Oddmund Berg

Empirical essays on the dynamics of consumption and saving

August 6, 2018

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

Acknowledgments

This thesis was written while I was employed as a PhD research scholar at Norwegian School of Economics. I am forever grateful to Gernot Doppelhofer for supervising both my master's and PhD thesis. You have been encouraging, challenging, supportive and easy to talk to.

I thank Kjetil Storesletten for being a friendly co-supervisor with opinions that greatly improved my work. I thank Morten Ravn for hosting my four month stay at University College London. Thanks to the Department of Economics, Skaugfondet and Wilhelm Keilhaus' Minnefond for providing funding for it.

I would like to express my sincere gratitude towards Elin Halvorsen who first opened the door for me at Statistics Norway in my master's thesis. Without your support, encouragement and co-authorship, there would not have been any thesis. I have truly enjoyed working with you, and look forward to becoming your colleague.

I would like to thank Tim Wyndham for being a fantastic co-author and colleague throughout all four years. Lets keep having full-day Skype sessions in the future, you taught me a lot.

I am grateful to Andreas Fagereng for welcoming me at Statistics Norway, answering my questions and being as helpful, encouraging and friendly as only you can be.

I thank Martin B. Holm and Thor Andreas Aursland for interesting discussions and eye-opening conversations. I hope to have many more of those in the future.

In my time writing the thesis, I have done a great deal of traveling back and forth between Bergen and Oslo. I would like to thank everyone at the Department of Economics in Bergen and at Statistics Norway in Oslo, for being welcoming, helpful and friendly. I feel at home no matter where I go. Traveling would not have been possible without the ones who have lent me a spare bedroom - especially Fredrik, who never said no and paid for it in full.

I am grateful for having a supportive and curious family. A special thanks to my parents for always supporting me the best way they can, showing me how to work hard, be patient and make concious choices. Indirectly, you have been more helpful to me in this work than you think.

There is no way to express my full gratitude towards my partner, Ina. No economist in the world has a higher alternative cost of being in the office, than me.

Contents

Introduction

Introduction

The study of economics does not seem to require any specialized gifts of an unusually high order. Is it not, intellectually regarded, a very easy subject compared with the higher branches of philosophy and pure science? Yet good, or even competent, economists are the rarest of birds. An easy subject, at which very few excel! The paradox finds its explanation, perhaps, in that the master-economist must possess a rare combination of gifts. He must reach a high standard in several different directions and must combine talents not often found together. He must be mathematician, historian, statesman, philosopher – in some degree. He must understand symbols and speak in words. He must contemplate the particular in terms of the general, and touch abstract and concrete in the same flight of thought. He must study the present in the light of the past for the purposes of the future.

John Maynard Keynes (Keynes (1924), pp.322 – 3)

Economics is a social science. In contrast to natural sciences, who study concrete objects and the laws of nature, social science is the study of human behavior, society and social relationships. In working on the current thesis, this has been a recurring, important and challenging theme. Human beings have free will, are creative, and respond to actions of other human beings. This makes it difficult to identify cause and effect, with an interpretation that is sensible and has external validity. Thus, in contrast to natural laws, there is no guarantee that the findings from such exercises will be relevant, or even true, when time passes by. My hope is that the findings produced in this thesis carry some knowledge of valuable character for the future.

The facts economists search for, are of the causal type. We seek to understand cause and effect from actions, laws and unexpected events. This is the closest thing to a solid ground one can achieve in social science. Therefore, it is not wonder that we search for such evidence to use as a base for decisions that affect welfare in our society. In the search for causality, economists use a range of methods, often with great resemblance to those of natural sciences. It is mostly the extensive use of mathematical and statistical methods one first think of, but also the aim to discover cause and effect through the use of experiments, occurring in the real world, or arranged in a lab. Therefore, many economists claim that it is the most scientific among social sciences (Colander, 2005). The thesis follows this tradition, analysing large amounts of data, documenting economic mechanisms through the use of traditional methods and quasi-experiments.

With this thesis, I aim to bring forward empirical evidence on policy relevant questions. Referring to Keynes description of an economist above, I agree that this requires a (surprisingly) large set of diverse skills. For example, in the two final chapters, I had to use a bit of my inner historian, statesman and philosopher, to find ways of identifying interesting relationships in the data. Building the data set used in the final chapter took some creativity and a bit of a computer scientist. While all three papers are policy relevant, the two first chapters are clearly aimed at the macroeconomic literature studying fluctuations, increasing our understanding of how households make their consumption and saving decisions when they face risk and uncertainty. The third chapter is concerned with alcohol policy and sick-leave, a heavy debated topic in Norway, with potentially high costs for society.

My focus has primarily been to provide insights relevant for macroeconomists. Hoover (2006) explains how the economics profession traditionally, and pedagogically, has separated between microeconomics and macroeconomics as two distinct and more or less independent parts of economics. Although some wanted to eliminate such a distinction (Lucas, 1987), it is still common to view them as different objects. One of the arguments for keeping a distinction, is that taking a world view, as is done in macroeconomics, do not require carrying all microeconomic details to achieve the same mechanisms and insights. Furthermore, as more details are added, one faces the trade off between complexity and tractability. Thus, if understanding the economy as a whole is the goal, one must make some sacrifices to fit everything into one picture.

From Keynes *General Theory* (1936), via Meade (1937) nine-equation system, to the IS-LM model by Hicks (1937), macroeconomic modelling experienced much development after the Great Depression in the 1930s. In the 1970s, the divide between salt-water and freshwater macroeconomics led to two distinct developments in the way academic macroeconomics is done: it came to be required that all theoretical models be based on an optimizing framework with model-consistent expectations (Wren-Lewis, 2018). In their historical account of macroeconomic modelling, Vines and Wills (2018) explain how incorporating such micro-foundations led to "the New Keynesian benchmark DSGE[1] model", represented by the models in Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007). Although these models could be regarded great achieve-

[1]Dynamic Stochastic General Equilibrium

ments and progress, the financial crisis hit in 2008, and macroeconomic modelling found itself under severe critique. These models could not explain what happened, and did not have clear advice on how to fix it. This spurred interest in the role of heterogeneity and distributional features for macroeconomic fluctuations; The need to better understand sources of risk and vulnerabilities at the micro level, and take into account aggregate implications from the fact that people are affected and behave in different ways.

While a subset of the macroeconomic literature have focused on heterogeneous agents and incomplete markets for nearly three decades[2], the aim has mainly been to investigate consumption and saving behavior, inequality, redistributive policies, economic mobility and other cross-sectional phenomena (Kaplan and Violante, 2018). Krusell and Smith (1998) solved an incomplete market model of heterogeneous agents with aggregate shocks to the economy. One of their findings was that, apart from computational costs, adding this type of heterogeneity did not matter much for outcomes compared to having one representative agent. This great piece of academic progress may have hindered further development on the course of heterogeneity, as the agenda seemed less worth to pursue when there where no apparent gains to be made. Fast forward twenty years, and modern macroeconomists talk about Heterogeneous Agent New Keynesian (HANK) models. These models allow for a more realistic set up of the household decision environment, with the potential for better dynamics and understanding of fluctuations and distributional features of the economy. In building such models, it is important to have a deep understanding of household behavior. This can be achieved through studies using micro data to document household behavior empirically, as is done in the chapters of this thesis.

Private consumption accounts for 50 percent of the Norwegian mainland gross domestic product.[3] Thus, how households contribute to and tackle fluctuations, is of the essence in the new macroeconomic research agenda. The academic literature on household consumption and saving is far from new, being based around the Permanent Income Hypothesis of Friedman (1957). The intuition of this model is that households make their current consumption decision based on their lifetime income. Having an idea of how much one can spend during the lifetime, but experiencing

[2]See Kaplan and Violante (2018) for a rich review of the heterogeneous agent incomplete-market models literature. Bewley (1983), İmrohoroğlu (1998), Huggett (1993) and Aiyagari (1994) are considered seminal contributions.

[3]Based on the Norwegian National Accounts (2017). Accessible from [https://www.ssb.no/en/nasjonalregnskap-og-konjunkturer/statistikker/knr]

that the income is typically low in the beginning of life, and growing before retirement, households smooth consumption by the use of debt and saving. Theoretical contributions to our understanding of risk and uncertainty in this context, goes all the way back to Sandmo (1969) and Dreze and Modigliani (1972). Sandmo (1969) showed how to solve the household problem when returns to savings are risky, and the household is risk averse. Dreze and Modigliani (1972) discussed optimal consumption behavior when there is uncertainty related to future resources. The literature has since explored and tested a range of different elements of the consumption problem that is internal or external to the consumer. For example, some studies look at the way consumer's preferences are defined and its implication for their behavior. Examples include that of temperance and prudence, habit formation and network effects (see Kimball (1990), Constantinides (1990) and Campbell and Cochrane (1999)). Others have worked on market imperfections that characterize the decision environment of the consumer, such as restricted credit access (see for example Zeldes (1989) and Deaton (1991) on liquidity constraints).

The empirical literature on the topic is extensive. Jappelli and Pistaferri (2010) give a review of the literature studying consumption responses to income changes. They highlight the important difference between expected and unexpected changes to income, which is a difficult thing to separate empirically. Furthermore, the literature is concerned with differences in the persistence of income shocks, since the Permanent Income Hypothesis predicts that only permanent shocks should affect consumption. A considerable amount of papers have studied the income process of households, and its relation to consumption and inequality (see for example Blundell, Pistaferri, and Preston (2008), Guvenen (2007), Guvenen and Smith (2014) and Guvenen, Karahan, Ozkan, and Song (2015)).

Attanasio and Weber (2010) surveys the literature on life cycle models and their implications for public policy. In their thorough account, they present some promising avenues for future research, highlighting the role of habits, temptations, beliefs and expectations in consumption and saving behavior. De Nardi and Fella (2017) give an overview of saving motives and conclude that more research on the relative importance of different motives, such as preference heterogeneity, rate of return heterogeneity, bequests and human capital, is still needed. Thus, the evidence presented in the chapters of this thesis seems welcome.

Recently, high frequency consumption data have allowed researchers to gain further knowledge on consumer behavior at an even more detailed level. For example, Baker, Johnson, and Kueng (2017) use high frequency data to look at shopping behavior and find evidence pointing towards shopping-trip fixed costs. Hinnosaar (2016) uses alcohol sales data and find that consumers act in accordance with time-inconsistent preferences, meaning people do not act according to plan when they are subject to temptations. Such insights can be important when considering consumption and saving behavior also on the aggregate level. Additionally, other literatures might benefit: The final chapter of this thesis uses knowledge about consumer behavior, such as the presence of travel costs and time-inconsistent preferences, to solve an identification problem in the literature on sick leave and alcohol consumption.

In the following, I give a non-technical summary of the chapters, their theme and results.

Chapter 1: Saving in good and bad times

This chapter is written with Elin Halvorsen, a researcher at Statistics Norway. We explore one of the proposed channels for reduced aggregate demand during bad times, namely that increased uncertainty and risk can make households reduce their consumption and save more. The problem encountered investigating this empirically, is that of quantifying the risk households perceive and allegedly react to. Since involuntary unemployment is proven to be the most important risk for household's income, we first exploit rich data on individuals' characteristics to estimate a model for unemployment risk. While this method has been used in the literature before by Carroll, Dynan, and Krane (2003), our data allow us to use the model to predict the probability of becoming unemployed in a given year, and see how this varies over time for different sub groups.

We then test whether the saving behavior of households with similar characteristics is different in times when they experience increases in risk, as opposed to times when they experience decreases in risk. Since we do not know how much assets households will hold for insurance purposes, pension and bequests, we assume that households with similar observable characteristics should have the same saving behavior, and test whether they save relatively more (or less) in periods where the job loss risk have increased (decreased). We find that decreases in risk has a negative effect on saving. Households are more willing, or able, to obtain debt when the job loss risk decreases. We test whether there are differences between young and old households, and find that the effect is primarily driven by the young. Theoretically, there are several reasons why this should be the case. For example, the consequence for life time earnings from a job loss incident is assumed to be higher for young people. The presence of liquidity constraints can magnify this. Finally, we show that the effect i asymmetric. This means that households seem to react differently to increases and decreases in the predicted job loss risk in Norway. One explanation for this result is that the generous welfare system in Norway reduces the impact an event of job loss has for future income and consumption, compared to other countries. If households are not worried about insuring against job loss, but have an easier time obtaining debt during good times, one could observe such asymmetry. Whether this is driven by preferences (willingness by consumer) or external factors (willingness by creditors), remains an open question.

Chapter 2: Gambling with the family silver

In this single authored chapter, I study household financial choices when public spending is subject to change and uncertainty. In October 2007, news emerged that eight Norwegian energy producing municipalities had sold up to ten years of future earnings from their hydro-electric power plants and invested it in high-risk financial products. Some of these municipalities lost more than 80 percent of the invested amount. This unexpected event is first and foremost a sad chapter for the involved parties, however, it also represents a unique opportunity to increase our knowledge on household behavior. The inhabitants are randomly selected to be participants in an experiment. Therefore, it is likely that the changes in behavior we observe is caused by the event. Furthermore, to the extent the affected municipalities are not special, the results are generalizable.

My main finding is that private consumption is sensitive to the economic condition of the local government. I claim that uncertainty related to future fiscal policy induce households to delay consumption.

The way I estimate this change in behavior, is by comparing the affected households to a group of households that is similar, but not affected. These are not treated by the experiment, and represent the control group. The crucial assumption is that the affected would have acted similarly as their control group, if it were not for the event. Estimating the difference in difference, I show that private consumption went down by 1.8 percent in the five worst affected municipalities the year after. While public spending decreased permanently in the years following the event, the consumption effect was temporary. I therefore argue that the behavior is not driven by the changes in public spending themselves, but must be caused by the uncertainty experienced in the worst affected municipalities. Before they received extraordinary transfers in November 2008, these municipalities struggled to deliver the legally required service level.

To further investigate whether uncertainty is a plausible explanation, I show that households in the affected municipalities rebalance their portfolios to holding a lower share of risky assets. If households have preferences that make them sensitive to the overall amount of risk they are subject to, it is optimal to reduce the amount of risky assets if public services seem more at risk. An alternative explanation for this behavior is that households in the affected municipalities are reminded that stock markets are risky, and therefore want to reduce their exposure. While this could be the case, it is not entirely clear why inhabitants in the affected municipalities should be more responsive to such news than people living outside, given that the event was largely covered in national newspapers. The finding may have important implications for our understanding of risk-taking behavior, such as participation in stock markets. This is a promising way to go forward.

Chapter 3: Outlet proximity, alcohol consumption and sick-leave.

The final chapter is written with Timothy G.A. Wyndham, a fellow PhD-student at Norwegian School of Economics. This chapter uses the fact that travel costs are likely to be important for consumer behavior, to establish a causal effect of alcohol consumption on sick leave. Collecting spatial data on distances, population composition and unique store level revenue and volume data from the Norwegian monopolist for stronger alcohol (Vinmonopolet), we exploit a widespread roll out of new stores in the period 2000-2016. In 1997, the government decided to increase the accessibility of legal alcohol in Norway by letting Vinmonopolet expand the number of outlets. This led to a steady increase in the access of alcohol within regions, through reduced travel distance to the nearest outlet.

We first show that there is a positive effect on regional alcohol sales from Vinmonopolet stores, when there is a new opening. Controlling for a set of other factors, like income and age composition, we show that this increase is driven by the fact that a range of people now have reduced travel distance to an outlet, and therefore buy more. There are two main channels that could explain such a result. First, if people are restricted by an actual cost of travelling to a store, this has in fact been reduced, making it is relatively cheaper to buy alcohol. A second explanation relates to time-inconsistent preferences. Consumers might think that they should not drink too much, however, when they are nearby, they get tempted and decide to purchase more than initially planned. With higher proximity, this becomes more prevalent.

Having established that sales go up, we aim to solve one of the difficult identification issues in the literature on alcohol consumption and its negative consequences. When researching effects of alcohol consumption, one faces the challenges of selection and reverse causality. For example, if people who are more prone to taking sick leave also drink more, there will be a positive association between the two, but it is not necessarily caused by alcohol consumption. Second, if being on sick leave leads to more drinking, the same problem arises. Since we have shown that increased proximity (which is wide-spread across the country and not targeted at a specific population) leads to higher sales, we check if sick leave in these regions is affected as well. In line with previous findings in the literature, we find that sick leave increases in the regions where alcohol sales increases. An increase in alcohol consumption by 1 percent per capita causes the number of people on sick leave to go up by 3 per hundred thousand. Although we cannot point to a mechanism at the individual level, by connecting individual consumption of alcohol to individual instances of sick absence, previous studies have shown that these associations exists. Increases in proximity leads to more heavy drinking (Halonen et al, 2013a) and heavy drinking is likely to affect sick leave (Halonen et al, 2013b).

References

Aiyagari, S. R. (1994): "Uninsured idiosyncratic risk and aggregate saving," *The Quarterly Journal of Economics*, 109(3), 659–684.

Attanasio, O. P., and G. Weber (2010): "Consumption and saving: models of intertemporal allocation and their implications for public policy," *Journal of Economic literature*, 48(3), 693–751.

Baker, S. R., S. Johnson, and L. Kueng (2017): "Shopping for lower sales tax rates," Discussion paper, National Bureau of Economic Research.

Bewley, T. (1983): "A difficulty with the optimum quantity of money," *Econometrica: Journal of the Econometric Society*, pp. 1485–1504.

Blundell, R., L. Pistaferri, and I. Preston (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98(5), 1887–1921.

Campbell, J. Y., and J. H. Cochrane (1999): "By force of habit: A consumption-based explanation of aggregate stock market behavior," *Journal of political Economy*, 107(2), 205–251.

Carroll, C. D., K. E. Dynan, and S. D. Krane (2003): "Unemployment risk and precautionary wealth: Evidence from households' balance sheets," *Review of Economics and Statistics*, 85(3), 586–604.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005): "Nominal rigidities and the dynamic effects of a shock to monetary policy," *Journal of Political Economy*, 113(1), 1–45.

Colander, D. (2005): "The making of an economist redux," *Journal of Economic Perspectives*, 19(1), 175–198.

Constantinides, G. M. (1990): "Habit formation: A resolution of the equity premium puzzle," *Journal of political Economy*, 98(3), 519–543.

De Nardi, M., and G. Fella (2017): "Saving and wealth inequality," *Review of Economic Dynamics*, 26, 280–300.

Deaton, A. (1991): "Saving and Liquidity Constraints," *Econometrica*, 59(5), 1221–1248.

Dreze, J. H., and F. Modigliani (1972): "Consumption decisions under uncertainty," *Journal of Economic Theory*, 5(3), 308–335.

Friedman, M. (1957): "A Theory of the Consumption Function," *NBER Books*.

Guvenen, F. (2007): "Learning your earning: Are labor income shocks really very persistent?," *American Economic Review*, 97(3), 687–712.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2015): "What do data on millions of US workers reveal about life-cycle earnings risk?," Discussion paper, National Bureau of Economic Research.

Guvenen, F., and A. A. Smith (2014): "Inferring labor income risk and partial insurance from economic choices," *Econometrica*, 82(6), 2085–2129.

Halonen, J. I., M. Kivimäki, M. Virtanen, J. Pentti, S. Subramanian, I. Kawachi, and J. Vahtera (2013a): "Living in proximity of a bar and risky alcohol behaviours: a longitudinal study," *Addiction*, 108(2), 320–328.

——— (2013b): "Proximity of off-premise alcohol outlets and heavy alcohol consumption: A cohort study," *Drug and alcohol dependence*, 132(1-2), 295–300.

Hicks, J. R. (1937): "Mr. Keynes and the" classics"; a suggested interpretation," *Econometrica: Journal of the Econometric Society*, pp. 147–159.

Hinnosaar, M. (2016): "Time inconsistency and alcohol sales restrictions," *European Economic Review*, 87, 108–131.

Hoover, K. D. (2006): "Microfoundations and the Ontology of Macroeconomics," in *The Oxford Handbook of Philosophy of Economics*.

Huggett, M. (1993): "The risk-free rate in heterogeneous-agent incomplete-insurance economies," *Journal of economic Dynamics and Control*, 17(5-6), 953–969.

İmrohoroğlu, S. (1998): "A quantitative analysis of capital income taxation," *International Economic Review*, pp. 307–328.

Jappelli, T., and L. Pistaferri (2010): "The consumption response to income changes," *Annual Revue of Economics*, 2(1), 479–506.

Kaplan, G., and G. L. Violante (2018): "Microeconomic Heterogeneity and Macroeconomic Shocks," *Journal of Economic Perspectives, Special Issue on The State of Macroeconomics a Decade After The Crisis*.

Keynes, J. M. (1924): "Alfred Marshall, 1842-1924," *The Economic Journal*, 34(135), 311–372.

——— (1936): *General theory of employment, interest and money*. Atlantic Publishers & Dist.

Kimball, M. S. (1990): "Precautionary Saving in the Small and in the Large," *Econometrica: Journal of the Econometric Society*, pp. 53–73.

Krusell, P., and A. A. Smith, Jr (1998): "Income and wealth heterogeneity in the macroeconomy," *Journal of Political Economy*, 106(5), 867–896.

Lucas, R. E. (1987): *Models of business cycles*, vol. 26. Basil Blackwell Oxford.

Meade, J. E. (1937): "A simplified model of Mr. Keynes' system," *The Review of Economic Studies*, 4(2), 98–107.

Sandmo, A. (1969): "Capital risk, consumption, and portfolio choice," *Econometrica: Journal of the Econometric Society*, pp. 586–599.

Smets, F., and R. Wouters (2007): "Shocks and frictions in US business cycles: A Bayesian DSGE approach," *American Economic Review*, 97(3), 586–606.

Vines, D., and S. Wills (2018): "The rebuilding macroeconomic theory project: an analytical assessment," *Oxford Review of Economic Policy*, 34(1-2), 1–42.

Wren-Lewis, S. (2018): "Ending the microfoundations hegemony," *Oxford Review of Economic Policy*, 34(1-2), 55–69.

Zeldes, S. P. (1989): "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, 97(2), pp. 305–346.

CHAPTER I

Saving in good and bad times:

Time varying unemployment risk and saving dynamics in Norway

Saving in good and bad times:

Time varying unemployment risk and saving dynamics in Norway*

Oddmund Berg† and Elin Halvorsen‡

Abstract

Higher job loss risk is associated with higher income uncertainty and lower future earnings. Increased job loss risk could therefore make households reduce consumption and increase saving. Exploiting a long panel with population wide Norwegian register data on income, wealth and unemployment, we construct individual-level time variation in job loss risk. We then estimate the effect of changes in job loss risk on active saving decisions of households. We find that a one percentage point increase in the job loss probability leads to an increase in saving of 9%. Furthermore, we find that the result is driven by young households increasing their debt when risk decreases. This asymmetry suggests that job loss risk affects households' willingness, or ability, to obtain debt for financing consumption in good times, rather than causing them to save during bad times.

Keywords: unemployment risk; consumption and saving; panel data; heterogeneity

JEL codes: E21; D15; E24

*We gratefully acknowledge comments and suggestions from Janine Aron, Gernot Doppelhofer, John Muellbauer, Arvid Raknerud and Kjetil Storesletten. We also thank seminar participants at Statistics Norway, Norwegian School of Economics, and The 39th Annual Meeting of the Norwegian Association of Economists. We would also like to thank Espen Bratberg and Thor Andreas Aursland for comments on an early version of the paper at the 2015 NHH-UoB Phd-workshop.

†Norwegian School of Economics, Oddmund.Berg@nhh.no

‡Statistics Norway, elin.halvorsen@ssb.no

1 Introduction

Recent studies of income risk highlight unemployment as the main source of income uncertainty.[1] Despite being a transitory shock, a large literature find that unemployment may have severe and long lasting effects on future income.[2] Incorporating the risk of unemployment into models with incomplete markets, several recent studies[3] show that time variation in households' precautionary saving is a channel that could contribute to fluctuations in aggregate demand, and even make recessions deeper. In this paper we focus on this psychological mechanism, namely that still employed households reduce their consumption and save more when job loss risk increases. Using a long panel of Norwegian households, we study the relationship between changes in job uncertainty and household saving in a new framework.

Previously, the literature has focused on either cross-sectional individual measures of job loss risk (e.g. Carroll, Dynan, and Krane, 2003, and Benito, 2006), or aggregated measures of income uncertainty/job loss risk in a dynamic setting (e.g. Carroll, Slacalek, and Sommer, 2012, and Hahm and Steigerwald, 1999).[4] Our first contribution is to specify an empirical strategy allowing us to identify the relationship between individual job loss risk and saving in a dynamic setting, thus bringing these two strands of literature together. Starting with a setup similar to Carroll et al. (2003) and Benito (2006), we estimate a model of job loss risk based on observable characteristics. Thus, our measure of job loss risk has the advantage of being objectively measured and possible to link directly to observed saving behavior on a disaggregated level. Acknowledging that the level of predicted job loss risk from this model is likely to be endogenous, in the sense that the choice of a certain job might reflect preferences that also affect consumption-saving decisions, we exploit the fact that over time, variation in job loss risk conditioning on a set of observable characteristics, is likely to be exogenous. In other words, using a long panel of Norwegian households with information on unemployment status, income, financial assets and a rich set

[1] Low, Meghir, and Pistaferri (2010), Guvenen, Ozkan, and Song (2014), and McKay and Papp (2011).

[2] Arulampalam, Gregg, and Gregory (2001), Gregory and Jukes (2001), for results from Norway, see Nilsen and Reiso (2014)

[3] For example, Challe, Matheron, Ragot, and Rubio-Ramirez (2017), Challe and Ragot (2016), McKay (2017), Ravn and Sterk (2014), and Krueger, Mitman, and Perri (2015).

[4] Having disaggregated data, but using a self-reported measure of future job loss expectations, Stephens Jr (2004) find that households with high expected probability of job loss, that eventually kept their job, increase their food consumption afterwards. Benito (2006) has shown that self-reported measures of job loss risk may not contain enough information to distinguish effects on consumption and saving, and that a better measure of job loss risk is a model-based predicted job loss likelihood.

of demographic variables, we estimate time and individual-specific probabilities of becoming unemployed and exploit the panel dimension of the data to construct exogenous individual-level fluctuations in unemployment risk. Our results show that a one percentage point increase in the job loss probability leads to an increase in saving of 9%. Alternatively, to get a sense of the magnitude, a one standard deviation (0.015) increase in the job loss risk would lead to 13.5% increased saving. Back of the envelope calculation suggests that the implication for consumption is a reduction of about 0.7% for the median household.[5]

Our second contribution is to exploit the richness of our data to decompose the effect and obtain a clearer view of the driving forces behind the overall relationship. We decompose effects along three dimensions; assets components, age, and whether there is an increase or a decrease in the predicted job loss risk. We find that debt is the asset that is most actively changed in the face of job loss risk. We test the theoretical implications for different age groups and show that the effect is driven by young households, consistent with theoretical predictions. We then proceed to show that in Norway, the mechanism is asymmetrical in risk. Our results are mainly driven by consumption increases in relation to risk reductions. Viewing our results in context of the Norwegian welfare system, which provide good insurance in the case of job loss, this finding may suggest that job loss risk affect households' willingness or ability to obtain debt for financing consumption during good times, more than causing them to save during bad times.

The outline of the paper is as follows: In the next section we show how the consumption-saving decision of the household relates to unemployment risk. Section 3 describes the empirical approach and defines the identification framework. Section 4 describes the data, while in Section 5 we present our results. In Section 6 we discuss our findings, and Section 7 concludes.

2 Theoretical background

After the Great Recession there has been a renewed focus on the relationship between income risk and consumption-saving decisions, but the literature on the subject is far from new. The Buffer-stock model with income uncertainty was presented by Carroll (1992) (see also Carroll and

[5] Although not directly comparable, our main results are very much in line with the results obtained by Campos and Reggio (2015). Analyzing the relationship between consumption growth and the change in unemployment at group levels, they find a decrease in consumption of around 0.7% in relation to a one percentage point increase in the unemployment rate.

Kimball, 1996, Kimball, 1990, Carroll and Kimball, 2005). While Carroll (1992) originally focused on the dynamic relationship between income uncertainty and consumption, the empirical literature has investigated the relationship using a wide specter of proxies for income uncertainty, studying the implications for a range of different outcomes such as optimal unemployment insurance, differences in wealth accumulation, inequality and life cycle behavior (for the latter, see Gourinchas and Parker, 2002). The theoretical foundation generating the channel from job loss risk to aggregate demand is well understood and derived directly from impatient and prudent households' intertemporal allocation under uncertainty. More recent, the tractable Buffer-stock model of Carroll and Toche (2009) highlights the mechanism in relation to unemployment risk particularly well, providing closed form solutions for optimal consumption.

To illustrate the effect changes in job loss risk have on optimal consumption-saving decisions, we consider the model in Carroll and Toche (2009). Here, unemployment risk is the probability of a multiplicative shock to the income process of the household, being an absorbing state. While this is an implausible assumption, the intuition and qualitative results carries over to more realistic models, both partial and general equilibrium,[6] that can only be solved numerically.

The consumer maximizes the discounted sum of utility from consumption c_{it}. Utility comes from an inter-temporally separable CRRA utility function $u(c_{it}) = \frac{c_{it}^{1-\rho_i}}{1-\rho_i}$, where ρ_i is the coefficient of risk aversion. The maximization is subject to the dynamic budget constraint:

$$a_{it+1} = (1 + r)(a_{it} - c_{it}) + y_{it}\epsilon_{it} \tag{1}$$

where next period's wealth a_{it+1} is the previous period's wealth net of consumption c_{it}, plus labor income y_{it}. ϵ_{it} is an indicator on the employment status. Define the incident of becoming unemployed $u_{it} = 1 - \epsilon_{it}$ which can happen with a probability of $Pr(u_{it}) = p_{it}$. To derive analytical results, Carroll and Toche (2009) assume that the state of unemployed is irreversible. The solution shows that there exist a steady state target wealth for the consumers, depending on the unemployment

[6] For a general equilibrium approach, see Challe et al. (2017). According to their model, the aggregate demand effect dominates the supply effect from increased saving. In the current paper, we abstract from such equilibrium and transmission effects, studying the reduced partial effect on saving from job loss risk and its heterogeneity.

risk, interest rate, expected income growth and preferences:

$$a^*_{it} = f(Pr(u_{it}), r, \Delta y_{it}, \beta_i, \rho_i) \tag{2}$$

where β_i is the discount rate and ρ_i is the coefficient of risk aversion. Target wealth is the result of insuring against income shocks, while taking lifetime income and inter-temporal considerations into account. Target wealth depends positively on unemployment risk, the interest rate, impatience and risk aversion, and is decreasing in the growth rate of wages.

In a static setting, it is difficult to identify saving mechanisms from this framework, as two otherwise equal consumers might hold different amounts of wealth, a^*_{it}, due to unobserved preferences (β_i, ρ_i). In a dynamic setting, such as the context of our yearly data, such unobserved time fixed effects can be controlled for. The empirical implications from the model is derived for the dynamic setting by assuming that households reach their target wealth by the end of each year. Thus, if unemployment risk increases from one year to another, target wealth by the end of the second period is higher than before. All else equal, any risk averse consumer will save a larger share of income than before in order to reach the higher target. The same mechanism applies for a reduction in unemployment risk. A decrease will reduce the target wealth, allowing the consumer to consume more in that period. Controlling for other factors that affect wealth accumulation (such as private pension saving, bequests, etc), the theory predicts a positive relationship between saving and changes in job loss risk.[7]

While the closed form solution is easy to interpret, the simple model does not give much insights about sources of heterogeneity. However, Carroll et al. (2003) solves a version of the model where unemployment is not an absorbing state, and where the consumer has finite lifetime. While this model is less tractable, in the sense that the solution is not analytical, it highlights two important features that the simple model do not address. First, the strength of the channel from changes in job loss probability to saving depends on the level of unemployment benefits provided by the public.[8] In Norway, the replacement rate is 0.64, reducing the expected effect of job loss risk on consumption compared to countries with lower replacement rates.[9] Having

[7]Conditional on other factors, saving is the change in target wealth by definition $\Delta a^*_{it} = s_{it}$.

[8]For example, if the replacement rate is equal to 1, the household receives the same income independent of job loss or not, removing any effect on consumption and saving.

[9]For a detailed overview of the Norwegian unemployment benefit system, see Appendix A.

established that the income loss associated with a job loss matters, one could imagine that there is heterogeneity across occupations. Even if the replacement rate is the same, the expected length of an unemployment spell may vary across occupations. Finally, if households have finite lives, the effect vary with age. Young households have more working years to be affected by a higher job loss probability, hence, the increase in risk affects a larger share of their lifetime income. This is also an empirical finding of Gourinchas and Parker (2002), who find that because of other wealth accumulation, households are only responsive to income risk until their mid-forties.

An important difference between the modeling framework and the empirical setting should be noted, namely the relationship between income and probability of job loss. Unlike income risk represented as a mean preserving spread, job loss risk carries two effects. First, increases in the job loss risk increases the variance of household's income. Thus, it increases uncertainty. Second, increases in job loss risk decreases expected life time income. Therefore, when observing an increase in the unemployment risk, there are two effects on the saving decision. To separate out a pure uncertainty effect (i.e. the precautionary effect) on target wealth from a change in the job loss risk, Carroll and Toche (2009) assume that income is subject to a growth factor that increases with the job loss probability. Having a compensating factor in income growth is necessary, since increasing the job loss risk also have a negative effect on the expected life time income. By compensating that reduction through the income growth factor, changes in job loss risk turns into a mean preserving spread of income, only affecting uncertainty. Empirically, such a compensating factor is implausible. Evidence suggests that unemployment risk and income growth is negatively related through scarring effects on wages (Arulampalam et al., 2001 and Nilsen and Reiso, 2014). Since a drop in income growth and increased uncertainty both lead to increased saving, one would have to make explicit assumptions on future income growth, to disentangle the two. In this paper, we aim to understand the overall effect of job loss risk on aggregate demand through consumption-saving decisions of households. We therefore abstract from this issue.

3 Empirical approach

The saving equation we estimate follows directly from the theoretical relationship between job loss risk and target wealth. Since our data is on a yearly frequency, we must assume that during the course of a year, households experiencing a change in their job loss risk adjust to their new target wealth by the end of that period by saving or dissaving:

$$s_{it} = \Delta a^*_{it}$$

where saving, s, is defined at the change in wealth adjusted for capital gains, i.e. active saving decisions. Studying changes in wealth, as opposed to accumulated assets, a_{it}, has several advantages empirically.[10] First of all, since households hold wealth for a number of different reasons in addition to risk, the empirical measure of a_{it} is likely to reflect long term economic considerations as well as insurance. In addition, it may be difficult to distinguish between wealth accumulated as part of long term life-cycle planning and wealth intended as a buffer against income shocks or unexpected expenditures. This complementarity of wealth is discussed in Blundell, Etheridge, and Stoker (2014), and underline that the relationship between wealth accumulation and risk is difficult to identify. Active saving is a natural choice of dependent variable for our purposes, reflecting current consumption-saving decisions that ultimately affect aggregate demand, while at the same time representing changes to target wealth in line with theory.[11]

The effect of time-varying job loss risk on active saving, s, is given by the parameter δ in the following equation:

$$s_{it} = \delta \Delta Pr(u_{it}) + \theta V_{it} + \mu_i + \alpha_t + \eta_{it} \tag{3}$$

where $\Delta Pr(u_{it})$ is a one-year change in job loss risk, V_{it} is a set of control variables capturing other saving motives, μ_i is an individual fixed effect, α_t is a time-fixed effect and η_{it} is the error term, assumed to be *iid*. Note that the relationship between observable characteristics and saving is assumed to be constant over time. This is key for identification, as beyond year fixed effects, it is assumed that households with similar characteristics act similarly, where it not for changes in job loss risk.

The first step is to quantify job loss risk. We do this by estimating the relationship between observed job loss and personal characteristics, following Carroll et al. (2003) and Benito (2006),

[10] In Section 4.2, our definition of active saving is presented in more detail.
[11] We also apply alternative definitions, such as saving to income rate, as robustness.

except that we use a linear probability model:[12]

$$Pr(u_{it}) = \varphi_t V_{it} + v_i + \varepsilon_{it} \tag{4}$$

where u_{it} is an indicator equal to one if the individual becomes unemployed during the course of a year.[13] V_{it} is a set of predictors believed to influence job loss risk and saving, such as age, labor market region, country of origin, sex, education, and labor income level. V_{it} consists of indicators of membership in a specific group, so there are no functional form assumptions involved. v_i is an unobserved fixed effect affecting job loss risk, such as risk aversion or other personality traits that affects job stability. With one time period of data, the presence of v_i is problematic, and one need to find an exclusion restriction to achieve identification when inserting (4) into (3).[14]

We exploit our long panel and estimate the probability model yearly, letting all parameters vary over time. After obtaining the parameter vector φ_t for each time period t, we compute time-variation in unemployment risk by simply taking the first differences to the predicted probabilities, removing v_i:

$$\Delta Pr(\hat{u_{it}}) = \hat{\varphi}_t V_{it} - \hat{\varphi_{t-1}} V_{it-1} + \varepsilon_{it} - \varepsilon_{it-1} \tag{5}$$

While the level of $Pr(\hat{u_{it}})$ might be affected by unobserved characteristics v_i, time differentiation removes the unobserved effect, leaving us with time-variation in job loss risk generated by changes in the fraction of workers with similar characteristics losing their job. Including year-fixed effects, movement in the unemployment rate, or the interest rate, and their effect on aggregate saving behavior, do not influence δ. V_{it} controls for life-cycle behavior, labor income differences, educational differences and family situation, which are all assumed to be constant for the time period at hand. If δ systematically explains variation in saving, it is because deviations from expected saving, given the current set of characteristics, co-vary with the individual time-variation in job

[12] At the outset we estimated all results using both a logistic and a linear probability model. Our results are not sensitive to the probability model chosen. We therefore present the linear model here, where identification is straight forward.

[13] Note the distinction between becoming and being unemployed. We consider only workers that have become unemployed within the period to avoid any concerns related to long term unemployment.

[14] Carroll et al. (2003) obtain their main results excluding regional dummies, while Benito (2006) favors lagged unemployment status, size of household and the household head's employer- and union status.

loss risk.[15] In the case where elements of V_{it} changes, one must assume that the previous set of characteristics V_{it-1} affect present saving only through its impact on the change in unemployment risk. We test this assumption in Section 5.2 by showing that our results are robust to the inclusion of lagged saving.[16]

One may argue a logistic model would be a better way of modelling the job loss probability. A known disadvantage with linear probability models is that they can predict negative probabilities, although this is less of a problem in our case as the specification in 4 consists entirely of indicators and not continuous variables. A clear advantage of the linear model is that it completely removes any unobserved fixed effects that will be endogenous in the saving regression. In Section 5.2 we present evidence that our results are not sensitive to the choice of probability model.

4 Data

4.1 Data description

The data is derived from a combination of administrative registers covering the whole Norwegian population for 22 consecutive years; 1993-2014. Data are assembled on the basis of annual tax records as well as other registers, such as the one administered by the Norwegian Labor and Welfare Administration. These data are of high quality as most information is third-part reported to the tax authorities, and very little is self-reported. Employers, banks, brokers, insurance companies and any other financial intermediaries are obliged to send both to the individual and to the tax authority, information on payment of earnings, the value of the asset owned by the individual and administered by the employer or the intermediary, as well as information on the income earned on these assets. Because of the reporting for tax purposes, obtaining a precise measure is unproblematic for most of the financial portfolio. Since we are interested in time varying saving, our measure of wealth at the outset is the sum of deposits, money market funds, stock market funds, bonds, stocks, and other financial assets.[17] Real wealth is less precisely measured in the tax records, often represented by inadequate ”tax values”, therefore we limit our measure of saving

[15] Appendix B provides a further discussion of the identification strategy.
[16] The results are also robust to inclusion of future income growth.
[17] The tax valuation of stocks, bonds and mutual funds were subject to different rules over the period 1993-2007, varying between 30 and 85 percent of the market values. As a first step we adjust all financial asset categories so that they represent market value, and not their tax values.

to financial wealth changes. Tax records are annual and report the stock of wealth as measured by December 31st.

Furthermore, the data set contains information on household education (level and type) from the National Education registry. We use a detailed classification of 50 educations (combinations of length and fields) established by Kirkebøen (2010), see Appendix C. Likewise, we use a thorough reclassification of economic labor market regions by Bhuller (2009), also described in detail in Appendix C. Finally, we use information about whether the individual received unemployment insurance to identify occurrences of unemployment. The definition of "becoming unemployed" is derived from a person having received unemployment benefits in year t, but not in year t-1 (as opposed to being unemployed, which is unconditional of the status in t-1).

In Norway, income taxes are levied at the individual level, while wealth is taxed jointly by couples. The information from the tax returns is combined with family identifiers from the population register in order to be able to aggregate income and wealth information at the family level. Therefore, we use information about wealth and saving at the family level. For education, income and other characteristics we use individual information, controlling for family composition and spousal income.

4.2 Active saving

In our data, the change in nominal financial assets from one year to the next consists of two parts; changes in the stock of asset and changes in the valuation of the asset. We do not want unrealized changes in the asset's price, i.e. unrealized capital gains and losses, to be part of our saving measure as they do not reflect the household's active saving behavior. Thus what we call "active saving" is the change in financial assets minus capital gains and losses. For stocks we have used the Oslo Stock Exchange index (OSE) to calculate gains and losses, for mutual funds we have used a combination of the OSE and the MSCI World index and for bond we have used the Treasury bill rate. For more details on the calculations we refer to Fagereng and Halvorsen (2017).

Measures of saving as the first difference in wealth tend to show high variance and extreme outliers. Several strategies can be chosen to avoid problems stemming from highly influential extreme values, such as deleting or manipulating the observations identified as problematic, or, by transforming variables so that the distribution of all variables have a lesser spread than the untransformed. One such possible transformation we can use is the inverse hyperbolic sine transformation $s = \ln(S + \sqrt{S^2 + 1})$ that behaves as $\pm\log(|S|)$ everywhere with the exception of in the neighbourhood of zero (Burbidge, Magee, and Robb, 1988). We also present results using saving rates instead of the inverse hyperbolic sine transformation.

4.3 Sample selection

Our sample consists of individuals in working age, i.e. between 24 and 60 years old. We exclude families whose joint annual labor income is lower than the basic unit of the National Insurance Scheme.[18] Since we base our identification on time differencing, we only include households that have a minimum of four consecutive observations. Based on the dependent variable, we trim the data set yearly excluding high and low savers, keeping those between the 1 and 99th percentile. Since we are worried that households buying and selling houses have erroneously large changes in their assets, we exclude year-observations where the household is registered as having moved during the year.

5 Results

5.1 First stage: The probability of becoming unemployed

We estimate the probability of becoming unemployed using the following linear specification:

$$u_{it} = \gamma_t + \varphi_t V_{it} + \varepsilon_{it} \tag{6}$$

where the probability of becoming unemployed is assumed to be determined by individual characteristics. V_{it} consists of indicators for age, labor market region, education, labor income quintile, spouses' labor income quintile, country of origin, sex and whether the worker is self-employed or not. To avoid simultaneous changes in probability and saving, we use lagged observations of income. However, for other variables, we believe it is important to reflect the current set of characteristics. Note that V_{it} is a large set of dummy-variables and we have therefore not restricted

[18] The tax regulations contain a number of amounts and amount limits which are directly linked to the basic unit in the National Insurance scheme. In 2017 the level of one basic unit is about 93,500 NOK (approximately 11,500 USD).

any of the observable characteristics to be in a linear relationship.[19] On average, the probability model has an adjusted R-squared of 0.015. In all specifications, the coefficients are allowed to vary over time, as explained in detail in Section 3. Thus, we capture differences in job loss risk between subgroups each year, and each of these subgroups will experience variation from year to year.

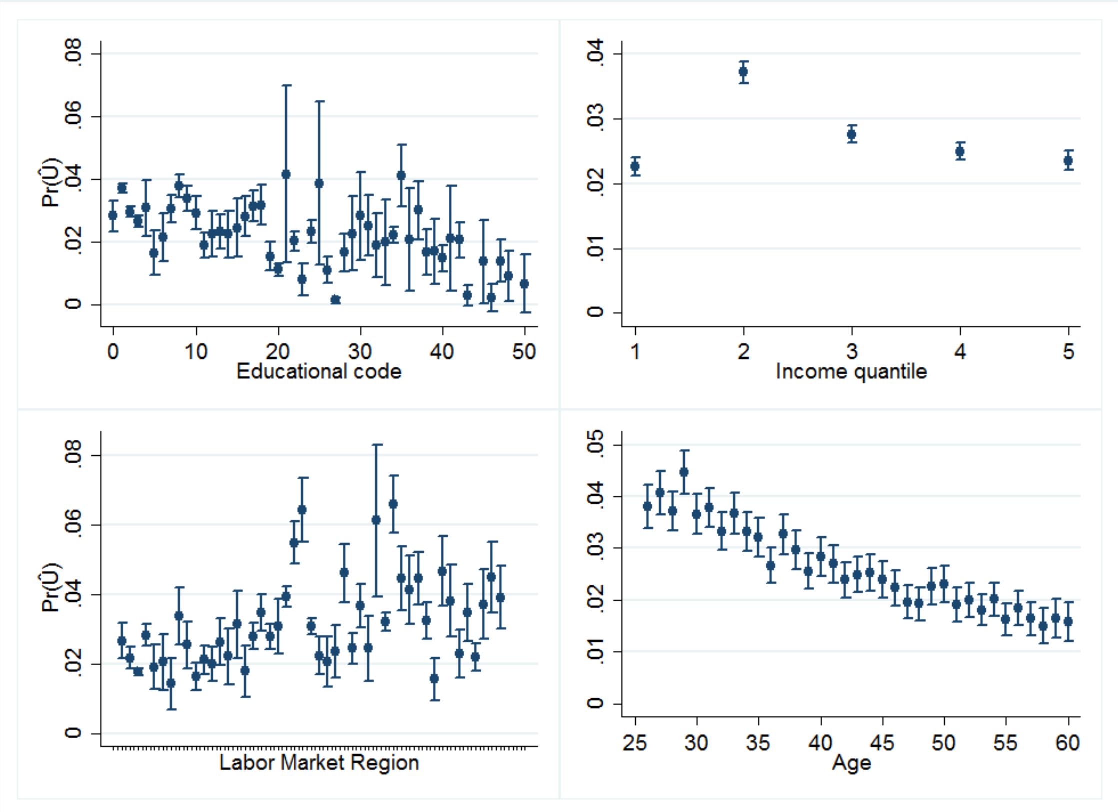

Figure 1: Margins plot of the probability of job loss, by local labor market region, age, education and labor income quintile for year 2000. Vertical lines represent confidence intervals at 95%.

We present some graphical evidence of the heterogeneity in job loss risk in Figure 1. The margins plots show a snapshot from the probability model in year 2000 for the coefficients of local labor market region, age, education group and labor income quintile. Overall there is considerable variation in job loss risk based on characteristics, although some groups are small, leading to large standard errors.

The educational categories are not a continuous scale, as the different groups are constructed on the basis of a combination of field of study and length. However, as we move to the right in the

[19]In the main specification, these amounts to 151 indicators per year.

figure, the educations are in general at a higher level. We note that these margins indicate that some groups with high education also experience quite high risk, such as those with a master level in humanities. The two groups with the largest standard errors are those with a degree in journalism, and those with a bachelor degree in maritime studies. This exemplifies that the differences seen between types of education, or professions, does not necessarily represent differences in risk between occupations, but may reflect other unobserved differences within educational groups. One could imagine that there is a greater variety of occupations within the group of bachelor level journalists, than for example the group of nurses. Some may even work in industries or occupations that is unrelated to their highest completed education, leading to more noise in the predicted values.

There is a great deal of variation between labor market regions. Some regions have a distinct lower probability, while the northernmost regions of Norway have a higher probability of becoming unemployed than southern regions. The probability of becoming unemployed is declining in age, from an average likelihood of about 4 percent at age 25 to around 1 percent at age 60. Last, we see that the probability falls in income quintile once the income is higher than the first quintile, confirming the notion that higher paid jobs are also less risky.[20]

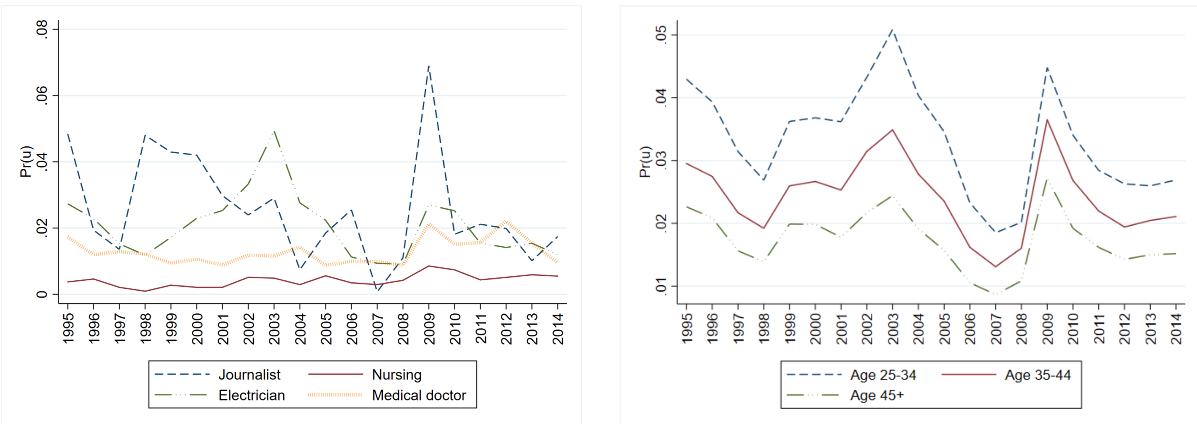

Figure 2: Margins plot of the probability of job loss over time, four education groups and three age groups

Although the cross sectional differences between different groups are interesting in themselves,

[20] The lowest labor income quintile is a very heterogeneous group that may contain persons with a combination of incomes, such as people who are partly on disability benefits, persons who combine self-employment with other labor income, and the very rich, who due to income shifting occasionally report very low annual labor income.

the most important variation in the job loss risk for our identification is variation from year to year. For illustration, we plot time variation in job loss risk for four selected educational groups, journalist, nursing, electrician and medical doctor, and three age groups in Figure 2. The figure shows that the differences we obtain between educational categories are largely as expected; medical doctors and nurses have low risk, stable jobs. Electricians and journalists, on the other hand, seem to be more susceptible to changing economic conditions over time. In addition to that, we highlight that different occupations are at risk at different times, as more journalists lose their job in the beginning of the period, while electricians experience a peak in 2003 (the years 1995 and 2003 represents periods of economic downturn in Norway). The large peak for journalists in 2009 reflects the financial crisis.[21] By age, the development over time is more similar, albeit at different levels. However, we see that the business cycle variations are more pronounced for younger households.[22]

Aggregating these data and looking at the distribution of changes over time, Figure 3 displays a box plot of the first differenced predicted job loss probability from the model from 1995 to 2014. Each box in the figure displays the median, 25th and 75th percentile, in addition to the upper and lower adjacent value, of changes in job loss risk over time. We see that the median change in job loss risk is close to zero from year to year, however, the distribution differs over the years. Going back to the previous argumentation regarding heterogeneity in job loss risk, the figure substantiates the fact that not all households experience the same development of job loss uncertainty over time, but that the model captures aggregate events like the build up and aftermath of the financial crisis.

5.2 Second stage: Saving response

Assuming that households adjust their preferred buffer to the level of uncertainty they face each period, the effect of job loss risk on saving can be found by estimating equation (3). We restrict the second stage regression to households who remain employed as we do not want our estimates to be affected by the change of behavior associated with an unemployment spell.[23] All variables

[21] Again, note the distinction between educational background and occupation: The peak may be driven by the share of people with journalism background having jobs that do not reflect their education, rather than a mass lay-offs in the newspaper industry.

[22] For the probability of job loss over time by labor income quintile, see Appendix D.

[23] For comparison we have also estimated (3) including individuals who become unemployed in the current period, and an opposite version where we exclude individuals who experience unemployment in the future in addition to individuals who become unemployed in the current period. The results from these robustness checks do not

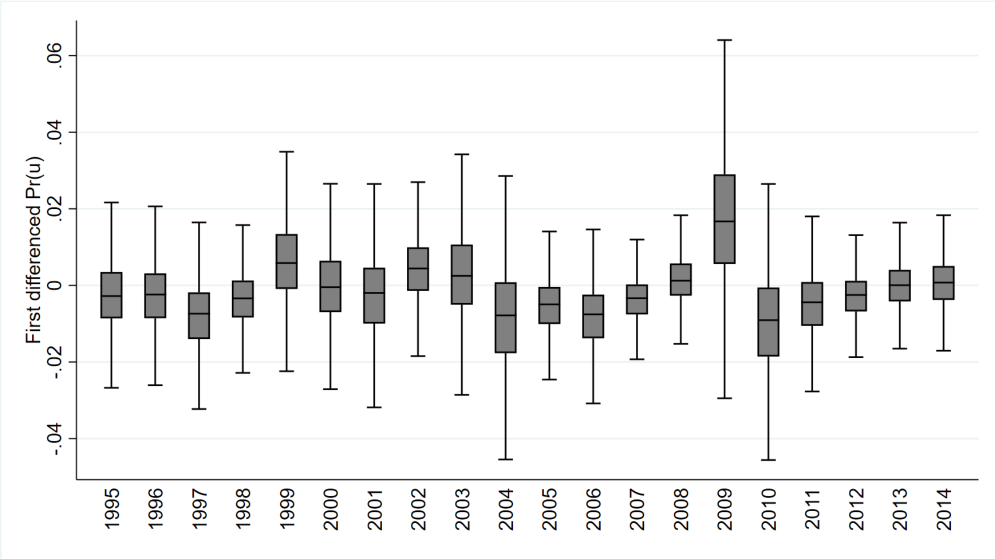

Figure 3: Box plot of first differenced predicted job loss probabilities, estimated in equation (6)
Each box displays the median, 25th and 75th percentile of the distribution, in addition to upper and lower adjacent values.

used to predict the job loss risk must be included as controls in the second stage, in addition to the predicted change in job loss risk. The coefficient of $\Delta Pr(\hat{u})$, δ, reflects changes in saving as a consequence of changes in the job loss rate generated by the fact that more or less households of a certain characteristic have lost their job in the period. Under the assumption that individual households cannot affect this development, and that there is a constant relationship between observed characteristics and saving, δ represents the effect on saving from an increase in the job loss risk.

Table 1 presents estimates of δ from a range of second stage saving regressions. The first row and first two columns present our baseline results, using the whole sample period (1995-2014) and hyperbolic sine transformed active saving as the dependent variable. We find a significant positive effect of an increase in job loss probability on saving. A one percentage point increase in the job loss probability leads to 9% increased saving. A one standard deviation increase in the job loss risk is 1.5% points, so an alternative way of viewing the results is that we find an increase in saving of 13.5% from a one standard deviation increase in the probability of job loss. Including fixed effects does not alter our results. Although not directly comparable, our main results are very much in line with the results obtained by Campos and Reggio (2015). They find

substantially change the findings presented here.

Table 1: Main results and alternative specifications - active saving 1995-2014

	Baseline		Incl. lagged dependent	
	OLS	FE	OLS	FE
(1) Benchmark:				
$\Delta Pr(\hat{u}_t)$	9.07**	8.67**	9.17**	9.71**
	(0.705)	(0.693)	(0.705)	(0.702)
(2) Logistic probability model:				
$\Delta Pr(\hat{u}_t)$	8.19**	7.77**	8.27**	8.60**
	(0.474)	(0.515)	(0.473)	(0.505)
(3) Timing, both periods:				
$\Delta Pr(\hat{u}_t)$	8.98**		9.10**	
	(0.695)		(0.694)	
$\Delta Pr(\hat{u_{t+1}})$	10.44 **		10.57**	
	(0.769)		(0.767)	
(4) Timing, next & previous period:				
$\Delta Pr(\hat{u}_t)$	8.96**		9.08**	
	(0.707)		(0.707)	
$\Delta Pr(\hat{u_{t+1}})$	10.41**		10.54**	
	(0.774)		(0.773)	
$\Delta Pr(\hat{u_{t-1}})$	0.17		0.18	
	(0.573)		(0.572)	
(5) Relative changes:				
$\Delta^R Pr(\hat{u}_t)$	0.0374**		0.0377**	
	(0.0054)		(0.0055)	
Alternative dependent variable:				
(6) Saving rate:				
$\Delta Pr(\hat{u}_t)$	0.88**		0.91**	
	(0.044)		(0.045)	

Note: The set of control variables include indicators for age, sex, marital status, children, education/type of work, local labor market region, birth country, self-employment, home-ownership, own labor income quintile, spouses' labor income quintile, and year.
Standard errors based on 100 two-stage estimations on 20 % random sample draws.
(These have approximately 4.5 million observations each.)
* $p < 0.05$ ** $p < 0.01$
Full results including all coefficients available upon request.

a clear decrease in consumption of around 0.7% in relation to a one percentage point increase in the unemployment rate. Remember that our dependent variable is household active saving, a variable that is a considerably smaller element of the household budget than consumption. To give a flavor of the real magnitude of our result, consider that the median saving in our estimation sample is roughly 21,000 NOK. Using the same imputation method as in Fagereng and Halvorsen (2017), median consumption in our sample is roughly 400,000 NOK. This implies a decrease in consumption of about 0.7% from a one standard deviation increase in the job loss risk. These numbers are of course not a precise measure of consumption, however, as a back of the envelope calculation, it suggests that our results are in the same ballpark as Campos and Reggio (2015). We also note that our estimates probably represent a lower bound, as measurement error in our predicted probabilities are likely to cause attenuation bias.

In the following we explore the relationship further by looking at alternative definitions of active saving and job loss probability, to enhance our understanding of the saving dynamics. These specifications also work as robustness checks, addressing some of the concerns one might have with the identification strategy.

Lagged saving

When we first difference the estimated probability, we essentially use V_{it-1} as an exclusion restriction, since these are predictors of job loss risk that is not included in the second stage. This is only a problem if the elements of $V_{it} \neq V_{it-1}$, since V_{it} is controlled for. We therefore test if our results are sensitive to the inclusion of lagged saving as an explanatory variable. The results are presented in the third and fourth columns of Table 1. The inclusion of lagged saving causes the coefficient on job loss risk to increase slightly, but there are no significant differences between the coefficients.

Logistic probability model

In the second row of Table 1, we present results using predicted changes in job loss risk generated by a logistic probability model in the first stage, rather than a linear model. As we argued above, there can be disadvantages using a linear probability model in terms of predictions of negative probabilities and a probability distribution that is perhaps too normally distributed. As seen throughout, there are no significant differences between the two approaches, but the logistic probability model have slightly lower point estimates.

Timing and dynamics of job loss risk and saving

Having a long panel with a yearly frequency, it is natural explore the timing and dynamics of job loss risk. Measuring the risk of job loss next year instead of the present can pick up valuable information, since households might increase their saving before job losses manifest themselves in the data, i.e. in anticipation of an increase in job losses. We build on the intuition that dismissing workers might be considered a last resort when firms already have experienced bad times for a while. Since it is during ”bad times” the presence of increased job uncertainty should affect the saving of the households, while job losses show up in the data later, there might be a lag between the observation of job loss risk and saving.

In order to study these effects, we estimate the first stage linear probability of job loss using next period's unemployment indicator as the dependent variable. We follow the same procedure as with our baseline results and predict the job loss risk based on future job losses, for the current period. We then first difference and estimate the second stage as before. The row labelled (3) in Table 1 shows the results of changing the timing of unemployment risk. From these results, we can conclude that changes in next year's job loss risk matter for current saving decisions.

In line with the results obtained by Basten, Fagereng, and Telle (2016), we find that households are forward looking and respond to future unemployment risk. Moreover, our results show that this response is present also among workers that do not become unemployed. The point estimate is of similar magnitude as the effect we find using current period's unemployment risk.

If there is persistence in the job loss risk, one might suspect that we are simply picking up the same variation through different variables. However, this is not the case. The correlation between the two risk measures is only 0.03, which is natural given that these measure separate job loss incidents. Unless the same group of people are fired every period, they should not be similar. Including both next period's job loss risk and previous period's job loss risk in the second stage (row (4) in Table 1) shows that current and future job loss risk matters for current saving, with an almost similar magnitude. Furthermore, previous period's risk changes does not affect current saving behavior, consistent with the theory. The third and fourth column of Table 1 shows that these results are robust to the inclusion of lagged dependent variable.

Relative changes in job loss risk

Next, we study an alternative definition of the job loss probability, namely relative changes. We define a relative change in the probability of job loss as

$$\Delta^R Pr(\hat{u}_t) = \frac{Pr(\hat{u}_t) - Pr(\hat{u}_{t-1})}{Pr(u_{t-1})}$$

One could argue that using relative changes instead of absolute changes has some advantages. Defining the job loss risk this way normalizes the changes by their level, meaning that a small change in an occupation with a generally low level of job loss probability will matter as much as a larger change in a riskier occupation.

The relative changes in job loss risk are highly correlated (0.6) with the absolute changes, but represent a slightly different distribution over time, being more positively skewed. The primary reason for this is that the denominator gets very close to zero for some observations. Since the OLS estimates are very sensitive to outliers, we remove the extreme observations before estimating. This reduces the standard deviation from .92 to .75 and the largest relative change goes down from 173.8 to 8.5. The row (5) of Table 1 show the results from estimating (4) using the relative changes in the probability as a measure of job loss risk. Relative changes in the predicted job loss risk also have a positive effect on the saving behavior of households. A change in the relative change of job loss risk equal to one standard deviation increase in absolute risk is associated with an increased saving of about 2.8 %. Comparing these estimates to the main specification, the implication for consumption is smaller. This is perhaps not so surprising as we are multiplying changes in job loss risk with the inverse of previous levels of job loss risk, and essentially down-weighting responses by individuals with high levels of job loss risk.

Alternative dependent variables: Saving rate

Finally, we check if our results are robust to other definitions of saving, such as the saving to income ratio. One might be concerned that we are not able to sufficiently control for the effect income has on differences in saving. Another concern might be the hyperbolic sine transformation we use to deal with large observations. To address both of these issues, we normalize the saving measure by the disposable income. We construct the saving to income ratio by simply dividing active saving by the family disposable income, S/Y. The results when using the saving rate as our dependent variable are shown in the final row of Table 1. We see that our results are robust to using this as the dependent variable, yielding largely comparable results to before. A one standard deviation increase in the job loss risk is associated with a larger saving to income ratio of about 0.8-1.0%. Alternatively, we can use the budget constraint of the household, $C/Y = 1 - S/Y$, and perform some back-of-the-envelope calculations. The median household in our sample has a saving to income ratio of about 0.05, meaning that the implied decrease in consumption is somewhere around 1%, an estimate close to the effect of the main specification.

5.3 Decomposing the saving response by assets and age

Our definition of active saving is the returns adjusted change in households' assets, as explained in Section 4.2. The level of detail in the data allows us to decompose the aggregate wealth measure into smaller components. We focus on debt and bank deposits in this section. There are several reasons for this. First, these are the assets of the household balance sheet where we have least reason to expect measurement error. Second, they do not contain high returns that distort the behavioral picture. Third, they are both fairly liquid, at least on a yearly basis, so we should expect that households use these actively when they adjust to changes in risk. Finally, it can be shown that the lions share of saving for most households is covered by these two assets.

The drawback of doing such a decomposition is that once we leave the whole identity of income, assets and consumption, we cannot infer whether a change in one of the assets is a reallocation between asset classes or a change in overall saving and consumption. An observed decrease in debt can for example be an extra down payment made from yearly income, where the household reduces its consumption, all other assets equal. Alternatively, the extra down payment can be made from drawing on deposits in a bank account, where the level of consumption is constant, but the total level of wealth has decreased. These two stories are observationally equivalent, and we are therefore not able to identify an underlying mechanism. We still believe it is useful to see how separate asset classes co-move with changes in the job loss risk, and the exercise gives us an impression of the relative importance of the two.

Table 2 shows the result of estimating the baseline relationship in equation (4) with decomposed dependent variables. The first row show the baseline result from Table 1, while the second row the

Table 2: Decomposed saving response to $\Delta Pr(\hat{u}_t)$ by asset class

Dependent variable:	OLS	(S.E.)	FE	(S.E.)
Total active saving	9.07**	(0.705)	8.67**	(0.693)
Saving in debt/mortgages	15.25**	(0.665)	14.11**	(0.654)
Saving in bank deposits	-4.74**	(0.653)	-5.18**	(0.645)
Observations	4 545 418		4 545 418	

Note: Same specification as benchmark in Table 1, except dependent variable.
The set of control variables include indicators for age, sex, marital status, children, education/type of work, local labor market region, birth country, self-employment, home-ownership, own labor income quintile, spouses' labor income quintile, and year.
* p < 0.05 ** p < 0.01

results using debt as the dependent variable. Saving in debt is constructed such that an increase in debt is defined as negative saving, i.e. a reduction of net wealth. A positive coefficient should therefore be interpreted as reduction of debt, or a down payment on a mortgage. According to Table 2, we find clear indication that when job loss risk increases, households reduce their leverage. Alternatively, one could interpret this the other way around: when the job loss risk is reduced, households are willing (or able) to increase their debt. We will return to, and study in more detail, possible asymmetrical responses to job loss risk in Section 5.4.

The last row of Table 2 show results using saving in bank deposits as the dependent variable. The amount of bank deposits is reported from the bank and is perhaps the most liquid asset the households have. Quite surprisingly, we see that households experiencing increased job loss risk, reduce the amount of deposits they have. Seen in combination with the movements in debt, we interpret this as households using bank deposits to reduce leverage when uncertainty is higher. Furthermore, there may be additional reallocation of funds to other financial assets not reported in the table. Thus, we conclude that we find evidence of both a reallocation within the portfolio as well as an overall increase in saving as job loss risk goes up.

Another interesting dimension to investigate the relationship between job loss risk and saving in, is age. Theoretically, the prediction is that younger households should be more responsive to changes in job loss risk. There are several reasons for this. Young households have usually not managed to build up large amounts of assets, and therefore they are less insured than wealthy

Figure 4: Decomposed saving response by asset class and age

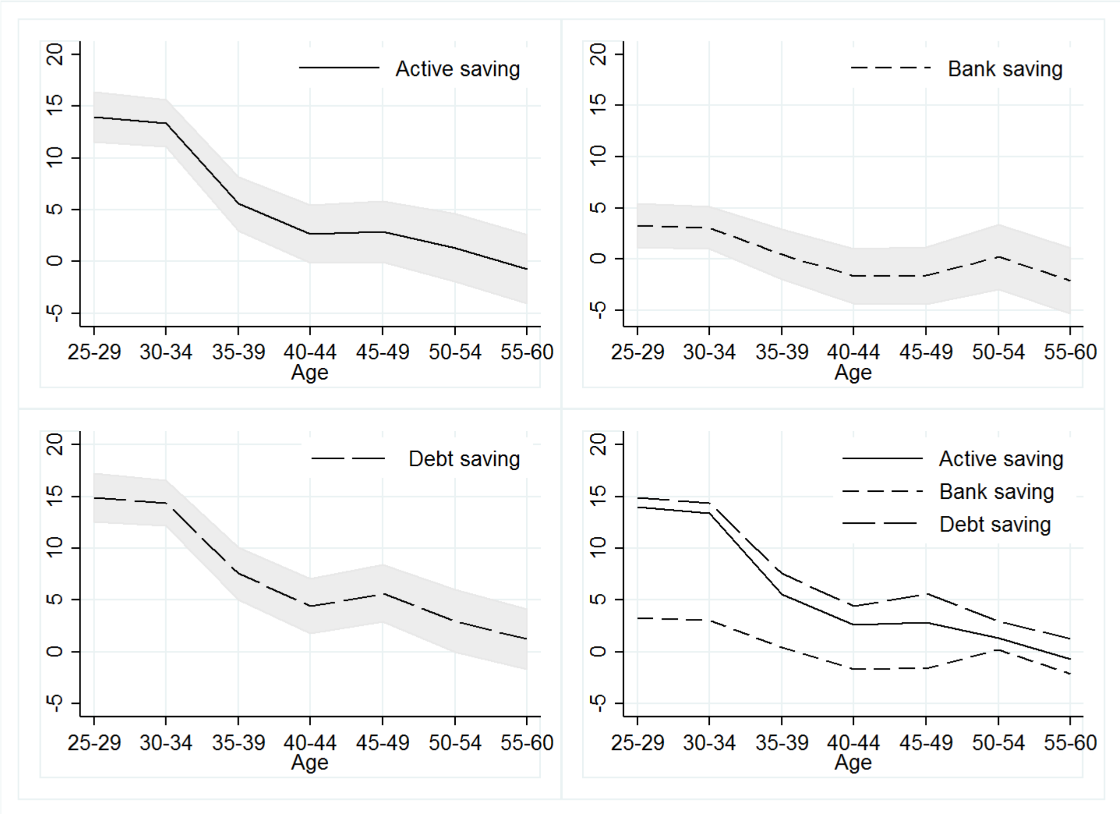

Note: Plots of the fixed effects coefficients and confidence intervals from the results in Table E.2 in Appendix E.

households.[24] This is often the case for older households, being so well insured that small changes in their job loss probability does not matter. Put differently, older households have already saved up most of their pension and will therefore not suffer large consequences if they lose their job. Since young households still have many years left before retirement, the consequence of unemployment is larger than for older households simply because a larger share of lifetime income is at risk. This intuition is reinforced by theories of the job ladder and unemployment scarring.[25] These theories suggest that entering unemployment at an early age has consequences for future unemployment risk and job market outcomes, amplifying the effect of job loss risk on life time earnings.

To investigate whether these groups behave differently, we split our sample into seven sub samples based on age, and perform the two steps in the main empirical specification outlined in Section 3. Furthermore, we also include the results decomposed by asset class, to get a better

[24] Some studies find that a significant share of households can be wealthy and credit constrained, so called wealthy hand-to-mouth (Kaplan and Violante, 2014).

[25] See for example Jarosch (2015) and Lise (2012) for interesting contributions.

overview of how they allocate their saving. The full set of results are presented in Appendix E.2, while the (fixed effects) coefficients are plotted in Figure 4.

Our results confirm that young households are more responsive to changes in job loss risk than older households. In fact, our results suggest that older households do not respond to changes in job loss risk at all. Households become unresponsive to job loss risk at the age of 40-44. This is completely in line with the findings of Gourinchas and Parker (2002), who find that households start accumulating significant amounts of wealth for retirement around age 40 and therefore stop acting as "buffer stock" agents (see Figure 6 in their paper).

Moreover, we find that young households respond primarily by adjusting their debt levels rather than by saving in bank accounts. As before, the two sets of coefficients does not add up to the coefficient for total saving, partly because they are a result of a combination of reallocation within these two assets and an increase in total saving, and partly because there may be reallocation to assets not presented in the figure (such as mutual accounts and other financial assets).

5.4 Asymmetric responses

Finally, we study if the saving response depends on whether there is an increase or a decrease in the predicted job loss risk. We remember from the Section 2 that the saving is determined by changes in the target wealth that a household wishes to keep in order to insure themselves from the potential spell of unemployment. This one-to-one mapping between levels of probability and target wealth implies that there should be the same amount of saving and dissaving for. However, while models typically assume such symmetry, loss aversion studies show that people react differently to equally sized losses and gains.[26] To test whether the responses are symmetrical, we simply interact an indicator being one if there is a positive change in the probability of job loss with the change, and run the baseline specification in equation (3). We remember from Figure 3 that in the aggregate, changes in job loss risk is nicely spread with a mean close to zero, meaning we have similar incidents of increases as decreases in the sample.

Table 3 shows the three coefficients of interest from this estimation. There is no level difference in saving between the groups experiencing increases and decreases in job loss risk, controlling for other factors. However, it is clear that our results are driven by decreases in job loss risk. The

[26] See, for example, Tversky and Kahneman (1991).

effect of job loss risk increases is not significantly different from zero.

Having previously seen that debt is the asset moving with job loss risk, in addition to the fact it is driven by young households, the asymmetric result reinforces the interpretation that job loss risk affect Norwegian households through ability or willingness to loan against future income. The finding that households are largely unresponsive to increases in job loss risk is surprising, but the generosity of Norwegian unemployment benefits could help us interpret these results: Since the replacement rate is high, the incentives for having a buffer against future income losses is reduced.

Table 3: Asymmetric active saving response to increases and decreases in $\Delta Pr(\hat{u}_t)$

	OLS	(S.E.)	FE	(S.E.)
$\Delta Pr(\hat{u}_t)$	19.44**	(0.990)	12.01**	(1.084)
$\Delta Pr(\hat{u}_t)$*Increase	-22.95**	(1.500)	-10.61**	(1.653)
Increase	-0.0287	(0.0162)	0.0246	(0.0176)
Observations	4 545 418		4 545 418	

Note: Same specification as benchmark in Table 1, except interaction with indicator for increase/decrease.
The set of control variables include indicators for age, sex, marital status, children, education/type of work, local labor market region, birth country, self-employment, home-ownership, own labor income quintile, spouses' labor income quintile, and year.
* $p < 0.05$ ** $p < 0.01$

6 Discussion

Our results indicate that households save more (consume less) when households with similar characteristics lose their job more often than in the previous period. Although our results are in line with earlier findings, our identification strategy is different in the sense that we control for confounding factors in a much more rigorous fashion than previous approaches. Compared to for example Campos and Reggio (2015), we have a risk measure that is more detailed, constructed from each individual's observable characteristics. Having connected a disaggregate risk measure to individual level saving adds confidence to what the mechanism behind our result is: Households respond to changes in their work prospects and future income, underlining that job security and income prospects are drivers for aggregate consumption. Studying the age profile, our results indicate that Norwegian households act as buffer stock agents until their forties.

Using our detailed data, we also understand *how* households save. In times when job loss risk increases, households are reluctant to take on new debt, and reduce the amount of liquid assets held in order to pay down their existing debt. It seems to be the case that changes in debt is a more common instrument to use than bank deposits. This might be a feature specific to the Norwegian context, as debt repayment contracts in Norway are flexible and easy to change. On the other hand, our results suggest that households are more willing (and able) to obtain debt for consumption in times when their jobs are safer. With that in mind, a relevant question is how much of the observed behavior is driven by choice, and how much is driven by variation in the presence of credit constraints? For example, Hendren (2017) shows that the reason why there is no private insurance market for unemployment risk is because of the information asymmetry between workers and insurance companies. If we believe that the same information asymmetry must exist between households and potential credit suppliers, there should be no role for time variation in credit constraints coming from changes in job loss risk. On the other hand, most banks require documentation of salaries and investigate the employment situation before handing out mortgages. It might therefore be the case that in times when households have safer jobs, credit is generally more accessible, which might explain why the effect is so strong for young households through changes in debt. Investigating whether these results can be explained by a liquidity channel in combination with job loss risk is a promising avenue for further research, but beyond the scope of the current paper.

7 Conclusion

In this paper we estimate the effect of job loss risk on saving and consumption. To do so, we propose a new way of identifying the mechanism, exploiting a long panel with data of Norwegian households' balance sheets. First, we estimate job loss probabilities based on observable characteristics. To avoid problems related to unobserved fixed effects and life cycle saving, we time differentiate the probabilities, as theory predicts that these changes should have a strictly positive relationship to households' saving.

We find that a one percentage point increase in the job loss probability leads to 9% increased saving. Back of the envelope calculations suggests this corresponds to a reduction in consumption of about 0.5% for the median household, which is similar to the results obtained by Campos and Reggio (2015). We further document that households are responsive to future job loss risk, which may be attributed to the fact that job loss uncertainty is present before actual job losses happen.

Decomposing the saving variable into its asset classes and focusing on debt and bank deposits, we document that households are more responsive through changes in their debt than changes in their bank deposits. We also study whether the effect is different for age groups, and find that young households act according to the buffer stock theory by being sensitive to changes in the job loss probability, an effect that is diminishing in age. In line with the results of Gourinchas and Parker (2002), we find that households age 45 and older are largely unresponsive to job loss risk.

Finally, we document that the effect is driven by households taking up debt in periods when job loss risk decreases. This suggests that in Norway, where the unemployment benefit system is very generous, job loss risk affects households willingness (or ability) to obtain debt financing during good times, rather than insuring against potential consumption drops in bad times.

References

Arulampalam, W., P. Gregg, and M. Gregory (2001). Unemployment scarring. *The Economic Journal 111*(475), 577–584.

Basten, C., A. Fagereng, and K. Telle (2016). Saving and portfolio allocation before and after job loss. *Journal of Money, Credit and Banking 48*(2-3), 293–324.

Benito, A. (2006). Does job insecurity affect household consumption? *Oxford Economic Papers 58*(1), 157–181.

Bhuller, M. S. (2009). Inndelingen av Norge i arbeidsmarkedsregioner. *SSB-notater*.

Blundell, R., B. Etheridge, and T. Stoker (2014). Precautionary Saving for Consecutive Life-Cycle Risks. *Working paper*.

Burbidge, J. B., L. Magee, and A. L. Robb (1988). Alternative transformations to handle extreme values of the dependent variable. *Journal of the American Statistical Association 83*(401), 123–127.

Campos, R. G. and I. Reggio (2015). Consumption in the shadow of unemployment. *European Economic Review 78*, 39–54.

Carroll, C., K. E. Dynan, and S. D. Krane (2003). Unemployment Risk and Precautionary Wealth: Evidence from Households' Balance Sheets. *Review of Economics and Statistics 85*(3), 586–604.

Carroll, C. D. (1992). The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence. *Brookings Papers on Economic Activity 1992*(2).

Carroll, C. D. and M. S. Kimball (1996). On the concavity of the consumption function. *Econometrica 64*(4), 981–992.

Carroll, C. D. and M. S. Kimball (2005). Liquidity Constraints and Precautionary Saving. *Manuscript, Johns Hopkins University*.

Carroll, C. D., J. Slacalek, and M. Sommer (2012). Dissecting Saving Dynamics: Measuring Wealth, Precautionary and Credit Effects. *ECB Working Paper No. 1474*.

Carroll, C. D. and P. Toche (2009). A tractable model of buffer stock saving. *NBER Working Paper No 15265*.

Challe, E., J. Matheron, X. Ragot, and J. F. Rubio-Ramirez (2017). Precautionary saving and aggregate demand. *Quantitative Economics 8*(2), 435–478.

Challe, E. and X. Ragot (2016). Precautionary Saving over the Business Cycle. *The Economic Journal 126*(590), 135–164.

Fagereng, A. and E. Halvorsen (2017). Imputing consumption from Norwegian income and wealth registry data. *Journal of Economic and Social Measurement 42*(1), 67–100.

Gourinchas, P.-O. and J. A. Parker (2002). Consumption over the Life Cycle. *Econometrica 70*(1), 47–89.

Gregory, M. and R. Jukes (2001). Unemployment and subsequent earnings: Estimating scarring among British men 1984–94. *The Economic Journal 111*(475), 607–625.

Guvenen, F., S. Ozkan, and J. Song (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy 122*(3), 621–660.

Hahm, J.-H. and D. G. Steigerwald (1999). Consumption adjustment under time-varying income uncertainty. *Review of Economics and Statistics 81*(1), 32–40.

Hendren, N. (2017). Knowledge of Future Job Loss and Implications for Unemployment Insurance. *American Economic Review 107*(7), 1778–1823.

Jarosch, G. (2015). Searching for job security and the consequences of job loss. *manuscript, Stanford University*.

Kaplan, G. and G. L. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica 82*(4), 1199–1239.

Kimball, M. S. (1990). Precautionary Saving in the Small and in the Large. *Econometrica 58*(1), 53–73.

Kirkebøen, L. J. (2010). Forskjeller i livsløpsinntekt mellom utdanningsgrupper.

Krueger, D., K. Mitman, and F. Perri (2015). Macroeconomics and Heterogeneity, Including Inequality. *Handbook of Macroeconomics (forthcoming)*.

Lise, J. (2012). On-the-job search and precautionary savings. *Review of Economic Studies 80*(3), 1086–1113.

Low, H., C. Meghir, and L. Pistaferri (2010). Wage risk and employment risk over the life cycle.

American Economic Review 100(4), 1432–67.

McKay, A. (2017). Time-Varying Idiosyncratic Risk and Aggregate Consumption Dynamics. *Journal of Monetary Economics 88*(June/July), 1–14.

McKay, A. and T. Papp (2011). Accounting for idiosyncratic wage risk over the business cycle. *Working paper, Boston University*.

Nilsen, Ø. A. and K. Reiso (2014). Scarring Effects of Early-Career Unemployment. *Nordic Economic Policy Review 1*, 13–45.

Ravn, M. O. and V. Sterk (2014). Job uncertainty and deep recessions. Technical report, Centre for Macroeconomics (CFM).

Stephens Jr, M. (2004). Job loss expectations, realizations, and household consumption behavior. *Review of Economics and Statistics 86*(1), 253–269.

Tversky, A. and D. Kahneman (1991). Loss aversion in riskless choice: A reference-dependent model. *The quarterly journal of economics 106*(4), 1039–1061.

Appendix A: The Norwegian unemployment benefit system

A person is eligible for receiving unemployment benefits if the following criteria are met:

The worker must

- have experienced a 50 % reduction in employment
- have earned at least 1.5 times the unemployment insurance basic unit (referred to as G) the previous year, or at least 3 G over the last three working years. In 2017 the level of one basic unit is about 93,500 NOK (approximately 11,500 USD).
- be registered as searching for a job, and document search activity every 14 days
- be defined as a job-searcher, which entails being willing and able to work anywhere in Norway

Some additional rules apply if the person is totally or partially laid off or unemployed due to bankruptcy. These are mainly clarifications regarding what is defined as a lay off and rules regarding wage insurance or severance payments.

Persons moving to or from another country or recently dismissed from initial services in armed forces are also subject to additional rules. These observations are presumably removed from our data set due to age- and living restrictions.

The person starts receiving unemployment benefits from the date he or she registers as seeking for a job, and will essentially have four weeks on getting all the documentation right before he or she starts losing any money. There is a 12 week delay before people that choose to resign from their job, or loses it as a consequence of own actions, becomes eligible for benefits.

The person receives 62.4 % of the gross income. The gross income is based on last year?s income, or the average of the past three years prior to applying. By gross income it is meant working income including any insurance benefits like sickness benefits, care benefits, pregnancy benefits and so forth. For any income exceeding 6 G, the person does not receive any unemployment benefits. Any working income during the unemployment spell will reduce the amount received. The length of which one can receive benefits is 104 weeks if the gross income exceeded 2 G, 52 weeks if the gross income was less than 2 G.

Appendix B: More on the identification strategy

We explain the identification strategy and its economic assumptions using the linear probability model, as this essentially becomes a regular two stage least squares with excluded time interactions. First we estimate the probability model yearly, letting all parameters vary over time:

$$u_{it} = \varphi_t V_{it} + v_i + \varepsilon_{it} \tag{7}$$

u_{it} is an indicator equal to one if the individual becomes unemployed. Note the distinction between becoming and being unemployed. We consider only workers that have become unemployed during the period, to avoid concerns related to long term unemployment. V_{it} is a set of predictors believed to influence job loss risk, such as age, labor market region, country of origin, sex, education/type of work, and income level. Lets further assume that V_{it} consist of indicators indicating membership in a specific group, so that there are no functional form assumptions involved. We estimate $\hat{\varphi}_t$ and use the model to predict the probability of becoming unemployed. Looking at a one period change in the predicted probability, we first point out that the unobserved fixed effect is removed.

$$\Delta Pr(\hat{u_{it}}) = \hat{\varphi}_t V_{it} - \hat{\varphi_{t-1}} V_{it-1} + \varepsilon_{it} - \varepsilon_{it-1} \tag{8}$$

Our identification relies on the fact that observed changes in predicted job loss risk consist of both exogenous and endogenous variation. Earlier studies that use level of job loss probabilities need an exclusion restriction in the V_{it} vector in order to obtain identification. Our identification comes from the fact that while the level of $Pr(\hat{u_{it}})$ to a large degree is endogenous, its evolution over time is not: How many people with a certain characteristic that lose their job is difficult to predict and affect for an individual. Since we observe how people with a given set of characteristics behave on average, we argue that we are able to control for them in the second stage saving regressions. Hence, what we measure in the second stage, is the effect of exogenous variation in job loss risk, through a set of characteristics, on saving.

The estimated change in the probability of job loss can be decomposed as in (9) to further understand what the sources of variation are.

$$\Delta Pr(\hat{u_{it}}) = \hat{\varphi}_t(V_{it} - V_{it-1}) + (\hat{\varphi}_t - \hat{\varphi_{t-1}})V_{it-1} + \Delta\varepsilon_{it} \tag{9}$$

We see that the one period change in probability of job loss consists of changes in the observable characteristics ($V_{it} - V_{it-1}$) scaled by coefficients, $\hat{\varphi}_t$, and changes in the coefficients of the observable characteristics ($\hat{\varphi}_t - \hat{\varphi_{t-1}}$) scaled by the set of characteristics V_{it-1}, and the first differenced error term. It is difficult to argue that changes in $Pr(\hat{u}_{it})$ coming from ($V_{it} - V_{it-1}$) are exogenous or unpredictable to the household, as these include changes in living region, education level or simply becoming older. These endogenous, or predictable, changes are all weighted by a set of exogenous "prices", namely $\hat{\varphi}_t$. As $\hat{\varphi}_t$ comes directly from the estimation of households losing their job, they are exogenous to a single household. For households to be responsive to these changes, we must assume that they have knowledge about them. The remaining variation in $\Delta Pr(\hat{u_{it}})$ comes from ($\hat{\varphi}_t - \hat{\varphi_{t-1}}$), which reflect the fact that individuals with different types of characteristics are at risk at different times. These changes are scaled by the set of household characteristics in the previous period, V_{it-1}, so only relevant changes in job loss risk are summed for the household.

If a household has fixed elements in V, the only variation comes through the exogenous changes ($\hat{\varphi}_t - \hat{\varphi_{t-1}}$). In that case, φ is clearly identified and will represent the difference in active saving between two otherwise equal households that can be explained by $\Delta Pr(\hat{u}_{it})$.

In the case where elements of V_{it} changes, one must assume that V_{it-1} affects present saving only through its impact on probability of job loss. One such element of V that changes each year is age. Every year, individual i becomes a year older, thus his or her job loss probability changes according to the estimated age profile. Since we estimate a year-specific age profile, there are differences in how the job loss risk changes depending on what year we are in. Consider for example two otherwise equal 30-year olds observed at time t and $t + n$. Since we include yearly indicators, in the second stage net of year specific effects, active saving should be equal for them independent of what year it is. What we pick up, is whether saving covary with job loss risk, conditional on these other features. Thus, the main threat to identification, is omitted variables that affect saving, but covary with time variation in individual job loss risk.

Appendix C: Data details

Table C.1 Classifications of educations

1	Elementary school	26	Bachelor level, Health care
2	Started junior high school	27	Bachelor level, Nursing
3	Completed junior high school	28	Bachelor level, Social work
4	Senior high school, Arts	29	Bachelor level, Therapeutical work
5	Senior high school, Chemical Technology	30	Bachelor level, Other
6	Senior high school, Media and Communication	31	Bachelor level, Humanities
7	Senior high school, Electricity and Electronics	32	Bachelor level, Social Sciences
8	Senior high school, Technical studies	33	Bachelor level, Science
9	Senior high school, Building and Construction	34	Other higher 3-4 year educations
10	Senior high school, Industrial Production	35	Master level,Humanities
11	Senior high school, Health care	36	Master level, Theology
12	Senior high school, Agriculture	37	Master level, Social Sciences
13	Senior high school, Transport	38	Master level, Law
14	Senior high school, Service	39	Master level, Economics and Business Admin.
15	Senior high school, Other	40	Master level, Engineering
16	College, Economics and Business Admin.	41	Master level, Architecture
17	Technical college	42	Master level, Science
18	College, Other	43	Master level, Medicine
19	Bachelor level, Pre-school Teacher	44	Master level, Dentistry, Veterinary, and Pharma
20	Bachelor level, Teacher	45	Master level, Agronomy
21	Bachelor level, Journalism	46	Military education
22	Bachelor level, Economics and Business Admin.	47	Other higher 5-6 year educations
23	Bachelor level, Siviløkonom	48	PhD, Science
24	Bachelor level, Engineering	49	PhD, Medicine
25	Bachelor level, Maritime Subjects	50	PhD, Other

Table C.2 Classification of labor market regions

11	Sør-Østfold	51	Sunnfjord
12	Oslo	52	Sognefjord
13	Vestfold	53	Nordfjord
14	Kongsberg	54	Søndre Sunnmøre
15	Hallingdal	55	Ålesund
21	Valdres	56	Molde
22	Gudbrandsdalen	57	Nordmøre
23	Lillehammer	58	Kristinsund
24	Gjøvik	61	Trondheim
25	Hamar	62	Midt-Trøndelag
26	Kongsvinger	63	Namsos
27	Elverum	64	Ytre Helgeland
28	Tynset/Røros	65	Indre Helgeland
31	Nordvest-Telemark	71	Bodø
32	Øst-Telemark	72	Narvik
33	Sør-Telemark	73	Vesterålen
34	Arendal	74	Lofoten
35	Kristiansand	75	Harstad
36	Lister	76	Midt-Troms
41	Stavanger	77	Tromsø
42	Haugesund	81	Alta
43	Sunnhordaland	82	Hammerfest
44	Bergen	83	Vadsø

Appendix D: Additional plots of job loss probability over time

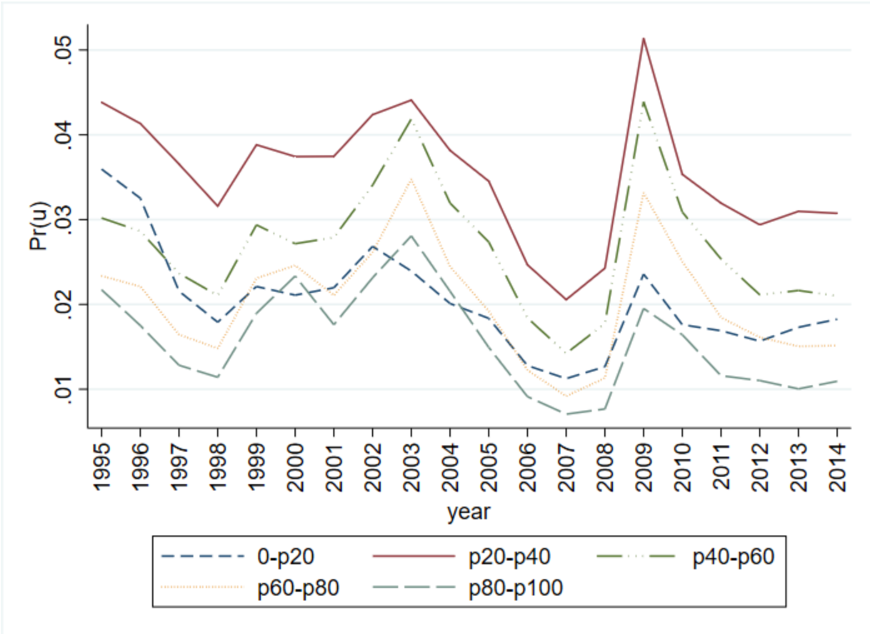

Margins plot of the probability of job loss over time, by labor income quantiles. The lower quantile behaves somewhat different due to labor income categorisation. Some individuals report low labor earnings, due to selfemployment and tax purposes. They seldom become unemployed, biasing the lowest earning group's job loss risk downwards.

Table E.2 Saving response by age and asset

	25-29		30-34		35-39		40-44		45-50		51-55		56-60	
	OLS	FE	OLS	FE	OLS	FE	OLS	FE	OLS	FE	OLS	FE	OLS	FE
							Active Saving							
$\Delta Pr(\hat{U}_t)$	20.05***	13.98***	18.87***	13.36***	8.560***	5.572***	6.069***	2.657	4.964***	2.882	3.714**	1.321	0.0391	-0.698
	(0.804)	(1.223)	(0.843)	(1.142)	(0.994)	(1.336)	(1.079)	(1.424)	(1.143)	(1.512)	(1.283)	(1.679)	(1.372)	(1.691)
							Saving in Bank Deposits							
$\Delta Pr(\hat{U}_t)$	3.609***	3.284**	3.605***	3.064***	0.676	0.446	-1.408	-1.653	-2.153*	-1.620	0.400	0.210	-1.906	-2.0900
	(0.696)	(1.082)	(0.754)	(1.045)	(0.906)	(1.240)	(1.002)	(1.349)	(1.059)	(1.421)	(1.213)	(1.607)	(1.318)	(1.636)
							Saving in Debt							
$\Delta Pr(\hat{U}_t)$	21.06***	14.87***	21.00***	14.36***	11.82***	7.570***	8.780***	4.414**	9.100***	5.642***	6.103***	3.014*	3.488**	1.231
	(0.784)	(1.189)	(0.820)	(1.111)	(0.962)	(1.291)	(1.030)	(1.358)	(1.067)	(1.407)	(1.173)	(1.533)	(1.207)	(1.484)
Lagged dep var	No	No	No	No	No	No	No	No	No	No	No	No	No	No
V_{it} included as controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	370,396	370,396	553,628	553,628	573,686	573,686	558,698	558,698	531,606	531,606	496,263	496,263	493,345	493,345

CHAPTER II

Gambling with the family silver

Household consumption and saving responses to fiscal uncertainty

Gambling with the family silver

Household consumption and saving responses to fiscal uncertainty*

Oddmund Berg†

Abstract

In the wake of the 2007 stock market crash, news emerged that eight Norwegian energy producing municipalities had sold up to ten years of future earnings from their hydro-electric power plants and invested it in high-risk financial products. Some of these municipalities lost more than 80 percent of the invested amount. I use a difference in difference analysis to show that this led to a reduction in private consumption of around 2 percent, the following year. I show that the response is driven by households who are the largest recipients of public services - the young and the elderly. The reduction in consumption is a result of households holding back income by saving more of their disposable income than before. Finally, households in the affected municipalities rebalance their portfolios to holding a lower share of risky assets. While public spending decreased permanently in the years following, the consumption effect was temporary. I therefore conclude that the effect is driven by households holding back consumption until uncertainty regarding fiscal outcomes is resolved.

Keywords: Fiscal uncertainty; consumption and saving; panel data; heterogeneity

JEL codes: D12; E21; E65; H31

*I gratefully acknowledge comments and suggestions from Gernot Doppelhofer, Andreas Fagereng, Jon H. Fiva, Elin Halvorsen, Martin B. Holm and Kjetil Storesletten. I also thank seminar participants at Statistics Norway, the NHH-UoB PhD-workshop and the Norwegian Business School PhD research seminar series.

†Norwegian School of Economics

1 Introduction

When the future path of government policy becomes more uncertain, households might delay consumption until this uncertainty has been resolved.[1] In this paper, I exploit an unexpected shock to the local government budget of eight Norwegian municipalities to investigate the effect of fiscal uncertainty on households' consumption and saving.

While a growing literature considers uncertainty as an explanation for macroeconomic fluctuations,[2] only a few studies have looked at the effect of uncertainty shocks on households' consumption and saving using micro data. A reason for this is that unexpected events involving policy uncertainty, and some degree of randomness, are rare.[3] One exception is Aaberge, Liu, and Zhu (2016), who study an event involving political uncertainty and reforms in China and find that households increase their saving. Another example is Giavazzi and McMahon (2012), who use a survey measure of subjective uncertainty to show that uncertainty related to the 1998 political election in Germany induced affected households to save and work more.

The current paper considers uncertainty related to the future fiscal policy of local governments in Norway. In the wake of the 2007 stock market crash, news emerged that eight Norwegian hydro-electric power producing municipalities had sold future power-income and invested it in high-risk financial products. It is well documented that these investments, and their high-risk profiles, were unknown to current local government electives and inhabitants (Hofstad, 2008).[4] Within a few months after the crash, the assets were sold with large losses, incurring deficits for the involved municipalities. According to the Norwegian law, these deficits had to be covered within four years, leading to a sharp tightening of the local government budget constraint. The following year contained uncertainty regarding future fiscal outcomes in the affected municipalities. This mirrors Bratberg and Monstad (2015), who study the same event as the current paper, looking at sick leave. They find that sick leave decreased in the affected municipalities in the weeks after the event, driven by the perception that jobs were more uncertain.

I use a household level measure of private consumption, imputed from Norwegian register

[1] Bloom (2014) gives a review of the literature on uncertainty.

[2] See for example Bloom (2009), Baker, Bloom, and Davis (2015), Caldara, Fuentes-Albero, Gilchrist, and Zakrajšek (2016) and Basu and Bundick (2017)

[3] Fuchs-Schündeln and Hassan (2016) provides an overview of natural experiments in macroeconomics, highlighting that changes in fiscal policy not motivated by business-cycle considerations are difficult to find.

[4] They had been made in the early 2000s and were subsequently managed in full by an external brokerage.

data, to show that inhabitants in the worst affected municipalities reduced their consumption by 1.8 percent in the first year following the scandal. This is in stark contrast to inhabitants in the municipalities who experienced losses, but managed to cover them relatively easily. In these municipalities, there was no effect on private consumption. Therefore, the consumption response seems to be related to economic turbulence caused by the event, and not the loss of assets itself. Furthermore, the fall in consumption is driven by age-groups who are the largest recipients of public goods - the young and the elderly. In other words, groups that where most likely to be affected from changes in public services, are the ones who act more careful.

To study the underlying actions that drive the consumption reduction, I look at saving in separate asset classes. The change in consumption comes from increased saving, meaning households spend less from their income than their comparison group. This is similar to the findings of Giavazzi and McMahon (2012) and Aaberge, Liu, and Zhu (2016), who both find that policy related uncertainty induce affected households to save more.[5]

To investigate whether uncertainty is the driver, I follow the literature on portfolio choice and background risk and study the share of risky assets held by the affected households (see for example Heaton and Lucas (2000), Palia, Qi, and Wu (2014) and Fagereng, Guiso, and Pistaferri (2017)). I find that households holding risky assets rebalance their portfolios after the event, holding a lower risky share one year after. Like the findings on consumption expenditures, this effect is only present in the worst affected municipalities.

Since the loss of publicly owned assets may carry other effects than uncertainty, I consider alternative explanations for the observed results. For example, a large literature considers the relationship between public spending and private consumption.[6] I provide a short theoretical framework to illustrate how local government spending decisions can affect private consumption, through combined insights about the local government budget constraint and the Permanent Income Hypothesis. Using municipality income statements, I show that public spending falls permanently over the period of analysis. If the changes in private spending were driven by actual changes in public spending, the decrease in private consumption should have been permanent as

[5]See for example Carroll and Kimball (2006) for a survey on the precautionary saving literature.

[6]Examples include Blanchard and Perotti (2002), Perotti (2007), Mountford and Uhlig (2009), Fatás, Mihov, et al. (2001), Galí, López-Salido, and Vallés (2007) using VAR-techniques. Ramey and Shapiro (1998) and Giavazzi and Pagano (1990) are examples using a narrative approach.

well. I therefore argue that the actual fiscal actions of the local governments are not the driver of the observed consumption drop. Considering other alternative explanations, I show that there is no effect on moving, employment or disposable income. I therefore conclude that the effect is driven by households holding back consumption until uncertainty regarding fiscal outcomes is resolved.

The paper proceeds as follows: Section 2 tells the story of the events, and describe how this is an uncertainty shock. Section 3 explains the data and imputation process, section 4 the empirical strategy. In section 5 I show the main results, provide a theoretical interpretation and study the underlying mechanism. I summarize and conclude in section 6.

2 Background and shock interpretation: The Terra-scandal

October 31st 2007, news emerged that eight Norwegian municipalities had sold ten years of future power income from their hydro-electric power stations, to gear and invest in high risk financial products. The deal was initiated and managed by two brokers from Terra Securities, who initially approached all 174 power producing municipalities in the early 2000s. From 2001-2007, the brokers invested in high risk financial products on behalf of the eight municipalities. Consequently, very few people had knowledge about the investments, and their risk profiles, when attention was brought to them in a national newspaper in 2007. By the end of 2007, news about the municipalities and their huge losses had reached international newspapers such as The New York Times and The Washington Post. By early spring and summer of 2008, most of the assets were sold, incurring large deficits. According to the Norwegian law, municipalities must cover such losses within a period of four years. This led to a tight economic situation for five of the eight affected municipalities.

Understanding what type of shock this is, and what consequences it carried for the involved municipalities, is essential for interpreting its effect on private consumption. Was this a shock sizeable enough to affect the inhabitants? Table 1 shows the amount invested and lost across municipalities, both in absolute terms and per capita. There is large variation, reflecting that municipalities vary in their population and financial situation.[7] Kvinesdal, Rana and Haugesund

[7]Since the hydroelectric power plants may range over several municipalities, some of them have shared ownership and split their taxes and fees between these. The fees are set by the authorities and paid separately by each plant,

managed to cover their losses relatively easily. Rana because of large income stream, Kvinesdal because of small exposure and large income, Haugesund as a combination of the two. Bremanger, Hattfjelldal, Hemnes and Narvik were put under state supervision shortly after the news broke in 2007. Vik was already under supervision two months before the news became public.

Being put under state supervision means that either the County Governor or the Ministry is to review the legality of the budget resolution passed by the municipal council or the county council, in addition to loan and financial leasing and long term rental contracts. Being under supervision decreases economic freedom, and is either a consequence of running deficits, or simply based on a general assessment of the economic state of the municipality. Between 2001 and 2017, between 42 and 118 out of Norway's 424 municipalities were under supervision. As a fairly common occurrence, it should not be used as an indicator of economic turbulence independently. On the contrary, the fact that three of the eight affected municipalities were not on the list after experiencing the deficits, is a strong indication of them being largely unaffected.

In addition to being put under state supervision, four of the municipalities received extraordinary transfers from the central Government to be able to uphold their legally required service level. This confirms that some municipalities had real economic struggles as a consequence of the losses. While these transfers reduced the economic importance of the event, they were not granted before mid November 2008, leaving households in a state of uncertainty regarding the economic future of the municipality for a long time. In this period, there was reason to fear that services could fall below the legally required level.

In the following, I split the analysis of the eight affected municipalities into two. I base this split on the narrative provided in the annual reports, news stories and the three columns to the right in Table 1. If the municipality received extraordinary transfers, it is clear that they struggled economically after the shock. I also include the municipality of Hemnes, which ended up under state supervision immediately after, indicating that even though it did not receive extraordinary transfers, it faced challenges.[8] Its loss per capita is the fourth largest. In the remainder of the paper, I label the five municipalities Bremanger, Hattfjelldal, Hemnes, Narvik and Vik "The five worst hit" and Haugesund, Kvinesdal and Rana "The three unaffected". This split gives an extra

making it complicated to provide an overview of the income from hydroelectric power plants by municipality.

[8]Moving Hemnes from the worst hit does not qualitatively change the results.

Table 1: Investment, loss and population stats

Municipality	Investment	Loss	Population	Loss per capita	Under state supervision (ROBEK)	Extraordinary transfers
Rana	297	222.5	25 124	8 900	-	No
Hattfjelldal	103	85	1 472	57 750	2008	Yes
Hemnes	84	78.3	4 500	17 400	2008	No
Narvik	242	188.3	18 391	10 240	2009	Yes
Vik	149	90	2 816	32 000	2007	Yes
Bremanger	350	217.7	3 903	55 700	2008	Yes
Haugesund	227	130	32 761	3 970	2010	No
Kvinesdal	43	18	5 622	3 200	-	No

Investment and loss are given in mill NOK. The conversion to dollar at the time was approximately 5.8.
Population was measured in the fourth quarter of 2007.
ROBEK - Register for Governmental Approval of Financial Obligations. For any registered municipality, either the County Governor or the Ministry is to review the legality of the budget resolution passed by the municipal council or the county council, in addition to loan and financial leasing and long term rental contracts.

layer of identification, allowing me to separate the news-shock element from uncertainty related to fiscal outcomes.

3 Data

The data comes from the Norwegian Tax Register, containing information on households' balance sheets in the period from 1994-2010. Households in Norway are subject to income and wealth tax, and are therefore required to report their wealth holdings and income to the tax authorities. This information is not self reported, but comes from third parties such as employers, banks, financial intermediaries and insurance companies, covering their complete wealth holdings. The data is therefore not subject to personal reporting and has the advantage of being precisely measured. The data gives a complete picture of asset ownership over time that can be decomposed into different categories. Furthermore, through personal identifiers, the income and wealth registers can be linked to other registers with information on personal characteristics, such as place of living, age, education, family status, number of children and immigration background.

The data is produced using conventional methods, making use of family identifiers from the population register to aggregate income and wealth to family level. Based on the address register, I keep households living in municipalities that are on the membership list of the National Association of Hydro-electricity in 2008.

I also include the Kostra-code, which is a grouping of municipalities based on population and financial status. The coding is used for comparing municipalities with each other, and is updated more or less every five years. I use the closest available update, which is 2008. One might worry that the treated municipalities are affected by the shock in 2007 in a way that moves them to another Kostra-group in 2008, however, this is the case only for Bremanger. Given that many of the potential control municipalities change code from 2003 to 2008, I stick with the 2008 version to have the most relevant classification for the year of treatment.

Finally, I use the annual reports and income statements of the municipalities, containing information on income, expenses and financial transactions broken down by categories, for the periods 2003-2012.

Imputing consumption from the administrative data

The method of imputing consumption from administrative tax records was pioneered by Browning and Leth-Petersen (2003) and is well documented and tested (see Fagereng and Halvorsen (2017) for details on the method using Norwegian data, and Eika, Mogstad, and Vestad (2017) for an assessment. Koijen, Van Nieuwerburgh, and Vestman (2014) tests quality using similar Swedish data.). Examples of applications with the Norwegian data are Kostøl and Mogstad (2015) and Fagereng, Holm, and Natvik (2016). The method has several advantages over using survey data, as expenditures are objectively measured and sample size covers the whole population. In the current application this is extremely important, as the number of households affected by the Terra-scandal is relatively small compared to the whole population.

The idea is to exploit the budget constraint of households:

$$A_t = (1 + r)A_{t-1} + Y_t - C_t$$

where the observed assets at time t, A_t, is the assets at the start of the period with interest, $(1 + r)A_{t-1}$, plus disposable income, Y_t, minus consumption expenditures, C_t. Rewriting, we see that the period consumption expenditures are simply the disposable income minus the change in the stock of assets and their returns:

$$C_t = Y_t - \Delta A_t + rA_{t-1} \tag{1}$$

This means that the consumption expenditure is everything the household does not actively save from their reported disposable income. Note that local government taxes and user fees are charged directly to the household and not captured in the disposable income. Ceteris paribus, an increase in real estate taxes or user fees will increase consumption expenditures.

Implementing this strategy is not completely straight forward. First, if we are not able to separate changes in A_t stemming from an active decision to save more of the income, from passive returns to existing assets rA_{t-1}, consumption will be underestimated in the case of large returns and vice versa. A case where this is problematic is stocks. Since we only observe the total value of stock-holdings at the end of the year, we do not know if an increase is due to the household buying more stocks, or if it is simply unrealized returns as a consequence of price changes in the stock market.

Another issue is the value of houses and other real estate. Since they are not transacted often, obtaining reliable estimates of their current market values can be difficult. These are registered in the tax records with a rough estimate that is meant to resemble about 25 % of its value. As mentioned, these values are seldom updated, and therefore imprecisely measured. Years with housing transactions are problematic, as the debt used to finance the purchase does not match the asset side of the equation, and we tend to get large and unrepresentative observations of the consumption expenditure. Since we want unrealized returns removed from the equation, one remedy is to remove observations where households make housing transactions.

Finally, due to the aggregation to family level, the imputation is sensitive to changes in family composition. These changes can to some extent be controlled for in the regressions, but may nonetheless create large variation in the consumption measure over time.

The aforementioned issues may lead to over- or underestimation of consumption. However, since I operate in a difference in difference environment, bias from measurement error will not influence my estimates under the assumption that the treatment and measurement errors are independent of each other. Since the sources of measurement error are global, it is reasonable to believe that this is the case.

The additional noise might reduce the ability to measure small effects. I therefore follow Fagereng, Holm, and Natvik (2016) and make exclusions to create a sample not suffering from extreme observations to obtain reasonable precision of my estimates. They show that the qualitative results are the same when these are included, however the estimates have larger standard errors, likely as a consequence of measurement error. The exclusions include removing negative imputed consumption values and removing the lowest and largest percentile.

Setting up control groups

The critical assumption in the difference in difference set-up is the parallel trend assumption, namely that the control group and treatment would have the same development in the outcome variable if it were not for the treatment. I ensure similarity between the treated and controlled by performing two simple steps: First, I follow Bratberg and Monstad (2015) and pick from the pool of households living in a municipality with income from hydro-power plants. About 173 out of 426 municipalities had income from hydro-power plants in 2008. In other words, this strategy ensures similarity between the municipalities in one important dimension, namely that financing of some municipal functions comes from power plant income that could have been invested and lost in the financial market.[9]

I also ensure that the population and financial status of the municipalities are comparable. The reason for this is clearly stated in Table 1, where we see the large differences between municipalities.[10] For example, Haugesund has ten times the population of Vik, and a different level of spending per capita. Therefore, for each treated municipality, I create a control group from the pool of power-producing municipalities consisting of households living in a municipality with the same Kostra-code as the treated group. The Kostra-code is based on population and financial status used for comparison purposes.[11] For example, when municipalities report their economic performance, they use their Kostra-group as a reference group.[12]

Table 2 shows descriptive statistics of the dependent and control variables of the final samples employed, disaggregated for the applied definitions of treatment and control groups. The control and treatment groups look similar in all observable dimensions.

[9] Another way of viewing it, is that I exclude all households that never had a positive probability of being treated.
[10] This is also highlighted in the paper by Aaberge and Langørgen (2003)
[11] See Aaberge and Langørgen (2011) for a description (in Norwegian)
[12] An alternative could be to perform matching, however, the large amount of data, and set of possible matching techniques, seem lees transparent than the current set-up.

Table 2: Descriptive statistics

Variable	Mean		Stdev		Min		Max		Obs	
	Control	Terra	Control	Terra	Control	Terra	Control	Terra	Control	Terra
The five worst hit										
Log Consumption	11.94	11.98	.70	.71	2.31	3.15	15.62	15.46	1 012 260	126 353
Consumption to income	1.13	1.13	.61	.62	.168	.169	6.13	6.09	992 317	123 708
Consumption to lagged income	1.11	1.11	.536	.543	.171	.173	5.89	5.86	639 495	78 154
Saving in debt to income*	-.055	-.067	.394	.485	-4.03	-4.04	.902	.877	994 449	123 655
Saving in financial assets to income*	-.126	-.136	.607	.619	-5.12	-5.08	.831	.830	992 317	123 708
Conditional Risky share	.256	.236	.263	.256	1.43e-06	1.20e-06	1	1	236 131	24 649
Log Disposable Income	11.91	11.95	.558	.454	2.30	2.43	15.62	15.46	1 012 260	126 353
Age	56.22	55.60	19.47	19.04	25	25	90	90	1 012 260	126 353
Education Length	3.33	3.46	1.57	1.61	0	0	9	9	1 012 260	126 353
Children	.28	.28	.73	.71	0	0	10	7	1 012 260	126 353
Family Size	1.83	1.79	1.19	1.15	1	1	12	14	1 012 260	126 353
Male	.46	.46	.498	.498	0	0	1	1	1 012 260	126 353
The three unaffected										
Log Consumption	12.03	12.02	.716	.727	2.511	2.344	15.77	15.97	685 847	247 673
Consumption to income	1.13	1.14	.62	.63	.169	.169	6.13	6.13	671 922	242 758
Consumption to lagged income	1.11	1.12	.547	.549	.171	.171	5.89	5.90	416 459	151 600
Saving in debt to income	-.068	-.070	.448	.548	-4.04	-4.03	.90	.90	670 424	242 177
Saving in financial assets to income	-.134	-.142	.623	.628	-5.12	-5.12	.831	.831	671 922	242 758
Conditional Risky share	.282	.281	.274	.272	1.22e-06	3.60e-06	1	1	173 671	61 805
Log Disposable Income	12.00	11.99	.558	.588	2.30	2.31	15.77	15.98	685 847	247 673
Age	52.82	53.42	18.60	18.84	25	25	90	90	685 847	247 673
Education Length	3.55	3.59	1.66	1.67	0	0	9	9	685 847	247 673
Children	.319	.305	.750	.741	0	0	11	8	685 847	247 673
Family Size	1.82	1.75	1.16	1.14	1	1	14	10	685 847	247 673
Male	.461	.460	.498	.498	0	0	1	1	685 847	247 673

4 Empirical strategy

To estimate the effect on consumption, I use the following specification,

$$ln(c_{it}) = \alpha_0 + \sum_{j=1995}^{2010} \theta_j year_j + \delta_0 T_{it} + \delta_1 T_{it} * D_{2008} + \delta_2 T_{it} * D_{2009} + \delta_3 T_{it} * D_{2010} + X_{it}\beta + \gamma M_j + \varepsilon_{it} \quad (2)$$

where α_0 is the intercept, θ_j is a set of yearly effects, T_{it} is an indicator being one if the individual lives in an affected municipality and 0 otherwise. $T_{it} * D_{2008-2010}$ are interaction variables for estimating the yearly treatment effects in the periods after the news emerged. The effect is given by δ_{1-3}. X_{it} is a set of control variables containing observable characteristics of the household, γM_j is a vector of municipality fixed effects and ε_{it} is the error term. In the main results, controls include indicator variables for education type[13], education length, age and family size. The log of disposable family income is entered as a linear control.

I also estimate the same specification, but instead of assuming parallel trends, I include interaction effects from 2003-2010, to give a better illustration of the parallel trend.

In a separate robustness section, I allow for differing linear trends between the treatment and

[13] A 10-category variable based on the Norwegian Educational standard.

Table 3: Interaction effects on log consumption

	The five worst hit	*The three unaffected*
T * D2008	-.0179***	-.00432
	(.00443)	(.00896)
T * D2009	-.00310	-.00595
	(.0061)	(.0211)
T * D2010	-.00904	.00290
	(.0093)	(.0154)
Observations	1 138 613	933 520

Note: Control variables: Age-indicators, education length- and type fixed effect, municipality fixed effect, year fixed effect, immigration category and household size.
Cluster robust standard errors on treatment level (municipality)
* $p<0.05$, ** $p<0.01$, *** $p<0.001$

control group. I vary the set of control variables, and perform the conventional placebo tests, estimating 1-year effects to assess whether significant differences occur when there should be none. Following Cameron and Miller (2015), I cluster all standard errors at the municipality level, since this is the level at which the treatment is assigned.[14]

5 Results

Table 3 shows the treatment effects from the formal difference in difference estimation. The first column shows the effect in the five worst hit, and the second shows the effect in the three unaffected. There is a 1.79 percent decrease in consumption in the five worst hit municipalities in 2008, however, there is no significant effect in 2009 and 2010. In the three unaffected municipalities, we see no significant effect, with point estimates no larger than 0.6 percent.[15] To assess the economic significance of the results, consider that the mean consumption in the five worst affected municipalities was 265 000 NOK in 2007, meaning the consumption reduction was 5000 NOK ($710) per capita.[16]

Since the parallel trend assumption is crucial, I re-estimate equation (2) and include interaction

[14]This increases the standard errors compared to clustering on the individual level.
[15]Some might worry that the difference in precision comes from differing number of clusters. The same conclusions apply when clustering on the individual level.
[16]The conversion to USD was 6.97 the 31st of December 2008.

Figure 1: Main results with pre-treatment interactions

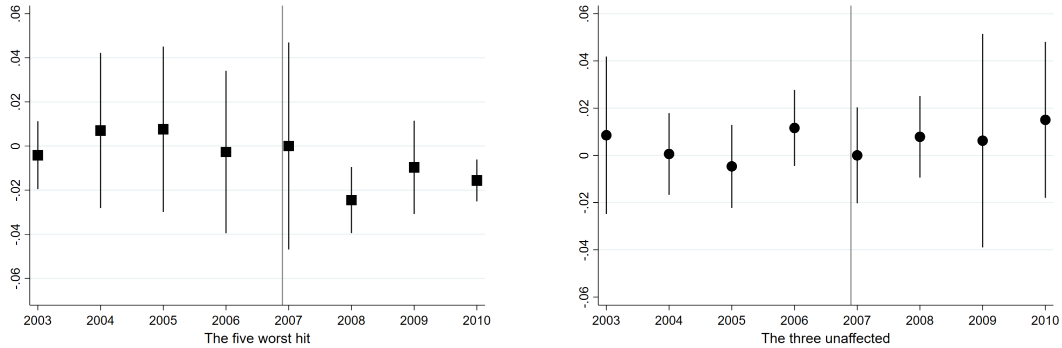

Note: Yearly interaction effects for 2003-2010, normalized to zero in 2007, 95% confidence interval.

terms for the pre-treatment periods 2003-2007. I plot the interaction coefficients and their 95 percent confidence intervals in Figure 1, normalizing the interaction effect in 2007 to 0. We see that in the periods leading up to the event, consumption was slightly higher than the average difference between the treatment and control group. In 2008, two things happen; There is a distinct fall in consumption, and precision increases. The reduction of noise could be sampling issues, but it could also be a consequence of household's behavior. If households generally act more carefully, in the sense that fewer large transactions occur, this could tighten the distribution and increase the precision of the estimates.[17] The figure shows that the drop in consumption compared to 2007 is more than 2 percent. The reason why the effect reported in Table 3 is lower, is that the difference between the treated and control group is slightly lower in 1995-2003 than that observed between 2003 and 2007. Compared to 2007, the reduction in 2010 is significant, suggesting that the sudden drop in 2008 may have had some persistence.

An important question is whether the drop in consumption is driven by changes in income. I therefore explore an additional specification where I use the ratio of consumption c_{it} to disposable income y_{it},

$$\frac{c_{it}}{y_{it}} = \alpha_0 + \sum_{j=1995}^{2010} \theta_j year_j + \delta_0 T_{it} + \delta_1 T_{it} * D_{2008} + \delta_2 T_{it} * D_{2009} + \delta_3 T_{it} * D_{2010} + X_{it}\beta + \gamma M_j + \varepsilon_{it} \quad (3)$$

[17]Another reason for the increased noise before the event could be the tax reform in 2006, where wealthy adjusted their income and wealth to avoid the introduction of tax on dividends. It is not apparent why this should lead to less precise estimates only among the five worst hit. I have also performed stricter trimming of the data, and the pattern remains.

Other than the change of dependent variable, the specification is equal to equation (2). Table 4 shows the effect on normalized consumption by disposable income. Since disposable income affects current consumption mechanically, I show that the same result appears when using lagged disposable income as dependent variable (Table 6 in the Appendix). Normalized consumption drops by more than 2.5 percent in the five worst hit municipalities. This confirms that the decrease in consumption is not a mechanical result from households experiencing lower disposable income. The estimation displays a similar pattern to the main specification, underlining that the effect is driven by households in the five worst hit consuming less of their income in the first year following the event. Furthermore, the zero-effect on income is confirmed in the top right panel of Table 4, using log disposable income as the dependent variable.[18]

To be able to further understand the mechanism underlying the consumption reduction, I decompose the imputed consumption measure (see equation (1)) and test for differences in the two main types of active saving, namely debt and financial assets: $\Delta Assets = \Delta Debt + \Delta Financial\_assets$. Because changes in the level of assets can be negative, the logarithmic transformation cannot be employed. Normalizing by income showed the same qualitative results as log consumption in the main specification, so I normalize saving in debt and financial assets by income to reduce the impact of extreme observations.[19] Other than the change of dependent variable, the specification is equal to equation (2). The results are shown in the bottom panel of Table 4, and show that the difference in consumption comes from differences in the saving in debt to income ratio. Households in the affected municipality seem to be less willing, or less able, to obtain debt. Alternatively, they reduce their existing debt by making extra down payments when the shock occurs. There is no significant effect on saving in financial assets.

Robustness of the main results is deferred to section 7.1. Here, Table 8 shows that there are no other significant differences between the groups at a 0.1 percent level in earlier periods. In other words, the drop in consumption is not estimated by chance. Furthermore, Table 9 shows that the effect is robust to inclusion of differing linear trends between the treatment and control group. Varying the set of control variables changes the point estimates slightly, however, the qualitative pattern remains.

[18]Somewhat puzzling, there seems to be a significant increase in disposable income among the three unaffected.
[19]Since this creates some outliers, I trim the ratios, removing the 1th and 99th percentile observations.

Table 4: Effect on consumption to income ratios and decomposed consumption

	Consumption to disposable income		Log disposable income	
	Five worst hit	*Three unaffected*	*Five worst hit*	*Three unaffected*
T * D2008	-0.0257***	-0.0101	-0.0202	0.0227*
	(.0051)	(.0117)	(.0166)	(.008)
T * D2009	-0.00237	-0.00441	-0.0104	0.0218**
	(.0061)	(.0191)	(.0095)	(.0061)
T * D2010	-0.00242	-0.00139	0.00038	0.0204*
	(.0133)	(.0197)	(.0063)	(.0069)
Observations	1 116 025	914 680	1 138 613	933 520

	Saving in debt to income		Saving in financial assets to income	
	Five worst hit	*Three unaffected*	*Five worst hit*	*Three unaffected*
T * D2008	0.0173***	0.0089	-0.0187	0.0031
	(.0040)	(.0147)	(.0529)	(.0124)
T * D2009	0.0089*	0.0083	-0.0469	-0.002
	(.0043)	(.0135)	(.0499)	(.0209)
T * D2010	0.0037	-0.0003	-0.0623	-0.0062
	(.0082)	(.0189)	(.0669)	(.0215)
Observations	1 118 104	912 601	1 116 025	914 680

Note: Control variables: Age-indicators, education length- and type fixed effect, municipality fixed effect, year fixed effect, immigration category and household size. Log disposable income is included linearly except where it is dependent variable.
Cluster robust standard errors on treatment level (municipality)
* p<0.05, ** p<0.01, *** p<0.001

Decomposing the effect over age-groups

Having established that there is an effect in the five worst hit, I continue to test if groups that benefit more from public services respond more to the shock. Specifically, I check for differences based on observed age. Figure 2 shows how the effect in the five worst hit municipalities is distributed over age-groups. These results are obtained by splitting the sample in ten-year birth cohorts and performing the exact same difference in difference analysis as before. We see that the effect is driven mainly by young households. Interestingly, but less clear, retired households (the group between 70 and 80) respond as well. There might be other explanations for differences across age-groups (for example, credit constraints, level of insurance). However, we note that the largest

Figure 2: Consumption effect by age (2008)

The five worst hit The three unaffected

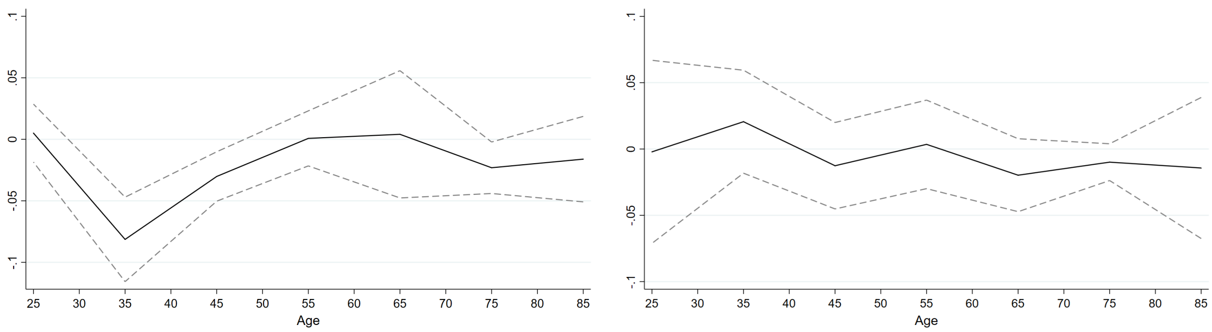

Note: Estimation based on splitting sample in 10-year cohorts. Control variables: Age-indicators, education length- and type fixed effect, municipality fixed effect, year fixed effect, immigration category and household size. Log disposable income entered as a control linearly. Dashed 95% confidence intervals based on cluster robust standard errors.

recipients of public goods and services are the ones who reduce their consumption. The effect is only present in the worst affected, while there is no effect for any age-group in the unaffected.

Risk-taking behavior of affected households

To investigate whether uncertainty is a probable driver of the consumption response, I do an additional test based on previous findings on risk-taking behavior. When households have preferences in accordance with prudence and temperance, they are sensitive to the overall risk exposure (Heaton and Lucas, 2000). Therefore, if households perceive the economic outlook of the municipality as more uncertain, they will reallocate their risky assets to safer ones. I follow Fagereng, Guiso, and Pistaferri (2017) and define the risky share of financial assets, RS, as the share of stocks and mutual funds to the overall sum of stocks, bonds, mutual funds, non listed stocks and bank deposits. Conditional on owning risky assets, I test whether there is a reallocation to holding less risky assets. I estimate the same specification as used in the main estimation for consumption, only changing the dependent variable to the risky share, RS, and conditioning on ownership:

$$RS = \alpha_0 + \delta_0 T_{it} + \sum_{j=1995}^{2010} \theta_j year_j + \delta_1 T_{it} * D_{2008} + \delta_2 T_{it} * D_{2009} + \delta_3 T_{it} * D_{2010} + X_{it}\beta + \gamma M_j + \varepsilon_{it} \quad (4)$$

Table 5 shows that the risky share decreases by 1.56 percent in the worst affected in 2008. This

Table 5: Effects on the share of risky assets held

	The five worst hit	The three unaffected
T*D2008	-.0156**	-.0009
	(.0054)	(.0117)
T*D2009	-.0044	-.0073
	(.0078)	(.0110)
T*D2010	-.0056	-.0019
	(.0073)	(.0114)
Observations	260 780	235 476

Note: Control variables: Age-indicators, education length- and type fixed effect, municipality fixed effect, year fixed effect, immigration category and household size.
Cluster robust standard errors on treatment level (municipality)
* p<0.05, ** p<0.01, *** p<0.001

is, again, in contrast to the zero effect found in the three unaffected. Following the same pattern as consumption, we see that there is no effect in 2009 and 2010. The result is consistent with a story of uncertainty, supporting the previous interpretation.[20] Households living in municipalities that struggled with financing their public services, and eventually had to be bailed out, did not only hold back on consumption by saving more, they also reallocated their assets to holding less risky ones. Once these municipalities received extraordinary transfers, there is no difference in risk-taking behavior.

One could argue that the reallocation to safer assets reflects that households in the affected municipalities are reminded of the risks in the stock market, and therefore reallocate. There is no way to test whether that is the case, however, if being reminded of stock market risks drive the results, we would expect there to be an effect in the three unaffected municipalities as well, since they were exposed to the news, and loss of public assets, in a similar way.

Unemployment, income expectations and moving

In this subsection, I address other effects in the municipalities, that potentially could affect or drive the observed results. Specifically, I show that there are no effects on unemployment and moving in the affected municipalities.

The issue of changes in disposable income was addressed in the main results section, and

[20]This is an interesting finding on its own, that will be investigated further in future versions.

Figure 3: Municipality unemployment rates before and after the event

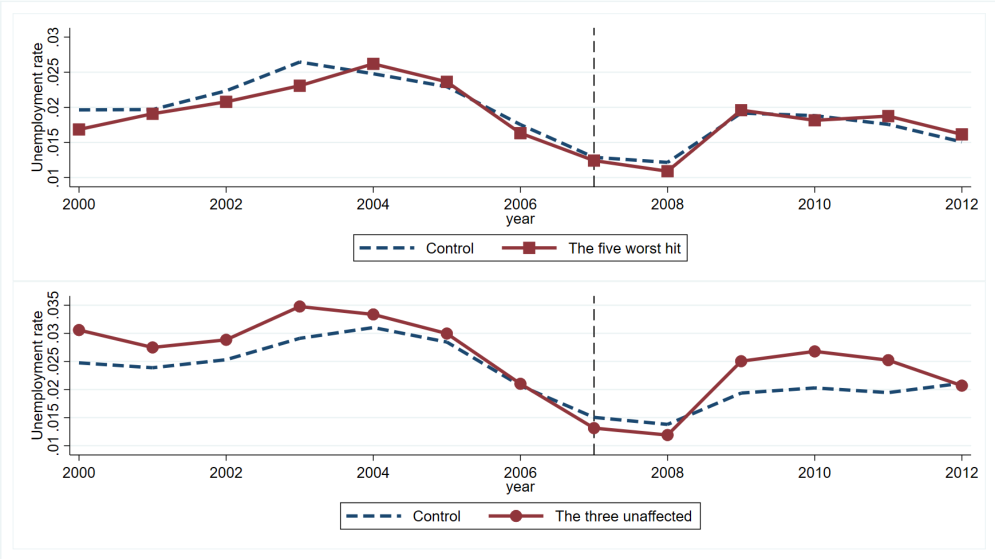

Note: Average yearly unemployment rate data from Statistics Norway. Rate averaged across municipalities in the treated and control group respectively.

is addressed in particular in the robustness section, but I return to the question with a different perspective here. Even though the decrease in consumption is not driven by decreases in the disposable income of households *within* my estimation sample, changes in *overall* employment could influence expectations and uncertainty regarding future income. Furthermore, Basten, Fagereng, and Telle (2016) show that households save more, anticipating future unemployment. Therefore, I check whether there was any changes to the labor market after the events. I collect yearly unemployment data on municipality level. Estimating the difference in difference on these data, I find no significant changes in unemployment after the shock. Figure 3 plots the unemployment rate in the treated and control group for the five worst hit and the three unaffected. We see that there is no difference in the five worst hit, while there is a small increase in the three unaffected. Note that the unemployment rate lies above in the pre-treatment period in the three unaffected. Therefore, the spike in 2009 might represent natural variation. To the extent unemployment risk could be a concern, or driver, there are no indications in the data of effects on labor market conditions in the five worst hit, represented through unemployment.

Another threat to identification I have not addressed, is moving. Given the need for strict sample selection to impute consumption (including the removal of observations where people

Figure 4: Percentage change in population before and after the event

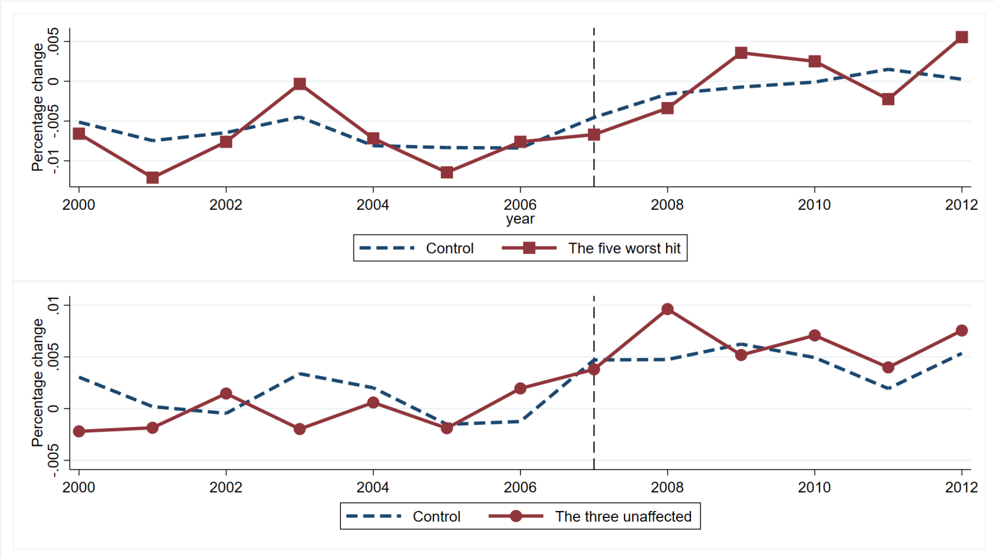

Note: Average yearly growth rate of population. Data from Statistics Norway. Rate averaged across municipalities in the treated and control group respectively, net of municipality fixed effects.

move), it could be the case that the effect I find is simply driven by a selection of stayers. To address this, I use publicly available data containing quarterly population data at the municipality level. Since there are large differences in population between municipalities, I use relative changes and estimate the difference in difference. I find no significant effect on changes in population, meaning that the event did not lead to people moving away.[21] Figure 4 plots the relative changes in population in the treated and control groups, net of municipality fixed effects. These indicate that changes to the population do not affect the results. This is also supported by the sample employed from the register data, where the number of observations per year is unaffected by the event.

A quick summary the main results: I have established that there is an effect on consumption in the five worst hit municipalities. The effect is driven primarily by the young, however, some households older than the retirement age respond as well. The effect is not driven by income changes or expectation of future unemployment, rather, households reduce the amount they consume out of their disposable income. This is done either through repaying existing debt, or taking on less debt compared to inhabitants in the control group.

[21] There are no indications of this mentioned in the annual reports of the municipalities.

Local governments and the relationship between fiscal policy and private consumption

Having seen a clear reduction in private consumption among the five worst hit municipalities, I have argued that fiscal uncertainty is the driver of the results. In that context, it is important to understand the actions of the local governments, and whether there are other effects than uncertainty about the economic changes that can rationalize household's choice of consumption and saving. For example, could the observed results be explained by private consumption and public spending being complements? Did private consumption go down because it became more expensive to live in the municipality after the events?

I formalize these arguments in the following section, by describing the local government budget constraint, and how this can affect private consumption in the framework of the Permanent Income Hypothesis. The point of this simple theoretical exercise, is to show how the actions of local governments matter, before I provide some illustrative empirical results of their actions, based on municipality level income statements.[22]

Local government budget constraint

First, consider the possible set of actions for the local government after the crisis. Local governments faces the following budget constraint in Norway when they make their economic decisions (Aaberge and Langørgen, 2003),

$$a_t + v_t = z_t + \sum_{s=1}^{S} p_{st} q_{st}, \tag{5}$$

where a_t are general grants-in-aid by the central government, v_t is local taxes and user fees, z_t is budget surplus or deficit and p_{st} and q_{st} are price and quantity in service sector s. The larger part of local government income comes from the central government through a_t and should be considered fixed. The local government has some freedom in setting the local property tax rates and user fees v_t.[23] The services provided, q_{st}, have a minimum level required by law (for example, health care and schooling is free and needs to meet certain requirements). Prices (i.e, cost) of providing

[22] This exercise bares some resemblance to Section 4 in Turnbull (1998), who focuses on flypaper effects and uncertainty of public services.

[23] For example, the maximum property tax rate is 6 %, and it is up to the local government to set the rate, if introduced at all. User fees of kindergartens is set locally, infrastructure services charged at cost.

services, p_{st} vary due to geographical and demographic differences and should be regarded fixed in this context. As long as total income exceeds the costs of minimum services, the local government is essentially free to provide whichever level or type of service they like. After the event, affected municipalities had to cover the deficit within four years. This could be done by increasing income on the left hand side, or by reducing the services provided to their inhabitants on the right hand side.

Permanent Income Hypothesis and consumption responses

Consider a household that lives in T periods and receives utility from consumption of publicly provided services g_t and private consumption goods c_t, discounted by β. It has a fixed income y_t, pays taxes τ_t to the government, and can save or borrow money freely at the risk free rate r. It can not leave debt or wealth behind ($b_{T+1} = 0$). In this standard permanent income framework, it is well known that optimal consumption follows the Euler equation,

$$u'(c_t, g_t) = \beta(1 + r)u'(c_{t+1}, g_{t+1}) \tag{6}$$

The way households react to changes in g_{t+1}, will depend on how their utility is defined. Assume that households have a utility function displaying constant elasticity of substitution between private consumption and public goods, such as $u(c, g) = (\alpha c^\rho + (1 - \alpha)g^\rho)^{1/\rho}$, where $\frac{1}{1-\rho}$ is the elasticity of substitution. If $\rho = 1$, these are perfect substitutes. If $\rho < 0$, there is complementarity between the two.

With income, public taxes and services are deterministic, $\rho = 1$, and $\beta(1 + r) = 1$, households achieve perfect smoothing:

$$c_t + g_t = c_{t+1} + g_{t+1}$$

This means that if households know at time t that the level of g_{t+1} drops, they can compensate by adjusting c_t and c_{t+1}. If households know at time t-1 that g_t will drop before g_{t+1} goes up again, they can smooth consumption by shifting consumption to period t. However, if $\rho < 0$, and we assume for simplicity that there are equal shares of public and private goods, the Euler equation becomes:

$$(\frac{c_t}{c_{t+1}})^{\rho-1} = \Big(\frac{c_{t+1}^{\rho} + g_{t+1}^{\rho}}{c_t^{\rho} + g_t^{\rho}}\Big)^{\frac{1-\rho}{\rho}}$$

Since marginal utility of private consumption now depends on the level of public services, it can be optimal to follow the pattern of public spending. Thus, if households expect future services to go down permanently, they will maximize utility by shifting consumption towards the present, when public spending is still high. If households expect a transitory dip in public services, they will hold back consumption in those periods, before increasing it when public services goes back up.

In the above, we have interpreted consumption movements through its relationship to public services, treated as an exogenous variable in their optimisation. However, we must also consider that local governments could hold the service level fixed, and finance the deficit through increases in local taxes and user fees.[24] To the extent they do so, this will directly affect the disposable income of inhabitants, and therefore their choice of consumption. In this scenario, inhabitants must expect lower disposable income in the future, which according to the permanent income hypothesis would lead to lower consumption levels already today. Remember from Section 3 that user fees and local taxes are charged directly to the inhabitants by the municipality, and therefore it is not included in the measure of disposable income from the state tax registers. Consequently, any changes to user fees or local taxes will show up in the imputed consumption measure as increases in private consumption, ceteris paribus. This means that if local governments were able to let households pay the bill, we would not see a drop in imputed consumption. This would, however, be traceable through the municipality income statements.

Finaly, to add uncertainty to the framework, it is common to assume one or more of the variables of the optimisation problem follows a stochastic process. Here, I simply add a random variable, $\varepsilon_{t+1} \sim \mathcal{N}(0, \sigma)$, to show the effect on consumption when the variance of this term increases, i.e., when the future becomes more uncertain. When the next period's marginal utility is uncertain, we get the expected Euler equation,

$$u'(c_t, g_t) = \beta(1 + r)E_t[u'(c_{t+1}, g_{t+1}, \varepsilon_{t+1})]. \tag{7}$$

[24] This is analogue to the classic tax policy experiment leading to Ricardian equivalence, except that here there is no positive income shock in the first period.

If households are prudent, meaning they have convex marginal utility, Jensen's inequality shows that if uncertainty increases (through increasing σ), the expected marginal utility of tomorrow's consumption increases. If households in the affected municipalities perceive that future consumption is more uncertain, they put less weight on its value. Therefore, they will save more and consume less in the current period. Future consumption could be regarded more uncertain if it is unknown what level of services g_{t+1} will be provided, or it could be affected indirectly if income and employment are regarded more uncertain, as argued by Bratberg and Monstad (2015).

Results from municipality income statements and further discussion

Since the financial choices of local governments can be of importance for private consumption, I use data from municipality annual reports between 2003 and 2012, and estimate changes to public expenditures per capita in the affected municipalities. Specifically, I use the same treated and control groups as before, and estimate yearly interaction effects the same way as was done for Figure 1. The specification includes municipality and year fixed effects. I also control for population. Figure 5 plots the interaction terms of the five worst hit and the three unaffected.[25]

Figure 5: Interaction effects on per capita public expenditures, normalized to zero in 2007

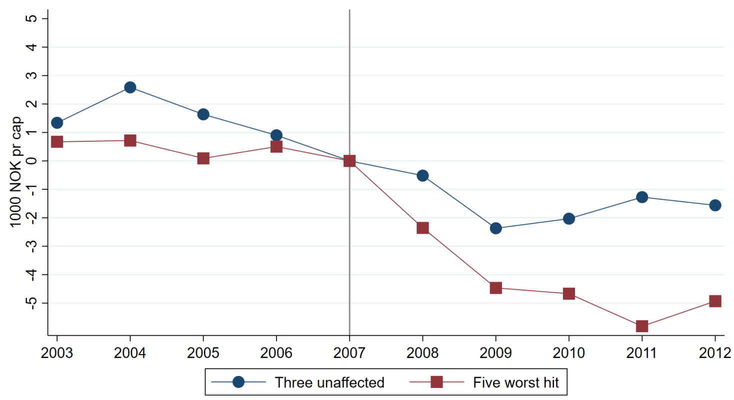

Note: Estimation of interaction effects done separately for the two plots. Controls include municipality fixed effects, year fixed effects and population size.

The plot for the five worst hit shows that there was a parallel trend prior to 2007. Public

[25]Note that only the individual interaction coefficient for the five worst hit in 2011 is significantly different from zero on a five percent level. Given the number of effects estimated relative to observations, it is no surprise that this estimation has low power.

expenditure per capita went down in 2008, and continued to do so for the remaining period of analysis. There was, in other words, a permanent drop in expenditures.

The plot for the three unaffected has a downward sloping trend, compared to their comparison group, from 2004 to 2008. This continued until 2009, when there was a small distinct drop, before it caught up marginally towards the end of the period. While the observed drop from 2007 to 2010 do not chime well with a narrative claiming these municipalities were unaffected, the measures taken looks less severe in these municipalities.

Table 7 gives a full overview of the results from estimating the average treatment effect in 2009-2012 on all elements of the income statement normalized by population. The estimation includes time- and municipality fixed effects. The results show that the most effective measures taken are on the cost side. Specifically, purchases of goods, social expenditures, wage expenditures and transfers were significantly reduced in the five worst hit. In the three unaffected, the reduction was mainly related to a reduction in wage expenditures. The local governments' ability to cover losses through increases in income is limited in this context. Although there is a point estimate increase in user fees, real estate taxes and other direct/indirect taxes, there is also a reduction in transfers offsetting those. None of these are statistically significant. In the three unaffected, there is a decrease in income similar to their decrease in costs.[26]

Given that the drop in public expenditures is permanent, and the effect seen on private consumption is temporary, it seems unlikely that changes in public services themselves are the sole reason for the drop in private consumption in 2008. While there are valid reasons why consumption should fall permanently, it seems unlikely that the initial drop can be fully explained by changes to public services. Interesting to note in this case, is the small decrease in the five worst hit in 2010, seen in Figure 1. Future research should include data for more years after the event to study longer term effects on consumption. An interesting extension would be to test if there is a persistent component to the consumption drop that can tell us something about the relationship between public and private spending as well.

Although uncertainty is difficult to quantify in this context,[27] several elements in the time line

[26] This might reflect that Haugesund tried to downsize their operational level in the period.

[27] The literature on aggregate uncertainty often uses the spread of forecasters or stock market returns to create measures of uncertainty. Baker, Bloom, and Davis (2015) provides such an index. On a local level, quantifying uncertainty is an even more challenging task.

point towards uncertainty as a plausible explanation for the sudden drop of private consumption in 2008. As the previous section showed, households in the worst hit municipalities had real reasons to be concerned about the fiscal future of the municipality, however, this was a concern that where greatly reduced when the central government provided additional funding for essential services. Thus, when the central government opened up the bail out, households understand that there are limits to how bad the cuts will be. Having experienced this reduction in uncertainty, households go back to normal behavior, despite the fact that local governments makes further cuts to spending. I therefore regard uncertainty as the most likely driver.

6 Summary and conclusion

I have used a difference in difference approach to evaluate how private consumption of households changed when their local government experienced large losses of public assets. Households in the five worst hit municipalities reduced their consumption by 1.8 percent the first year after the event. In the three municipalities that did not struggle economically afterwards, no effect is found on private consumption. The effect is temporary, meaning I only find an effect the first year after the event.

The fall in consumption is driven by age-groups who are the largest recipients of public goods - the young and the elderly. In other words, groups that where most likely to be affected from changes in public services, are the ones who act more careful. I show that the change in consumption comes from increased saving, meaning households spend less of their income than their comparison group. To investigate whether uncertainty is the driver, I follow the literature on portfolio choice and background risk, and study the share of risky assets held by the affected households. I find that households holding risky assets rebalance their portfolios after the event, holding a lower share of risky assets one year after.

Since the loss of publicly owned assets may carry other effects than uncertainty, I consider alternative explanations for the observed results. I provide a short theoretical framework to illustrate how my results can be interpreted in the light of local government spending decisions. Using municipality income statements, I show that public spending falls permanently over the period of analysis. If the changes in private spending were driven by actual changes in public spending, the decrease in private consumption should have been permanent as well. I therefore argue that the actual fiscal actions of the local governments are not the driver of the observed consumption drop. Considering other alternative explanations, I show that there is no effect on moving, employment or disposable income.

I therefore conclude that the effect is driven by households holding back consumption until uncertainty regarding fiscal outcomes is resolved.

References

Aaberge, R., and A. Langørgen (2003): "Fiscal and spending behavior of local governments: Identification of price effects when prices are not observed," *Public Choice*, 117(1-2), 125–161.

Aaberge, R., and J. Langørgen (2011): "Gruppering av kommuner etter folkemengde og konomiske rammebetingelser 2008 - [Grouping of municipalities by population and economic conditions 2008]," *Statistics Norway, Reports 08*.

Aaberge, R., K. Liu, and Y. Zhu (2016): "Political uncertainty and household savings," *Journal of Comparative Economics*.

Baker, S. R., N. Bloom, and S. J. Davis (2015): "Measuring economic policy uncertainty," Discussion paper, National Bureau of Economic Research.

Basten, C., A. Fagereng, and K. Telle (2016): "Saving and portfolio allocation before and after job loss," *Journal of Money, Credit and Banking*, 48(2-3), 293–324.

Basu, S., and B. Bundick (2017): "Uncertainty shocks in a model of effective demand," *Econometrica*, 85(3), 937–958.

Blanchard, O., and R. Perotti (2002): "An empirical characterization of the dynamic effects of changes in government spending and taxes on output," *The Quarterly Journal of Economics*, 117(4), 1329–1368.

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77(3), 623–685.

——— (2014): "Fluctuations in uncertainty," *Journal of Economic Perspectives*, 28(2), 153–176.

Bratberg, E., and K. Monstad (2015): "Worried sick? Worker responses to a financial shock," *Labour Economics*, 33, 111–120.

Browning, M., and S. Leth-Petersen (2003): "Imputing consumption from income and wealth information," *The Economic Journal*, 113(488).

Caldara, D., C. Fuentes-Albero, S. Gilchrist, and E. Zakrajšek (2016): "The macroeconomic impact of financial and uncertainty shocks," *European Economic Review*, 88, 185–207.

Cameron, A. C., and D. L. Miller (2015): "A Practitioners Guide to Cluster-Robust Inference," *Journal of Human Resources*, 50(2), 317–372.

Carroll, C. D., and M. S. Kimball (2006): *Precautionary Saving and Precautionary Wealth*chap. The New Palgrave Dictionary of Economics.

Eika, L., M. Mogstad, and O. Vestad (2017): "What can we learn about household consumption from information on income and wealth," Discussion paper, Mimeo.

Fagereng, A., L. Guiso, and L. Pistaferri (2017): "Portfolio Choices, Firm Shocks, and Uninsurable Wage Risk," *The Review of Economic Studies*, 85(1), 437–474.

Fagereng, A., and E. Halvorsen (2017): "Imputing consumption from Norwegian income and wealth registry data," *Journal of Economic and Social Measurement*, 42(1), 67–100.

Fagereng, A., M. B. Holm, and G. J. Natvik (2016): "MPC heterogeneity and household balance sheets," *Disucssion Paper, Statistics Norway*.

Fatás, A., I. Mihov, et al. (2001): *The effects of fiscal policy on consumption and employment: theory and evidence*, vol. 2760. Citeseer.

Fuchs-Schündeln, N., and T. A. Hassan (2016): "Natural experiments in macroeconomics," *Handbook of Macroeconomics*, 2, 923–1012.

Galí, J., J. D. López-Salido, and J. Vallés (2007): "Understanding the effects of government spending on consumption," *Journal of the European Economic Association*, 5(1), 227–270.

Giavazzi, F., and M. McMahon (2012): "Policy uncertainty and household savings," *Review of Economics and Statistics*, 94(2), 517–531.

Giavazzi, F., and M. Pagano (1990): "Can severe fiscal contractions be expansionary? Tales of two small European countries," *NBER macroeconomics annual*, 5, 75–111.

Heaton, J., and D. Lucas (2000): "Portfolio choice in the presence of background risk," *The Economic Journal*, 110(460), 1–26.

Hofstad, M. (2008): *Eventyret om TERRA og de åtte små kommunene som dro til Wall Street for å bli rike - ["The Adventure of TERRA and the eight small Municipalities that went to Wall Street to get rich"]*. Hegnar Media.

Koijen, R., S. Van Nieuwerburgh, and R. Vestman (2014): "Judging the Quality of Survey Data by Comparison with" Truth" as Measured by Administrative Records: Evidence From Sweden," in *Improving the Measurement of Consumer Expenditures*, pp. 308–346. University of Chicago Press.

Kostøl, A. R., and M. Mogstad (2015): "Earnings, Disposable Income, and Consumption of Allowed and Rejected Disability Insurance Applicants," *The American Economic Review*, 105(5), 137–141.

Mountford, A., and H. Uhlig (2009): "What are the effects of fiscal policy shocks?," *Journal of Applied Econometrics*, 24(6), 960–992.

Palia, D., Y. Qi, and Y. Wu (2014): "Heterogeneous Background Risks and Portfolio Choice: Evidence from Micro-level Data," *Journal of Money, Credit and Banking*, 46(8), 1687–1720.

Perotti, R. (2007): "In Search of the Transmission Mechanism of Fiscal Policy," *NBER macroeconomics Annual*, 22, 169–249.

Ramey, V. A., and M. D. Shapiro (1998): "Costly capital reallocation and the effects of government spending," in *Carnegie-Rochester Conference Series on Public Policy*, vol. 48, pp. 145–194. Elsevier.

Turnbull, G. K. (1998): "The overspending and flypaper effects of fiscal illusion: Theory and empirical evidence," *Journal of Urban Economics*, 44(1), 1–26.

7 Appendix

Table 6: Effect on consumption to income ratios

	Consumption to disposable income		Consumption to lagged income	
	(1)	(2)	(3)	(4)
	Five worst hit	*Three unaffected*	*Five worst hit*	*Three unaffected*
T * D2008	-.0257***	-.0100	-.0283*	-.0043
	(.0051)	(.0120)	(.0136)	(.0118)
T * D2009	-.0024	-.0044	.0077	-.0341
	(.0061)	(.0191)	(.0148)	(.0165)
T * D2010	-.0024	-.0014	-.0069	-.0016
	(.0133)	(.0197)	(.00623)	(.0231)
Observations	1 116 025	914 680	717 649	568 059

Note: Control variables: Age-indicators, education length- and type fixed effect, municipality fixed effect, year fixed effect, immigration category and household size. Log disposable income included as a linear control.
Cluster robust standard errors on treatment level (municipality)
* p<0.05, ** p<0.01, *** p<0.001

Table 7: Treatment effects (2009-2012) from municipal annual reports (1000 NOK pr capita)

	The five worst hit		The three unaffected	
	ω_2	St.error	ω_2	St.error
User fees	0.431	(0.301)	-0.0577	(0.328)
Sale and rental	-0.434	(0.505)	-1.507*	(0.726)
Transfer with claims	-1.895***	(0.649)	-0.916	(0.839)
Grants	0.323	(1.052)	-0.931	(1.530)
Other state transfers	1.010	(1.815)	0.238	(0.309)
Other transfers	-1.810***	(0.530)	-0.457	(0.310)
Income taxes	-0.837	(0.865)	0.141	(0.513)
Real estate taxes	0.529	(0.549)	0.0281	(0.824)
Other direct/indirect taxes	1.226	(0.847)	-0.0602	(0.265)
Sum operating income	-1.456	(2.418)	-3.522	(2.858)
Wage expenditures	-1.718*	(0.950)	-2.646**	(0.934)
Social expenditures	-0.370***	(0.109)	-0.100	(0.461)
Purchase of goods	-2.664***	(0.424)	-0.114	(1.010)
Purchases of services	-0.163	(0.648)	0.290	(0.504)
Transfers	-1.277***	(0.299)	-0.351	(0.477)
Depreciation	0.0706	(0.176)	0.123	(0.256)
Distributed costs	1.025**	(0.404)	-0.810	(1.007)
Sum operating expenses	-5.095***	(1.719)	-3.608	(2.423)
Gross operating profit	3.405**	(1.479)	0.653	(0.741)
External financial income	1.535	(1.317)	0.453	(0.415)
External financial expenditures	-4.988**	(1.917)	-0.614	(0.711)
External financial transactions	6.523**	(3.160)	1.057*	(0.496)
Net operating profit	9.999**	(4.384)	1.834*	(0.860)
Observations		783		126

Note: All specifications include year and municipality fixed effects in addition to controlling for population.
Cluster robust standard errors on treatment level
* p<0.05, ** p<0.01, *** p<0.001
All variables in 1000NOK

7.1 Placebo analysis and robustness checks

I perform placebo and robustness checks. First, the placebo exercise is displayed in Table 8. Estimating a one-year effect using 1999-2008 as on year treatment periods, shows that measuring an effect at 0.1 percent significance level on log consumption is rare in these samples. I find no significant effect at 0.1 percent level in the three unaffected in any of the years. I also perform the same placebo exercise for consumption to income ratio and log disposable income. The specification is similar to the one used for the main results, except that I estimate 1-year effects excluding the future sample. This means that for the first column of Table 8, the treatment year is $T = 1999$, and the sample employed is 1995-1999.

$$x_{it} = \alpha_0 + \sum_{j=1995}^{T} \theta_j year_j + \delta_0 T_{it} + \delta_1 T_{it} * D_T + X_{it}\beta + \gamma M_j + \varepsilon_{it} \tag{8}$$

Second I do robustness checks. These are displayed in Table 9 and show estimation of the main specification varying controls (columns 1-3). The qualitative pattern is the same across all the specifications: No effect in the three unaffected, and a significant effect in the five worst hit. The point estimate is sensitive to the controls included, which is to be expected given that there are time-varying differences between the municipalities in the period that affects consumption. Specification (4) includes a separate linear trend between the treated and control group. None of the interaction terms for different trends are statistically significant, and the point estimates are close to zero for all the groups. Although the treatment effect becomes less precisely estimated, it is still significant at the 5 percent level. The point estimate of consumption reduction is now 3.18 percent for the five worst hit, which is larger than in the main specification without linear trend terms. There is still no effect in the three municipalities with no change in their public spending, consistent with my main results. Specification (5) shows that inclusion of treatment specific income controls reduces the point estimate.

Table 8: Placebo exercise - Only coefficient of treatment effect (δ_2) reported

	Post treatment period, T =									
	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008
All Terra Municipalities										
Log Consumption	0.0158*	0.00903	0.00805	0.0154***	-0.00282	-0.00000657	-0.00242	0.00202	-0.000411	-0.0115*
	(0.00603)	(0.00560)	(0.00537)	(0.00454)	(0.0113)	(0.00867)	(0.00923)	(0.00769)	(0.00831)	(0.00567)
Consumption to income ratio	0.0158**	0.0146**	0.0149*	0.0148*	-0.00537	0.00222	-0.00713	0.00133	-0.00309	-0.0163*
	(0.00521)	(0.00523)	(0.00693)	(0.00739)	(0.0118)	(0.00885)	(0.00897)	(0.00954)	(0.0100)	(0.00743)
Log disposable income	-0.00289	-0.00919	-0.00884	-0.0108	-0.00855	0.00280	-0.00491	-0.00873	0.00503	-0.000404
	(0.0105)	(0.00634)	(0.00452)	(0.00630)	(0.00545)	(0.00612)	(0.00464)	(0.00590)	(0.00968)	(0.00862)
The five worst hit - T5										
Log Consumption	0.0128	-0.00480	0.00677	0.0125*	0.00685	0.0174	0.0165	0.00478	0.00727	-0.0177***
	(0.00774)	(0.00640)	(0.00463)	(0.00501)	(0.00765)	(0.0172)	(0.0169)	(0.0155)	(0.0193)	(0.00448)
Consumption to income ratio	0.0129	0.00157	0.00929	0.0104	-0.00378	0.0211	0.00352	0.00534	0.0131	-0.0255***
	(0.00903)	(0.00609)	(0.00521)	(0.00990)	(0.00780)	(0.0172)	(0.0161)	(0.0251)	(0.0177)	(0.00513)
Log disposable income	-0.0106*	-0.0138	-0.0121	-0.0108	-0.00980*	-0.00120	-0.0115*	-0.0136*	-0.0107	-0.0201
	(0.00473)	(0.00937)	(0.00675)	(0.00631)	(0.00442)	(0.00362)	(0.00516)	(0.00570)	(0.00774)	(0.0166)
The three unaffected - T3										
Log Consumption	0.0123	0.0201**	0.00596	0.0119	-0.00934	-0.0163*	-0.0197*	-0.00187	-0.0135	-0.00429
	(0.00854)	(0.00499)	(0.00934)	(0.00607)	(0.0159)	(0.00618)	(0.00764)	(0.00939)	(0.0102)	(0.00893)
Consumption to income ratio	0.0112*	0.0239***	0.0152	0.00903	-0.00969	-0.0174**	-0.0214	-0.00538	-0.0228	-0.00999
	(0.00503)	(0.00474)	(0.0131)	(0.00981)	(0.0178)	(0.00573)	(0.0108)	(0.00832)	(0.0123)	(0.0118)
Log disposable income	0.00765	-0.00163	-0.0000762	-0.00420	-0.00230	0.00925	0.00211	-0.00211	0.0200	0.0223*
	(0.0161)	(0.00876)	(0.00571)	(0.00872)	(0.00748)	(0.00935)	(0.00648)	(0.00828)	(0.0130)	(0.00787)

Note: Control variables: Age-indicators, education length- and type fixed effect, municipality fixed effect, year fixed effect, immigration category and household size. Log disposable income enters linearly as control except when being the dependent variable.
Cluster robust standard errors on treatment level (municipality)
* p<0.05, ** p<0.01, *** p<0.001

Table 9: Main results with varying set of control variables

	The five worst hit				
	(1)	(2)	(3)	(4)	(5)
T*D2008	-.0458**	-.0342*	-.0179***	-.0318*	-.013*
	(.0148)	(.0146)	(.0044)	(.0129)	(.0051)
T*D2009	-.0263**	-.0125	-.0031	-.0189	.0026
	(.0096)	(.0088)	(.0061)	(.0122)	(.0058)
T*D2010	-.0215	-.0083	-.0090	-.0267	-.0028
	(.0161)	(.0145)	(.0093)	(.0245)	(.0105)
Municipality fixed effect	Yes	Yes	Yes	Yes	Yes
Year fixed effect	Yes	Yes	Yes	Yes	Yes
Age fixed effect	No	Yes	Yes	Yes	Yes
Other characteristics	No	No	Yes	Yes	Yes
Separate linear trend	No	No	No	Yes	No
Separate income control	No	No	No	No	Yes
Observations	1 138 613	1 138 613	1 138 613	1 138 613	1 138 613

	The three unaffected				
	(1)	(2)	(3)	(4)	(5)
T*D2008	.0133	.0083	-.004	-.0007	-.0027
	(.0128)	(.0157)	(.009)	(.0115)	(.0094)
T*D2009	.0106	.0062	-.0059	-.0018	-.0040
	(.0236)	(.0296)	(.0211)	(.0198)	(.0200)
T*D2010	.0227	.0163	.0029	.0075	.0049
	(.0195)	(.0267)	(.0154)	(.0159)	(.0144)
Municipality fixed effect	Yes	Yes	Yes	Yes	Yes
Year fixed effect	Yes	Yes	Yes	Yes	Yes
Age fixed effect	No	Yes	Yes	Yes	Yes
Other characteristics	No	No	Yes	Yes	Yes
Separate linear trend	No	No	No	Yes	No
Separate income control	No	No	No	No	Yes
Observations	933 520	933 520	933 520	933 520	933 520

Note: Other characteristics : Indicators for Education length- and type, immigration category and household size. Log disposable income enters linearly.
Cluster robust standard errors on treatment level (municipality)
* p<0.05, ** p<0.01, *** p<0.001

Table 10: Full results from main specification

	The five worst hit		*The three unaffected*	
Variable	Coefficient	Standard error	Coefficient	Standard error
T*Year=2008	-.0179384***	(.0044273)	-0.00432	(0.00897)
T*Year=2009	-.0030807	(.0061107)	-0.00595	(0.0211)
T*Year=2010	-.0090366	(.0092757)	0.00286	(0.0154)
T=1	-.0358463***	(.0007517)	0.00412	(0.00245)
Year=1994	-.120284***	(.0054784)	-0.146***	(0.00852)
Year=1995	-.1170369***	(.0070317)	-0.152***	(0.00763)
Year=1996	-.1131295***	(.004476)	-0.133***	(0.00670)
Year=1997	-.1103248***	(.0056616)	-0.134***	(0.0102)
Year=1998	-.0857439***	(.0055135)	-0.107***	(0.0103)
Year=1999	-.1160272***	(.0054098)	-0.127***	(0.0105)
Year=2000	-.0799354***	(.0062698)	-0.107***	(0.0113)
Year=2001	-.078754***	(.0055581)	-0.0945***	(0.0126)
Year=2002	-.0672315***	(.0039775)	-0.0756***	(0.00824)
Year=2003	-.0590211***	(.0060827)	-0.0800***	(0.00769)
Year=2004	-.0461364***	(.0039343)	-0.0551***	(0.00854)
Year=2005	-.0728492***	(.0041392)	-0.0808***	(0.00930)
Year=2006	-.026162***	(.0037554)	-0.0401***	(0.00691)
Year=2007	0	.	0	.
Year=2008	.0046423	(.0040291)	-0.0248*	(0.00976)
Year=2009	-.0420549***	(0057247)	-0.0679***	(0.0104)
Year=2010	-.0138301**	(.0045893)	-0.0378*	(0.0130)
Log disposable income	.7452732***	(.0044803)	0.759***	(0.00532)
Age=25	0		0	
Age=26	-.0174847***	(.0041238)	-0.0160***	(0.00349)
Age=27	-.0248699***	(.004767)	-0.0203**	(0.00634)
Age=28	-.024698***	(.0047917)	-0.0306***	(0.00368)
Age=29	-.0271885***	(.0060674)	-0.0289***	(0.00581)
Age=30	-.0367847***	(.005387)	-0.0372**	(0.00949)
Age=31	-.0407553***	(.0052691)	-0.0252*	(0.00895)
Age=32	-.0330198***	(.0061442)	-0.0293***	(0.00376)
Age=33	-.0435534***	(.0057564)	-0.0422***	(0.00538)
Age=34	-.041346***	(.0057851)	-0.0394***	(0.00629)
Age=35	-.0426793***	(.0054616)	-0.0496***	(0.00808)
Age=36	-.0506713***	(.0052755)	-0.0451***	(0.00919)
Age=37	-.04455***	(.0053432)	-0.0398***	(0.00720)
Age=38	-.0500981***	(.0059894)	-0.0525***	(0.00853)
Age=39	-.0527731***	(.0048551)	-0.0503***	(0.00723)
Age=40	-.043085***	(.0060621)	-0.0542***	(0.00885)
Age=41	-.0531906***	(.0052523)	-0.0506***	(0.0107)
Age=42	-.0566646***	(.0052159)	-0.0537***	(0.00718)
Age=43	-.0505491***	(.0058246)	-0.0421**	(0.0101)
Age=44	-.047203***	(.0055876)	-0.0522***	(0.00935)
Age=45	-.0439403***	(.005417)	-0.0375***	(0.00841)
Age=46	-.0405389***	(.0055619)	-0.0387***	(0.00637)
Age=47	-.0356618***	(.0049859)	-0.0414***	(0.00900)
Age=48	-.0355295***	(.0061278)	-0.0406***	(0.00633)
Age=49	-.0420633***	(.0066317)	-0.0390***	(0.00821)
Age=50	-.0368304***	(.0061596)	-0.0256**	(0.00653)
Age=51	-.0304367***	(.0054204)	-0.0261*	(0.00941)
Age=52	-.0331119***	(.005182)	-0.0268**	(0.00722)
Age=53	-.0380445***	(.0054033)	-0.0301***	(0.00498)
Age=54	-.0286734***	(.0047996)	-0.0237*	(0.00792)
Age=55	-.0418***	(.005454)	-0.0343***	(0.00778)
Age=56	-.0287622***	(.0057046)	-0.0252**	(0.00687)
Age=57	-.031259***	(.0057941)	-0.0313**	(0.00797)
Age=58	-.0260882***	(.0075324)	-0.0297***	(0.00588)
Age=59	-.0317721***	(.0066997)	-0.0366***	(0.00694)
Age=60	-.040365***	(.0053694)	-0.0180*	(0.00734)
Age=61	-.0270308***	(.0050714)	-0.0161**	(0.00443)
Age=62	-.0328059***	(.0061594)	-0.0262**	(0.00752)
Age=63	-.0151357*	(.0066247)	-0.00341	(0.00510)
Age=64	-.0177499**	(.0053958)	0.0100	(0.00618)
Age=65	-.0274711***	(.0060446)	-0.0112	(0.0110)
Age=66	.0004119	(.0059494)	0.00159	(0.00931)
Age=67	-.0344252***	(.0069343)	-0.00727	(0.00511)
Age=68	-.0305188***	(.0049218)	-0.0107	(0.00803)
Age=69	-.0312929***	(.0049054)	-0.00998	(0.00573)
Age=70	-.0237178***	(.0056804)	-0.00782	(0.00402)
Age=71	-.0326435***	(.0058932)	-0.0103	(0.00799)
Age=72	-.0292321***	(.0059577)	-0.0196**	(0.00477)
Age=73	-.0315004***	(.0049932)	-0.0235***	(0.00370)
Age=74	-.0414585***	(.0057702)	-0.0217***	(0.00398)
Age=75	-.0362503***	(.0052815)	-0.0222***	(0.00387)
Age=76	-.0466856***	(.0054033)	-0.0324***	(0.00537)

Age=77	-.0496838***	(.0053669)	-0.0334***	(0.00696)
Age=78	-.0498062***	(.0048859)	-0.0419***	(0.00524)
Age=79	-.0568583***	(.0056252)	-0.0511***	(0.00466)
Age=80	-.0527337***	(.0051894)	-0.0388***	(0.00617)
Age=81	-.0623773***	(.0055442)	-0.0591***	(0.00632)
Age=82	-.0704677***	(.0049254)	-0.0621***	(0.00564)
Age=83	-.0711324***	(.0050979)	-0.0599***	(0.00730)
Age=84	-.0798396***	(.0046442)	-0.0629***	(0.00782)
Age=85	-.0874148***	(.0058625)	-0.0672***	(0.00703)
Age=86	-.0876697***	(.0054509)	-0.0780***	(0.00786)
Age=87	-.0891694***	(.0063961)	-0.0887***	(0.00779)
Age=88	-.1026602***	(.0051798)	-0.0796***	(0.0109)
Age=89	-.0985077***	(.0057729)	-0.101***	(0.00811)
Age=90	-.0976277***	(.0061135)	-0.0871***	(0.0107)
Non-/pre school	0		0	
Elementary	.0144474	(.0182318)	-0.0117	(0.0209)
Secondary	.0063745	(.016478)	0.0105	(0.0198)
High School Basic	.0515164**	(.0169074)	0.0542*	(0.0208)
High school	.0560955**	(.0175359)	0.0626*	(0.0218)
High school + add on	.0762848***	(.0162819)	0.0682**	(0.0222)
University/college lower	.0793507***	(.0177999)	0.0765**	(0.0218)
University	.1076615***	(.019468)	0.0870**	(0.0219)
Research	.1132591***	(.0261574)	0.0994**	(0.0276)
Not given	.017758	(.0108874)	0.0206	(0.0128)
General	0		0	
Humanities Art	-.0311863***	(.0051734)	-0.0220**	(0.00553)
Teacher pedagogy	-.0055095	(.0041063)	-0.00764*	(0.00308)
Social Legal	.0026311	(.0078826)	-0.00495	(0.00712)
Econ Adm	-.0007322	(.0038893)	0.00186	(0.00319)
Natural engin	-.0087095**	(.0028323)	-0.0138**	(0.00363)
Health sports	.0005921	(.0026872)	-0.00413	(0.00362)
Primary	-.0256651***	(.0047917)	-0.0277***	(0.00640)
Transport service safety	.004334	(.0049268)	0.00597	(0.00482)
Not given	-.0264575	(.01367)	-0.0265	(0.0131)
Immigr:A	0		0	
Immigr:B	-.0012241	(.0040899)	-0.0226**	(0.00685)
Immigr:C	-.0230954	(.0401478)	0.000743	(0.0242)
Immigr:E	.0234705	(.0168392)	-0.0290***	(0.00579)
Immigr:F	.004326	(.0061182)	-0.00709	(0.00489)
Immigr:G	.0003855	(.0124339)	-0.0241**	(0.00786)
Children=0	0		0	
Children=1	.0498309***	(.0026963)	0.0375***	(0.00281)
Children=2	.1016081***	(.0033053)	0.0720***	(0.00627)
Children=3	.1261921***	(.0064106)	0.0944***	(0.00972)
Children=4	.1600423***	(.0148656)	0.115***	(0.0134)
Children=5	.1707393***	(.0242581)	0.0652	(0.0419)
Children=6	.198247***	(.048908)	0.177*	(0.0661)
Children=7	.4848973***	(.1175907)	0.324*	(0.117)
Children=8	.3691062	(.2405031)	0.235*	(0.0861)
Children=9	.4768396	(.4106495)	0.488**	(0.157)
Children=10	.305546***	(.0637233)	0.0344	(0.0857)
Children=11			0.907***	(0.0693)
Family size= 1	0		0	
Family size= 2	.068592***	(.0023317)	0.0684***	(0.00324
Family size= 3	.066746***	(.0026339)	0.0865***	(0.00503
Family size= 4	.096362***	(.0042366)	0.142***	(0.00613
Family size= 5	.0853685***	(.0050717)	0.144***	(0.00953
Family size= 6	.0627678***	(.0094748)	0.119***	(0.0101)
Family size= 7	.070423**	(.021981)	0.104**	(0.0281)
Family size= 8	.0224686	(.0336154)	0.112	(0.0602)
Family size= 9	-.1702951*	(.0823804)	0.0255	(0.0490)
Family size= 10	-.2129183*	(.1068991)	-0.0811	(0.0936)
Family size= 11	-.0143542	(.0634729)	-0.236	(0.228)
Family size= 12	-.2905356	(.4101984)	-0.175*	(0.0715)
Family size= 13			0.160*	(0.0700)
Family size= 14			-0.0695	(0.0859)
Constant	3.119949***	(.0493425)	2.951***	(0.0542)
Observations	1 138 613		933520	

Table 11: The five worst hit

Municipality name	Municipality code
Audnedal	1027
Balestrand	1418
Bardu	1922
Beiarn	1839
Berg	1929
Berlevåg	2024
Bindal	1811
Bremanger	1438
Brønnøy	1813
Bygland	938
Tysfjord	1850
Engerdal	434
Etnedal	541
Evenes	1853
Flå	615
Folldal	439
Fyresdal	831
Gildeskål	1838
Grane	1825
Gratangen	1919
Grong	1742
Kåfjord	1940
Hamarøy	1849
Hattfjelldal	1826
Hemnes	1832
Hjartdal	827
Hjelmeland	1133
Hol	620
Hornindal	1444
Høyanger	1416
Iveland	935
Jondal	1227
Kvalsund	2017
Kvam	1238
Kvinnherad	1224
Kvænangen	1943
Lebesby	2022
Lesja	512
Lierne	1738
Luster	1426
Lærdal	1422
Marnardal	1021
Masfjorden	1266
Meråker	1711
Målselv	1924
Namsskogan	1740
Narvik	1805
Nissedal	830
Nome	819
Nord-Aurdal	542
Nord-Fron	516
Norddal	1524
Nordreisa	1942
Nore og Uvdal	633
Notodden	807
Oppdal	1634
Rauma	1539
Rendalen	432
Rindal	1567
Roan	1632
Rollag	632
Saltdal	1840
Sauda	1135
Sel	517
Snillfjord	1613
Snåsa	1736
Sortland	1870
Spydeberg	123
Steigen	1848
Storfjord	1939
Suldal	1134
Surnadal	1566
Sørfold	1845
Tokke	833
Tolga	436
Tynset	437
Ullensvang	1231
Ulvik	1233
Vaksdal	1251
Valle	940
Vang	545
Vestre Slidre	543
Vik	1417
Vinje	834
Voss	1235
Åmli	929
Åseral	1026

Table 12: The three unaffected

Municipality name	Municipality code
Alta	2012
Bod	1804
Fauske	1841
Haugesund	1106
Kvinesdal	1037
Lenvik	1931
Meløy	1837
Odda	1228
Rana	1833
Sarpsborg	105
Sunndal	1563
Sør-Varanger	2030
Tinn	826
Årdal	1424

CHAPTER III

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 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.

[1] Supermarkets 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.

[2] Previous 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.

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 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

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

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.

[4]Although 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.

[5]Individuals 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.

[6]We 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.

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.

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

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.

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
Liqour (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 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 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 prox-

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

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

Figure 1: Outlet locations in 2000 and 2015

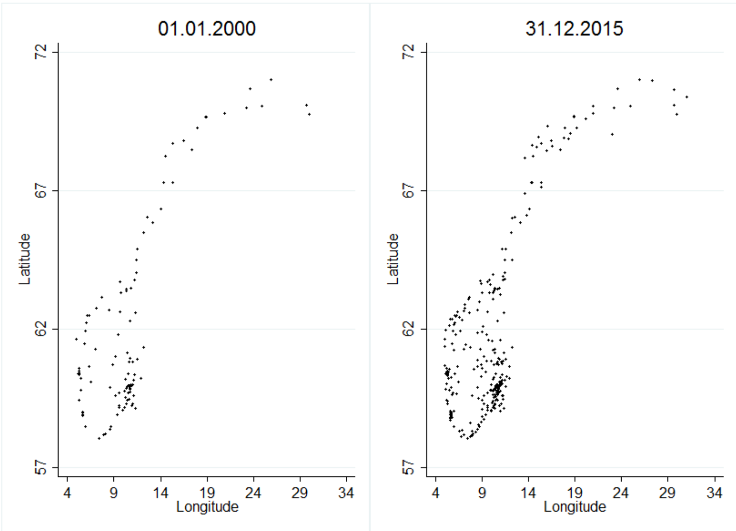

Figure 2: Number of stores over time

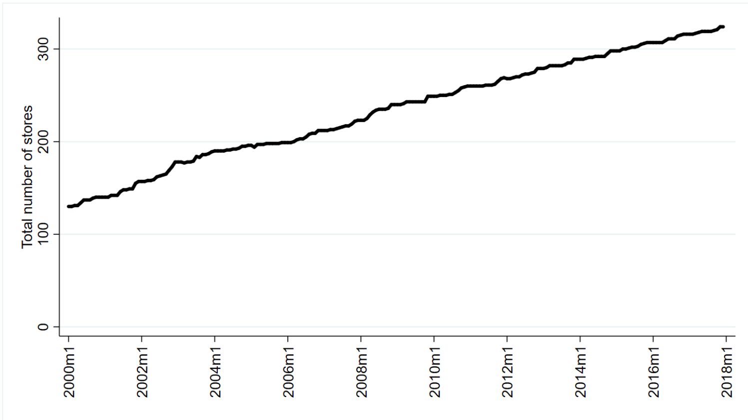

Figure 3: Frequency count of outlets in a LMR

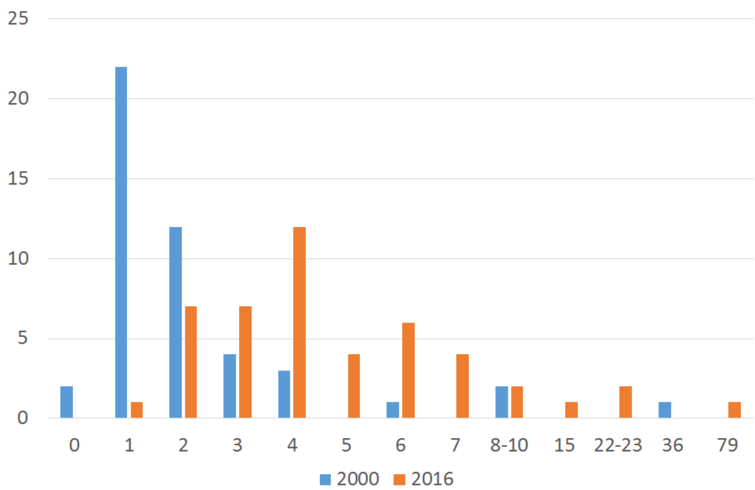

Note: Figure 1 shows the GPS-coordinates of outlets in 2000 and 2015. Figure 2 plots the total number of outlets on a monthly basis over the same time period. Figure 3 counts the number of LMRs by their total number of outlets. In 2000, 22 LMRs had only one store. In 2016, only one LMR had one store.

imity 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 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.

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

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

[10] Vinmonopolets 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 4: Average distance per capita, population and grid data.

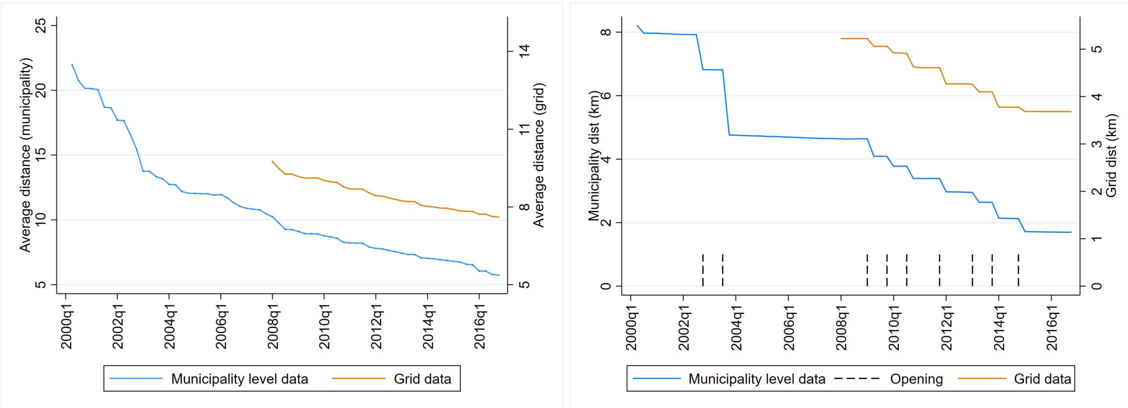

(a) Full sample

(b) LMR of Stavanger

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 of Stavanger.

with the error-term in (1), we adopt a two stage IV-approach where alcohol is instrumented by proximity in (2),

$$S_{rt} = \alpha_r + \gamma A_{rt} + \beta X_{rt} + \varepsilon_{rt} \tag{1}$$

$$A_{rt} = \pi_r + \pi_1 AD_{rt} + \delta X_{rt} + u_{rt} \tag{2}$$

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] 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,

[11] The results are almost identical using volume.

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.

[12]Recall 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***	-0.0150***	-0.0149***	-0.0147***	-0.0145***
	(0.00301)	(0.00291)	(0.00281)	(0.00285)	(0.00278)
Quadratic:					
Average distance	-0.0242***	-0.0240***	-0.0238***	-0.0229***	-0.0224***
	(0.00496)	(0.00466)	(0.00477)	(0.00432)	(0.00449)
Average distance squared	0.000181	0.000194	0.000194	0.000178	0.000170
	(0.000131)	(0.000126)	(0.000129)	(0.000120)	(0.000127)
Linear with grid data :					
Average grid distance	-0.0218***	-0.0218***	-0.0224***	-0.0233***	-0.0225***
	(0.00508)	(0.00509)	(0.00460)	(0.00487)	(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.

[13]NOK and USD in 1. January, 2015

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 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.

[14] The 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 lag	2 lags	3 lags
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
All	29.07*	32.89*	30.38*	34.08*	30.92	16.70*	15.58*	15.33
	(13.06)	(14.43)	(15.26)	(15.19)	(15.98)	(7.982)	(7.773)	(8.033)
L1.Sickness absence						0.500***	0.467***	0.465***
						(0.0202)	(0.0190)	(0.0193)
L2.Sickness absence							0.0769***	0.0762***
							(0.0204)	(0.0204)
L3.Sickness absence								-0.00939
								(0.0232)
Men	33.43*	34.47*	31.74*	35.96*	33.01*	14.41*	13.91	14.03
	(13.41)	(14.11)	(15.04)	(15.04)	(15.82)	(7.121)	(7.220)	(7.287)
Women	25.69	27.62	28.71	31.27	29.87	20.67*	19.54	18.29
	(14.83)	(16.06)	(17.00)	(16.62)	(17.87)	(10.53)	(10.15)	(10.45)
First stage:								
Average driving distance	-0.0158***	-0.0150***	-0.0151***	-0.0149***	-0.0150***	-0.0141***	-0.0141***	-0.0138***
	(0.0029)	(0.0029)	(0.0029)	(0.0029)	(0.0029)	(0.0027)	(0.0027)	(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.

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

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, 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 variable[15] and match on pre-treatment observations 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,

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

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 where 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

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

[16] Since 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
$\alpha post_t$	-0.0003	0.0013
	(0.002)	(0.00297)
ηT_{rt}	0.0159	0.0021
	(0.009)	(0.002)
$\omega post_t * T_{rt}$	0.0479 ***	-0.0011
	(0.0140)	(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

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

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,

[17] Another 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 5: Relative rankings of effects

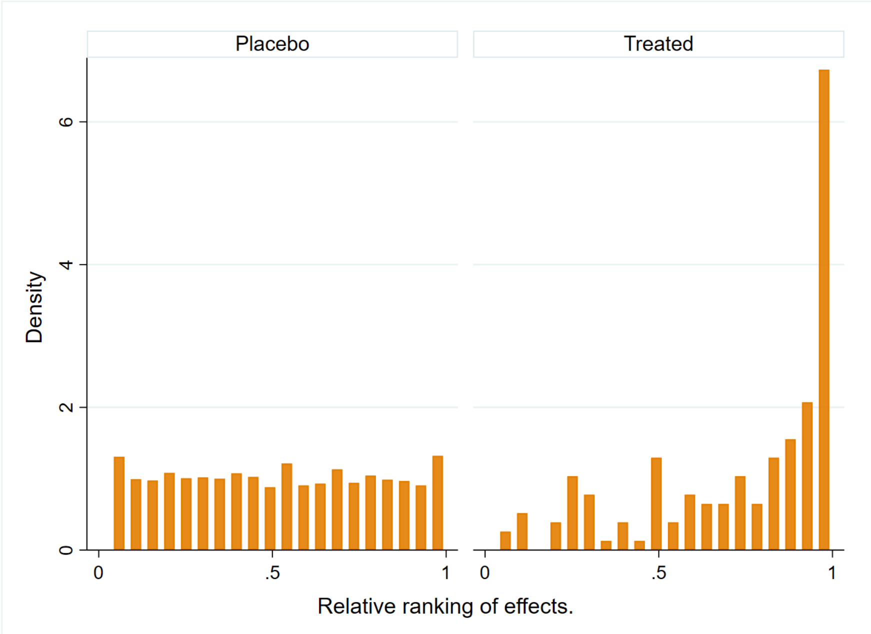

Note: Histogram showing the distribution of the relative rank, by treatment status. The left panel shows the relative rank of the effect estimated for the placebo openings. The right panel shows the relative rank of the estimated effect from the actual openings.

Figure 6: SCM-effects against distance reduction

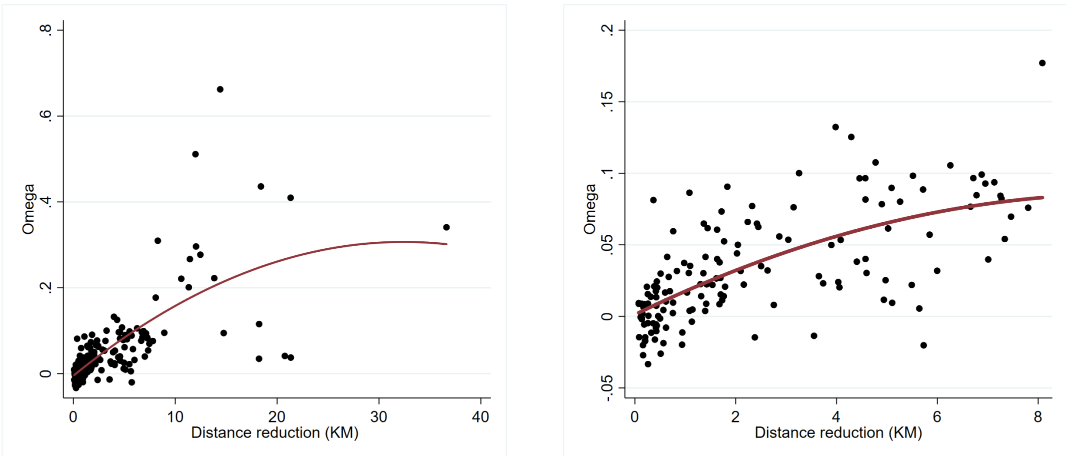

(a) All **(b)** 0-90th percentile distance reductions

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.

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 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.

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

	Baseline specification				
	(1)	(2)	(3)	(4)	(5)
Quadratic first stage					
All	28.91*	32.57*	30.28*	32.86*	29.70
	(12.81)	(13.88)	(14.76)	(14.85)	(15.81)
Men	31.49*	32.10*	29.16	32.89*	29.59
	(13.52)	(13.97)	(15.08)	(15.02)	(16.10)
Women	27.36	34.60*	32.47	34.77*	31.09
	(14.46)	(15.75)	(16.72)	(16.84)	(17.90)
Observations	2948	2948	2948	2948	2948
Excluding large distance reductions					
All	31.93**	36.62**	35.30*	33.35*	31.49*
	(11.80)	(12.84)	(13.93)	(14.04)	(15.13)
Men	35.32**	35.35**	33.50**	32.95*	31.15*
	(11.48)	(11.34)	(12.55)	(13.03)	(14.08)
Women	29.34*	39.36*	37.78*	35.73	33.07
	(14.77)	(17.24)	(18.62)	(18.43)	(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. * $p< 0.05$, ** $p< 0.01$, *** $p< 0.001$

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 Californias 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 superhighway 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," .