* p<0.10, ** p<0.05, *** p<0.01

An alternative market definition

We re-estimate Model (3) where we include only our 12 regional markets, excluding the potentially different national market. This reduces our estimation sample by 11%. The results are presented in Table 6.

Generally, we find similar results to those shown for the whole sample in Table 4. We lose only around 1 percentage-point of explanatory power, and we mostly keep the significance structure on our parameters. The AR(2) variables come in significant with very similar values as for the whole sample. The same is true for the difference-in-difference variables and elasticities. Only the long-run elasticities for the 'other' group are between one and two percentage points lower in magnitude. The difference-in-difference elasticities lose some significance, but are still significant on a 5%-level. The share variable, though similar in numbers, is no longer significant. Apart from this, we fully confirm the results from Table 4 where we estimate across the whole sample, also including the national market as one of our 13 defined markets.

Table 6: Difference-in-difference results Model (3), logarithm of hits, unique sessions and unique visitors, robust standard errors, all regional markets, national market excluded.

	Log of hits b/se	Log of sessions b/se	Log of visitors b/se
$\ln(x_{i,j,t-1})$	0.496*** (0.027)	0.400*** (0.026)	0.370*** (0.026)
$\ln(x_{i,j,t-2})$	0.195*** (0.013)	0.210*** (0.017)	0.205*** (0.016)
$Post_{i,j,t}^L$	-0.048*** (0.014)	-0.055*** (0.009)	-0.057*** (0.012)
$Post_{i,j,t}^O$	-0.025** (0.012)	-0.030** (0.012)	-0.031** (0.013)
$Share_{-i,j,t}$	0.033 (0.021)	0.034 (0.022)	0.036 (0.022)
Long-term: Largest	-0.155*** (0.044)	-0.140*** (0.022)	-0.135*** (0.027)
Long-term: Other	-0.080** (0.037)	-0.076** (0.030)	-0.073** (0.029)
Constant	3.384*** (0.319)	3.846*** (0.330)	4.007*** (0.333)
r2	0.462	0.368	0.331
N	16 402	16 402	16 402
Fixed effect; news outlet	Yes	Yes	Yes
Fixed effect; time	Yes	Yes	Yes

* p<0.10, ** p<0.05, *** p<0.01

As an additional robustness test of our market definition, we estimate our Model (3) where we leave out NRK regional sites from the control group. Since NRK is publicly funded and might have a different regional and national role, they might differ from the other commercial regional news outlets. Thus, in this model, the control group consists of all other news outlets in our sample. The results are tabulated in Table 7. We find very similar results as we have in Table 4, the explanation power increases for all three models, and all our major predictions come through.

Table 7: Difference-in-difference results Model (3), logarithm of hits, unique sessions and unique visitors, robust standard errors, all regional markets, NRK-sites excluded.

	Log of hits b/se	Log of sessions b/se	Log of visitors b/se
$\ln(x_{i,j,t-1})$	0.545*** (0.025)	0.436*** (0.032)	0.393*** (0.033)
$\ln(x_{i,j,t-2})$	0.197*** (0.014)	0.231*** (0.016)	0.223*** (0.016)
$Post_{i,j,t}^L$	-0.038*** (0.011)	-0.043*** (0.007)	-0.047*** (0.010)
$Post_{i,j,t}^O$	-0.028*** (0.010)	-0.025*** (0.009)	-0.025** (0.010)
$Share_{-i,j,t}$	0.030 (0.019)	0.013 (0.018)	0.017 (0.020)
Long-term: Largest	-0.148*** (0.044)	-0.128*** (0.022)	-0.122*** (0.026)
Long-term: Other	-0.110*** (0.037)	-0.076*** (0.026)	-0.064*** (0.024)
Constant	2.879*** (0.221)	3.287*** (0.275)	3.615*** (0.343)
r2	0.544	0.459	0.402
N	15 991	15 991	15 991
Fixed effect; news outlet	Yes	Yes	Yes
Fixed effect; time	Yes	Yes	Yes

* p<0.10, ** p<0.05, *** p<0.01

Summing up our robustness section, our results do not change much, neither when we focus on an event-window around the paywall introductions, nor when we scrutinize the market definition. The national market does not appear to be very different from the set of regional markets, and changing the control group by excluding NRK does not change our conclusions.

Discussion and conclusions

The current paper analyses the impact of introducing paywalls on the demand for online news. We estimate models for number of hits, unique sessions and unique visitors. Applying a difference-in-difference framework, we utilize very detailed weekly data across 122 news outlets from January 2012 to December 2015 on the usage of electronic news outlets before and after the introduction of paywalls. Of these, 69 introduced paywalls during our sample period. We contrast this to the consumption of news from a number of news providers that offered open access to their online

content throughout the sample period. The national public service broadcaster (NRK) was providing freely attainable news throughout the whole period. Given the Norwegian media topography, we are able to allocate all our news outlets into 13 well defined local (and a national) market(s). NRK was providing online news tailor made for these regions throughout our data period.

Using a dynamic autoregressive distributed lag framework we are also able to provide both short and long run results on the effects of introducing payment walls.

Analysing a relatively large number of markets enables us to study heterogeneity with regard to both how differences in type of news outlet (relative size in their local markets) affect the paywall introductions, and how heterogeneity in the timing of paywall introductions affects demand responses.

We find that the short run average impact of a paywall introduction on the number of hits is negative and, between 3 and 4%, which is smaller than in previous studies. However, the effect is found to be much higher in the long run and when we control for news outlet heterogeneity. After a paywall introduction, the long-run reduction in demand for the average news outlet is between 9 and 11%. This difference between short- and long run impact may seem surprising at first glance. We find that the longer the paywall exists, the stronger the impact from its reader base will be. Readers' habits take some time to change. In our model, the short run is defined as only a couple of weeks, and though readers will spend less time visiting the news outlet, it will take some time before a new long-run equilibrium is reached. First, knowledge about suitable alternatives will increase as time passes from the introduction of the paywall. While some readers will emigrate to alternative news outlets immediately, others are less knowledgeable about alternatives, and will emigrate later. Second, in some instances, the paywall was extended over time: That is, the first weeks after a news outlet introduced a paywall, a fairly small fraction of articles published was located behind the paywall. As time went by, this fraction increased, and readers without a subscription were forced to reduce their consumption of news on this news outlet. Thus, both demand (learning) and supply (paywall-introduction structure) can explain the higher long run responses in demand. This is also supported by the findings in our event study, where the effects across a five-week window are higher than our estimated short run effects, but lower than our estimated long run effects.

Turning to the relative ranking and size of the news outlets, we also find some evidence for heterogeneity in the paywall responses. The largest news outlet within its regional (national) market experiences larger effects than the other news outlets. The largest news outlets face a long-run reduction in demand between 13 and 15% after paywall introductions, as compared to the others that experience between 8 and 11% decrease in demand.

As competing news outlets install paywalls, the share of freely available online news is reduced. Thus, we control for the competitors' share of hits, unique sessions and unique visitors when estimating the effect of paywall introductions. As the percentage of news content behind the competitors' paywall increases by 10%, we find a general and significant increase in online consumption between 3.6 to 4.4%.

We also analyse the extent to which the timing of introducing a paywall affects demand. By allowing the effect to differ according to the amount of the market that is behind the competitors' paywalls, we generally find that the effect does not change significantly. However, for the largest news outlets, we find some indication of an increased negative demand effect as more and more competing news outlets introduce payment walls.

We show that our results are robust to autocorrelations, and we also replicate our results in an event study where we only focus on the five weeks before and after the introduction of paywalls.

Our results seem to suggest that paywalls do indeed reduce demand, but to a lesser extent than what is found in other studies. Compared to these, our study has more detailed data over a longer period of time for many more news outlets and paywall introductions. One weakness in our data is that heterogeneity beyond relative size and timing is not controlled for beyond fixed effects. A future study should aim to also include additional information on readership and reader demography.

References

- Aral, S., and Dhillon, P. (2017). Digital Paywall Design: Effects on Subscription Rates & Cross-Channel Demand. Working paper.
- Casadesus-Masanell, R., and Zhu, F. (2010). Strategies to fight ad-sponsored rivals. *Management Science*, 56(9), 1484-1499.
- Chiou, L., and Tucker, C. (2013). Paywalls and the demand for news. *Information Economics and Policy*, 25(2), 61-69.
- Genzcow, M., and Shapiro J.M. (2010) .What drives media slant? Evidence from U.S. daily newspapers. *Econometrica*, 78(1):35-71.
- Gumus, M., Kaminsky P., and Mathur S. (2016). The impact of product substitution and retail capacity on the timing and depth of price promotions: Theory and evidence. *International Journal of Production Research*, 54(7), 2108-2135.
- Hognaland, I., and Saebi, T. (2015). Business Model Selection and Innovation: An Empirical Investigation of the Norwegian Newspaper Industry. Beta. *Scandinavian Journal of Business Research*, 29(01), 50-73.
- Høst, S. (2016). Papiravisene lever, og de satser digital. Status ved utgangen av 2015. Høgskulen i Volda (Volda University College).
- Kim, H., Song, R., and Kim, Y. (2018). Free-to-Paid Transition of Online Newspapers: The Impact of Paywalls on Website Pageviews. Working Paper, October 2018
- Lambrecht, A., and Misra, K. (2016). Fee or Free: When Should Firms Charge for Online Content?. *Management Science*, 63(4), 1150-1165.
- Lacy, S. and Sohn A.B. (1990). Correlations of newspaper content with circulation in the suburbs: A case study. *Journalism & Mass Communication Quarterly*, 67(4): 785-793.
- Litman, B.R. and Bridges J. (1986). An economic analysis of daily newspaper performance. *Newspaper Research Journal*, 7, 9-26.
- Mazumdar, T., Raj S.P., and Sinha I. (2005). Reference price research: Review and propositions. *Journal of Marketing* 69(4), 84-102.
- Mitchell, A., Holcomb J., and Page D. (2015). Local news in a digital age. Pew Research Center.
- Nagle, T.T., Hogan J.E., and Zale J. (2016). *The Strategy and Tactics of Pricing: A Guide to Growing More Profitably*, 5th ed. (Routledge, New York).
- Pattabhiramaiah, A., Sriram, S., and Manchanda, P. (2018). Paywalls: Monetizing Online Content. Georgia Tech Scheller College of Business Research Paper No. 17-8.
- Sjøvaag, H. (2016). Introducing the Paywall: A case study of content changes in three online newspapers. *Journalism Practice*, 10(3), 304-322.

Appendix

Table A1: News outlets and their regional markets, averages 2011-2015.

Paper	Introduction		Hits	Visitors	Market Share
	Year	Week			
National Market					
Verdens Gang	2012	24	9 927 362	1 338 830	30.93
NRK			7 742 327	996 398	24.13
Dagbladet	2013	11	5 368 873	807 286	16.73
Nettavisen			2 907 840	504 571	9.06
TV 2			2 345 140	442 629	7.31
Aftenposten	2013	47	1 950 088	357 144	6.08
Hegnar Online			1 048 468	70 514	3.27
Dagens Næringsliv	2014	11	710 115	123 562	2.21
Nationen	2014	13	49 040	10 930	0.15
Vårt Land			42 041	11 206	0.13
Morgenbladet			900	390	0.00
Klassekampen	2015	38	46	17	0.00
Hedmark and Oppland					
NRK Hedmark og Oppland			270 921	121 387	55.00
Oppland Arbeiderblad	2015	5	76 340	14 123	15.50
Østlendingen	2015	16	62 111	13 746	12.61
Glåmdalen	2015	17	54 005	8 730	10.96
Avisen Hadeland			17 200	2 859	3.49
Avisa Valdres			12 018	2 891	2.44
Østfold					
NRK Østfold			210 621	106 363	30.38
Fredrikstad Blad	2015	17	270 934	31 858	39.08
Moss Avis	2015	16	80 999	14 179	11.68
Sarpsborg Arbeiderblad	2015	17	62 419	9 489	9.00
Halden Arbeiderblad	2015	16	37 766	5 828	5.45
Smaalenenes Avis	2015	16	24 548	4 814	3.54
Rakkestad Avis	2015	24	5 939	976	0.86
Østlandssendingen					
NRK Østlandssendingen			228 214	128 633	31.76
Romerikes Blad	2015	16	250 956	43 662	34.92
Budstikka.no			120 486	27 808	16.77
Østlandets Blad	2015	19	47 621	10 032	6.63
Dagsavisen	2014	46	36 074	15 457	5.02
Akershus Amtstidende	2015	19	10 976	2 279	1.53
Indre Akershus Blad	2015	16	7 058	1 852	0.98
Eidvoll Ullensaker Blad			6 658	2 120	0.93
Vestby Avis	2015	25	5 866	968	0.82
Enebakk Avis	2015	25	2 487	398	0.35
Ås Avis	2015	25	2 174	461	0.30
Østajells					
NRK Østajells			346 089	197 086	31.29
Drammens Tidende	2014	38	186 134	35 859	16.83
Tønsbergs Blad	2014	39	141 894	29 856	12.83
Telemarksavisa	2015	35	101 697	19 739	9.19
Østlandsposten	2015	6	61 821	9 991	5.59

Varden	2014	46	52 603	10 590	4.76
Sandefjords Blad	2014	38	45 022	7 295	4.07
Ringerikes Blad	2015	5	40 497	6 687	3.66
Lågendalsposten	2014	38	32 226	6 084	2.91
Porsgrunns Dagblad	2015	24	19 286	3 551	1.74
Hallingdølen			16 197	3 458	1.46
Gjengangeren	2015	6	16 115	3 308	1.46
Bygdeposten	2015	6	12 377	2 498	1.12
Telen	2015	6	10 413	2 239	0.94
Kragerø Blad Vestmar	2015	4	7 694	1 895	0.70
Røyken og Hurums Avis	2015	24	5 585	1 460	0.50
Jarlsberg Avis	2015	6	3 574	748	0.32
Svelviksposten	2015	25	1 804	456	0.16
Llierposten	2015	25	1 766	484	0.16
Øyene	2015	25	1 666	279	0.15
Sande Avis	2015	25	1 289	363	0.12
Eiker Avis	2015	24	400	140	0.04
Sørlandet					
NRK Sørlandet			219 366	100 365	32.56
Fædrelandsvennen	2012	11	304 193	50 285	45.15
Agderposten	2013	50	92 429	17 262	13.72
Lindesnes Avis			18 072	3 962	2.68
Llister24			18 035	4 658	2.68
Tvedestrandsposten	2015	24	11 122	2 020	1.65
Aust-Agder Blad			10 505	1 569	1.56
Rogaland					
NRK Rogaland			237 651	119 350	25.32
Aftenbladet	2013	21	531 353	90 499	56.62
Haugesunds Avis	2014	38	169 418	28 459	18.05
Hordaland					
NRK Hordaland			293 648	159 309	16.68
Bergens Tidende	2013	42	882 458	157 262	50.11
Bergensavisen	2015	49	537 260	90 110	30.51
Dagen			16 043	5 029	0.91
Kvinnheringen	2015	24	10 568	2 240	0.60
Hardanger Folkeblad	2015	24	9 580	2 140	0.54
Avisa Nordhordland	2015	24	7 757	1 814	0.44
Strilen			1 050	380	0.06
Bygdanytt			975	449	0.06
Vestnytt			904	363	0.05
Askøyværingen			716	245	0.04
Sogn and Fjordane					
NRK Sogn og Fjordane			737 691	126 390	87.92
Firda	2014	39	64 668	11 549	7.71
Firdaposten	2014	39	21 385	3 164	2.55
Fjordens Tidende	2015	19	8 494	2 036	1.01
Fjordingen	2015	19	6 771	1 484	0.81
Møre and Romsdal					
NRK Møre og Romsdal			282 274	107 948	38.92
Sunnmørsposten	2014	45	207 099	36 422	28.56
Romsdals Budstikke			135 233	20 189	18.65
Tidens Krav	2015	4	60 307	10 664	8.32
Vikebladet-Vestposten			12 235	2 625	1.69
Aura Avis	2015	24	7 608	1 478	1.05

Driva			7 112	1 932	0.98
Åndalsnes Avis			5 855	1 424	0.81
Mørenytt			5 770	1 564	0.80
Sunnmøringen			1 758	598	0.24
Trøndelag					
NRK Trøndelag			323 871	149 487	22.44
Adresseavisen	2014	25	868 379	136 861	60.18
Trønderavisa	2015	2	101 495	23 529	7.03
Namdalsavisa	2015	3	49 189	7 954	3.41
Stjørdalens Blad			30 947	4 553	2.14
Fosnafolket			16 952	3 456	1.17
Avisa Sørtrønderlag			14 224	3 279	0.99
Hitra Frøya			13 986	2 670	0.97
Trønderbladet			8 896	2 220	0.62
Opdalingen			8 643	1 686	0.60
Malvikbladet			4 242	840	0.29
Innherred			2 139	543	0.15
Nordland					
NRK Nordland			358 855	177 717	45.83
Avisa Nordland	2015	8	170 177	26 814	21.73
Lofotposten	2014	45	58 779	6 888	7.51
Rana Blad	2015	6	57 050	8 194	7.29
Fremover	2015	3	45 676	7 390	5.83
Vesterålen online	2014	50	44 678	6 519	5.71
Helgeland Arbeiderblad			30 520	5 680	3.90
Brønnøysunds Avis			17 247	3 026	2.20
Troms and Finnmark					
NRK Troms og Finnmark			388 414	174 204	45.81
Nordlys			294 425	57 590	34.72
itromsø			60 230	12 358	7.10
Harstad Tidende	2015	10	32 538	6 305	3.84
Altaposten			23 237	5 436	2.74
iFinnmark			19 977	3 794	2.36
Troms Folkeblad			18 769	3 966	2.21
Framtid i Nord			10 351	2 470	1.22

Note: Market shares calculated from number of hits

na

Table A2: Difference-in-difference results (Model 2), logarithm of hits, unique sessions and unique visitors, clustered standard errors, all markets.

	Log of hits b/se	Log of sessions b/se	Log of visitors b/se
$\ln(x_{i,j,t-1})$	0.500*** (0.026)	0.400*** (0.026)	0.367*** (0.026)
$\ln(x_{i,j,t-2})$	0.198*** (0.012)	0.209*** (0.017)	0.202*** (0.016)
$Post_{i,j,t}$	-0.014 (0.034)	0.016 (0.029)	0.035 (0.029)
$Trend_t$	0.010** (0.005)	0.019*** (0.007)	0.022*** (0.007)
$Post_{i,j,t} \cdot Trend_t$	-0.007 (0.008)	-0.016** (0.008)	-0.021** (0.008)
$Share_{-i,j,t}$	0.010 (0.022)	-0.008 (0.023)	-0.010 (0.024)
Constant	3.385*** (0.313)	3.926*** (0.340)	4.132*** (0.346)
Short-run effect	-0.030* (0.016)	-0.025* (0.013)	-0.018 (0.012)
Long-run effect	-0.101* (0.052)	-0.063* (0.033)	-0.041 (0.029)
r2	0.475	0.386	0.347
N	18 451	18 451	18 451
Fixed effect; news outlet	Yes	Yes	Yes
Fixed effect; time	Yes	Yes	Yes


* p<0.10, ** p<0.05, *** p<0.01

Table A3: Difference-in-difference results Model (3), logarithm of hits, unique sessions and unique visitors, clustered standard errors, all markets.

	Log of hits b/se	Log of sessions b/se	Log of visitors b/se
$\ln(x_{i,j,t-1})$	0.500*** (0.026)	0.404*** (0.025)	0.372*** (0.025)
$\ln(x_{i,j,t-2})$	0.198*** (0.012)	0.212*** (0.016)	0.206*** (0.016)
$Post_{i,j,t}^L$	-0.027** (0.012)	-0.034*** (0.008)	-0.034*** (0.012)
$Post_{i,j,t}^O$	-0.037 (0.022)	-0.032* (0.019)	-0.035* (0.020)
$Post_{i,j,t}^L \cdot Share_{-i,j,t}$	-0.115 (0.071)	-0.121*** (0.039)	-0.136*** (0.047)
$Post_{i,j,t}^O \cdot Share_{-i,j,t}$	0.008 (0.050)	-0.011 (0.052)	-0.001 (0.054)
$Share_{-i,j,t}$	0.036* (0.020)	0.041** (0.020)	0.045** (0.021)
Short-term; Largest	-0.051*** (0.015)	-0.059*** (0.010)	-0.063*** (0.012)
Short-term; Other	-0.035** (0.014)	-0.035*** (0.011)	-0.036*** (0.011)
Long-term; Largest	-0.170*** (0.049)	-0.154*** (0.024)	-0.149*** (0.027)
Long-term; Other	-0.116** (0.046)	-0.090*** (0.030)	-0.085*** (0.028)
Constant	3.411*** (0.310)	3.905*** (0.326)	4.104*** (0.330)
r2	0.475	0.383	0.343
N	18 451	18 451	18 451
Fixed effect; news outlet	Yes	Yes	Yes
Fixed effect; time	Yes	Yes	Yes

* p<0.10, ** p<0.05, *** p<0.01

Figure A1: Histogram of VARSOC lag suggestions for the autoregressive process based on Akaike information criterion



NHH is one of the leading business schools in Scandinavia. Over 3,000 students study across a range of Bachelor, Master and PhD programmes.

NHH has a long reputation for its high academic level and contributions to the international research community. A large number of our faculty hold a PhD from institutions outside of Norway, in particular top US universities. This creates a diverse and stimulating academic environment.

The PhD student body is made up of around 100 men and women working within different fields of specialisation. The programme encourages close interaction between students and faculty in a social/academic climate where students are regarded as junior colleagues.

The PhD programme offers courses over a wide range of topics within six specialisations: Accounting; Economics; Finance; Management Science; Professional and Intercultural Communication; and Strategy and Management. The programme aims at giving the graduate a solid training in performing high quality scientific research in these areas, making use of state of the art empirical and theoretical techniques. This prepares the student for employment in national and international policy institutions, within research centres, business enterprises, and for the international academic job market. The entire programme is held in English. It runs over three years, with the first year consisting primarily of course work. The next two years are then devoted to independent research and the writing of a doctoral thesis, under the supervision of an advisor appointed from the NHH faculty.

NHH

