STOCKHOLM SCHOOL OF ECONOMICS
Department of Economics
BE551 Degree project in economics
Spring 2026

Putting a price on religion: Exploring causality between church tax increases and church exits

Niklas Saft (25819)

Abstract: This paper explores various possibilities of identifying a causal relationship between an increase in church tax and an increase in the number of people leaving the Church of Sweden (exits). By building upon previous work on the topic and applying both traditional two-way fixed-effects (TWFE) regressions and modern difference-in-differences (DiD) estimators to a panel dataset consisting of Swedish pastoral units, I analyse the variation in local church tax rates, and the number exits from pastoral units. My sample consists of 582 pastoral units during the period of 2014-2023, out of which 312 increased the church tax at least once. The empirical results show that when improving previous work by adjusting the traditional TWFE regression and applying a modern DiD estimator, the parallel trends assumption becomes considerably stronger. Nevertheless, even with this improved identification, I do not find any robust evidence that an increase in church tax leads to an overall increase in exits. While some specifications suggest effects on the timing of treatment effects, most coefficients are imprecise and highly sensitive to the choice of model specification. The findings suggest that changes in church tax play a very limited role in exit behaviour, although these results may be limited by a lack of statistical power, among other issues. As a result, I fail to reject the null hypothesis that there is no causal effect of increases in church tax on the number of exits.

Keywords: Church tax, Church of Sweden, Heterogenous difference-in-differences, Staggered treatment timing, Tax sensitivity

JEL: C23, H24, H31, Z12

Supervisor:	Sampreet Goraya
Date submitted:	11 May 2026
Date examined:	19 May 2026
Discussants:	Jens Ek Deneberg and Lovisa Öhrn
Examiner:	Johanna Wallenius

Acknowledgements

I would like to express my utmost gratitude to Gustav Svärdhagen who not only helped me with understanding this topic but also with getting hold of the necessary data to complete my empirical analysis.

I am equally indebted to my supervisor Sampreet Goraya, whom I would like to thank for his excellent insight into modern econometrics and constructive feedback, which were crucial for my analysis to be empirically robust.

A special thanks also goes to Christian Hesse whose understanding of the Church of Sweden and its internal dynamics was very helpful for my theoretical understanding of the topic.

Finally, I am offering my deepest gratitude to my parents, my friends and everyone else who offered their steadfast support and invaluable engagement throughout the process of completing this thesis.

Thank you so much,

Niklas

Table of Contents

1 Introduction 4
2 Institutional background 4
2.1 Organisational structure of the Church of Sweden 4
2.2 The church tax 5
3 Literature review 7
3.1 Church tax 7
3.2 Church tax and migration 8
3.3 Empirical analysis 10
3.4 Significance of this study 12
4 Hypothesis 13
5 Data and methodology 13
5.1 Data 13
5.1.1 Data selection 13
5.1.2 Data processing 14
5.2 Descriptive statistics 15
5.3 Methodology 17
5.3.1 Replicating the methodology by Svärdhagen 17
5.3.2 Improving the model by Svärdhagen 18
5.3.3 Using did_multiplegt_dyn 19
5.3.4 Internal validity 21
6 Empirical results 23
6.1 Replicating Svärdhagen 23
6.1.1 Empirical results by Svärdhagen 23
6.1.2 Results from replicating Svärdhagen 23
6.2 Results from improving Svärdhagen's model 24
6.3 Results from using did_multiplegt_dyn 26
7 Discussion 29
7.1 Discussion of results 29
7.2 Limitations 31
8 Conclusion and further research 32
8.1 Conclusion 32
8.2 Further research 33
Bibliography 35
Appendix 38

1 Introduction

In recent decades, the overall trend in the Church of Sweden clearly indicates a long-term decline in membership (Svenska kyrkan, 2026a). As a result, the Church of Sweden must face a reality where its revenues from church tax ("kyrkoavgift"), which represent 60% of total revenues (Svenska kyrkan, 2026b), might decline significantly over the coming decades, ultimately leading to worse liquidity and potentially cuts in financing for various services. This forces the church to rethink its financing situation fundamentally and contemplate whether taxes should be raised or not. Therefore, looking at the data surrounding previous changes in church tax and their effects on membership and thus revenues is highly relevant for the long-term sustainability of the various activities of the church.

In this thesis, I write about the church tax in the Church of Sweden and its effect on membership. By estimating how much the number of people leaving the church changes when the church tax changes, one could increase the clarity around the potential outcomes and risks of changing the church tax. I build upon a paper by Gustav Svärdhagen titled "Kyrkoavgiftsförändringar och utträden – finns det ett samband?" ("Changes in church tax and decline in membership – is there a correlation?). Due to a lack of parallel trends in his difference-in-differences (DiD) setup, Svärdhagen was unable to find a significant treatment effect (Svärdhagen, 2024).

This work examines whether the crucial parallel trends assumption can be fulfilled and valid treatment effects can be found by adjusting the analysis by Svärdhagen (2024) in various ways: First of all, I shall replicate the methodology by Svärdhagen on my dataset. Having done that, I shall add various controls and robustness checks, such as adding diocese-year fixed effects and separating increases and decreases in church tax. Because the church tax in the dataset experiences changes at different points in time and in some cases multiple times, employing a traditional two-way fixed effects DiD model, just like Svärdhagen did (Svärdhagen, 2024), may be inappropriate since this may create faulty coefficients (de Chaisemartin and D'Haultfœuille, 2020). Therefore, I attempt to address this issue with the help of the R package did_multiplegt_dyn, which allows for multiple, continuous or binary treatments and staggered implementation, based upon de Chaisemartin and D'Haultfœuille (2024), which provides the theoretical background behind the package (Naqvi, 2024).

This thesis is merely about the economics of the church tax and its effects on the number of exits from the church. A discussion about the theological and political feasibility of setting the church tax based on the resulting number of exits would go beyond the realms of a bachelor thesis in economics.

2 Institutional background

2.1 Organisational structure of the Church of Sweden

The organisational structures, internal processes and overall regulations of the Church of Sweden are all governed by the Church Ordinance (Kyrkoordningen (KO)) of the Church of Sweden (*Kyrkoordningen*, Inledningen). The general structures and regulations are also stipulated in the state legislation "Lag om Svenska kyrkan", which prescribes (among other things) the relationship

between the Church of Sweden and the Kingdom of Sweden (*Lag om Svenska kyrkan*, 1998). In the following, when referring to the KO, I shall only be referring to its latest version (January 1st, 2026) for explaining the organisational aspects that are important to know about for understanding this topic.

When I write about "the church", unless otherwise specified, I usually refer to the Church of Sweden as a whole, and not to any specific legal unit within it.

The Church of Sweden has an Evangelical-Lutheran confession of Christianity, with the Augsburg Confession being one of the central documents regulating its theology (*Kyrkoordningen*, Första avdelningen, Inledning). The Church of Sweden covers all of Sweden, which means that all people in Sweden have a local church that they can attend. On a local level, the legal entity that represents the church is called a parish ("församling"). The parish is the primary unit of the church, and it is where the primary purposes of the church are fulfilled (worship, diakonia, teaching and mission) (*Kyrkoordningen*, Andra avdelningen, Inledning).

There are also some parishes that are located outside of Sweden ("Svenska kyrkan i utlandet") (*Kyrkoordningen*, 2 kap. 12 §), but in this thesis I am not going to cover these.

One or more parishes are always part of a pastoral unit ("pastorat" in Swedish, also known as "kyrklig samfällighet"), in which case the pastoral unit has the overall responsibility that all its parishes fulfil their duties. Also, the pastoral unit will have the economic responsibility for the parishes (*Kyrkoordningen*, 2 kap. 5 §). For both parishes and pastoral units, the highest decision-making body is the church council ("kyrkofullmäktige") (*Kyrkoordningen*, 3 kap. 1 §). One of the duties of the church council is to decide upon the church tax ("kyrkoavgift") that will apply in the pastoral unit (*Kyrkoordningen*, 3 kap. 2 §).

On a regional level, the Church of Sweden is represented by several dioceses ("stift"). In addition, each diocese is divided up into a number of deaneries ("kontrakt"), which are made up by local parishes and act as an intermediate regional authority. Unlike pastoral units or parishes, deaneries do not have any economic significance (*Kyrkoordningen*, Tredje avdelningen, Inledning). In total, the Church of Sweden has thirteen dioceses (*Kyrkoordningen*, 36 kap. 6§). In addition to the local church tax, the dioceses also decide upon a regional church tax, which, together with the local church tax, adds up to the total church tax (*Lag om Svenska kyrkan*, 8§).

2.2 The church tax

"Thou shalt truly tithe all the increase of thy seed, that the field bringeth forth year by year."

- Deuteronomy 14:22 (*The Bible*. King James Version, 1611)

In this verse from Deuteronomy (Fifth book of Moses) in the Torah, Moses prescribes to the Israelites that they must give an annual so-called tithe of their produce. A tithe, literally meaning a tenth, was a form of taxation in ancient Israel and Judaism where the Israelites were obliged to pay a tenth of their agricultural earnings to the Levites and the Priests, who did not own any land themselves, so that they could still maintain the temple cult (Metzger and Coogan, 1994, p. 745).

As Christianity grew out of Judaism, the idea of financing religious activities via a tithe still lingered on and, in spite of some Christian teachers being sceptical about its Pharisaic appearance, it soon

became a moral duty for Christians to pay 10% of their personal income to the church authorities, something which became a legal obligation by the sixth century. By the eighth century, the tithe had become a well-established form of financing for ecclesiastical authorities all over Western Europe (Brundage, 2003). Sweden was no exception to this. In fact, the revenues that the church received from tithes (known as "tionde" in Swedish) not only financed priests and bishops, but also various forms of social welfare, such as economic relief for poor people and medical assistance (Skatteverket, 2026b).

Nowadays, the Church of Sweden still finances itself through a tax ("kyrkoavgift"), albeit not with the exact same structure as the medieval tithes. Instead, the church tax is levied on all church members via a fixed rate that each pastoral unit can decide upon itself. Anyone who is a member of a local pastoral unit after November 1st of any given year will have to pay the church tax the following year. The revenues from the church tax in 2026 amounted to 17.0 billion SEK and made up around 60% of total revenues of the Church of Sweden. The average church tax for 2026 is 1,05% of taxable income, the largest part of which gets decided upon by the local pastoral unit and the smaller part by the regional diocese. Importantly, everyone must pay the church tax of the pastoral unit which they are members of. And you can only be a member of the pastoral unit where you have your civil registration, with only some exceptions for members of the five non-territorial pastoral units (Svenska kyrkan, 2026b; *Kyrkoordningen*, 35 kap.). Therefore, in order to affect your church tax rate, you would either have to leave the church or move your residence to the geographical area of another pastoral unit. Skatteverket, the Swedish tax agency, helps the Church of Sweden (along with other religious groups) with collecting the church tax (Skatteverket, 2026a). The tax base for the church tax is the same as for the municipal tax for personal income (*Kyrkoordningen*, 42 kap 2 §).

If a pastoral unit wishes to change its church tax rate and wants to implement this change in the following year, it must report this change internally to the national-level church board ("kyrkostyrelsen") by the 14th of November. For dioceses, this date is the 30th of November (*Kyrkoordningen*, 42 kap. 3 §). The church board shall also each year by the 8th of March inform all pastoral units and dioceses which church tax will apply during the year, and it shall by the 10th of September provide preliminary information on the tax base for the church tax for the current income year to the pastoral units and dioceses (*Kyrkoordningen*, 42 kap. 4 §), thus providing the possibility of reacting to short-term fluctuations in the tax base.

In addition to the church tax at the diocesan and pastoral unit level, there is also a so-called "burial tax" ("begravningsavgiften"). This burial tax is the source of financing all burials in Sweden. The Church of Sweden is the legal entity which bears the responsibility for all burials in Sweden, except for the municipalities of Stockholm and Tranås, and thus receives most funding for these burials from the burial tax. However, the burial tax is not part of the church tax, and it must be paid by all Swedish resident taxpayers, regardless of religious affiliation (Svenska kyrkan, 2026d). Therefore, despite being associated with the church, the burial tax will not be included in this analysis.

In this thesis, unless otherwise specified, the church tax rate will be given in percent.

3 Literature review

3.1 Church tax

Although not the most usual theme in economic research, there are some papers which have looked at topics related to the church tax. Riegel et al. (2019) analyse the role of the church tax with regards to people leaving the church in Germany. They found that there is a positive correlation between income and the likelihood of leaving the church, that leaving the church is most likely when individuals enter the labour market and thus have to pay the tax the first time, that the likelihood remains high during the period when individuals have their highest tax burden and that men are more likely to leave than women. Importantly, this paper shows that financial incentives are not unimportant and have to be considered (Riegel *et al.*, 2019).

Birkelbach (1999) further reinforces these conclusions by looking at panel data of German students surveyed through their life cycle from the age of 15 to 43. Like Riegel et.al., Birkelbach found that the existence of a church tax significantly increases the probability of leaving the church, especially when entering employment. He also shows that an increase in the overall tax rate is associated with an increased probability of leaving the church. These correlations are stronger for Protestants than for Catholics, the latter group experiencing a stronger bond to its church according to Birkelbach. For example, Catholics are found to be much less affected by entering the labour market, with regards to their propensity to leave the church. Thus, Protestants leave more often and earlier than Catholics, indicating that there are fundamentally different mechanisms going on for each denominational group (Birkelbach, 1999).

As mentioned above, my analysis is mainly inspired by Svärdhagen (2024). Svärdhagen used a dataset consisting of the number of exits and the church tax values for all pastoral units in the Church of Sweden during 2014-2023, as well as the number of 20-34-year-olds in each of those pastoral units. He dropped any pastoral units that increased the church tax multiple times or decreased it at some point. It is unclear how he treated pastoral units that were merged at some point during the time frame. Svärdhagen attempted to estimate the effect of increases in church tax on the number of exits from pastoral units by using a difference-in-difference analysis (DiD), with all the pastoral units that increased the tax as the treatment group and all the pastoral units that left the tax unchanged as the control group. However, because the results showed that some pre-treatment coefficients were significant, Svärdhagen was unable to fully confirm parallel trends and thus could not establish a causal relationship between an increase in tax and an increase in exits (and thus a stronger decrease in membership) (Svärdhagen, 2024).

Looking at this topic from the perspective of the church, Hentze (2023) finds that, in spite of a projected nominal increase in church tax revenues of 1.9 billion Euros between 2019 and 2027, the real value of church tax revenues in Germany is projected to decline by 1.3 billion Euros. This decline is due to a substantial decrease in membership, with the annual decline in membership estimated to be around 2-3%. In 2022 alone, membership declined by 1.3 million people. As a result, German churches may have to face a reality of growing financial constraints and a need to adjust financing and expenditures. He also shows that 62.2% of church tax revenues in Germany come from the households sitting in the top 20% personal income bracket, thus making German churches highly vulnerable to behavioural responses among high-income groups (Hentze, 2023).

There have been other studies looking at the relationship between church tax and the probability of leaving the church on an individual level. Unlike Svärdhagen (2024), some of them have been

able to find significant effects. For example, Lyytikäinen & Santavirta (2013) looked at the propensity of individuals to leave the Finnish state church by analysing the effect of variation in tax liability on the propensity to leave the church on an individual level. They did this using a linear probability model including individual fixed effects. In the results, they found the price elasticity of demand of membership in the church to be -0.01 and that a 1 standard deviation increase in church tax leads to a 0.5-1 percentage point decrease in the likelihood of being a member of the church. They also show that the timing of leaving the church tends to be toward the end of each calendar year, which is the deadline before which you have to leave in order to avoid paying church tax the following year. Finally, they also found that transaction costs seem to moderate the probability of leaving, because when leaving the church was made administratively easier, there was a large increase in the number of exits (Lyytikäinen and Santavirta, 2013).

When looking at individual-level variation in Sweden, Jonsson & Babajan (2020) found in a working paper that there is a significant positive relationship between the level of church tax and the probability of leaving the Church of Sweden. In fact, they found that a 0.1 percentage point increase in church tax is associated with a 1.25 percentage point increase in the probability of leaving the church (Jonsson and Babajan, 2020). In the paper that they published in 2021, they did not include this result, but they did show several other indicators predicting a higher exit probability, such as having a higher income, having a higher education, being male, living in Stockholm, moving from one pastoral unit to another and not being confirmed (Babajan and Jonsson, 2021).

My analysis is looking at the same events (tax increase and exits), except by aggregating the individual decisions to the level of the pastoral unit, thereby answering the question whether the significant individual-level effects also are significant on an aggregate level, something which is particularly important to know about from the perspective of decision-making bodies in the pastoral units.

3.2 Church tax and migration

One of the issues that could be raised when looking at the relationship between exits and tax increases is that, instead of leaving the church altogether, people could decide to switch pastoral units by migrating from one geographical area to another. If this effect was real, it would considerably worsen the prospects of getting clear empirical results. Therefore, it is important to make sure that this is not a significant effect. Upon evaluating this issue, I have not been able to find any research looking at whether a change in church tax affects the propensity to migrate from one pastoral unit to another. Hence, there seems to be no direct evidence supporting the claim that people will be switching pastoral units upon experiencing an increase in church tax.

Nevertheless, there is some literature looking at the relationship between taxation and migration in more general terms, from which one could make (hypothetical) deductions regarding the case of church taxes and pastoral units. Kleven et al. (2020) present a thorough summary of previous research on domestic and international mobility in response to certain tax levels and tax changes, albeit mainly for certain groups, such as high-income individuals, and specific countries. One of its core results is that people do respond to differences in tax levels and there is some mobility, but again, this evidence mainly includes results from very specific groups. It has largely not been able to show similar results from more diverse demographic groups (Kleven *et al.*, 2020).

Feldstein and Wrobel (1998), which is one of the few papers analysing middle and low-income individuals, also find significant effects that taxation can have on migration, but in some specifications, they were unable to reject that there is no effect (Feldstein and Wrobel, 1998).

Regarding the top demographic groups, the evidence is also mixed: For example, whilst Moretti and Wilson (2017) find that the top 5% of inventors are strongly elastic regarding migration in response to tax differences (mobility elasticity = 1.8) (Moretti and Wilson, 2017), Young et.al. (2016) find that US millionaires show only a weak elasticity (mobility elasticity = 0.1) (Young *et al.*, 2016). Kleven et.al. (2020) also highlight that there is a big difference between domestic and international mobility. Akcigit et.al. (2018) show an example of this in the fact that the elasticity of mobility of US inventors is much higher for inventors not coming from the state of their residence (1.23) than for inventors coming from that state (0.11) (Akcigit *et al.*, 2018).

Akcigit, Baslandze and Stantcheva (2016), Kleven, Landais, and Saez (2013) and Kleven et al. (2014) show similar results regarding the difference between domestic and international mobility, although it has to be noted that these differences are likely partly driven by mechanical differences between the number of foreign and domestic individuals in a particular country (and thus the relative impact of domestics or foreigners migrating to a country) (Kleven, Landais and Saez, 2013; Kleven *et al.*, 2014; Akcigit, Baslandze and Stantcheva, 2016).

Looking further at migration effects across demographics in response to local taxes in Indiana and Maryland, Yushkov (2025) finds that county-level taxation actually reduces migration overall. In fact, he finds that both migration inflows and outflows have some coefficients which show a significant decrease in response to increased tax rates. It also appears that – at least in Indiana – the number of inflows significantly decreases more than the number of outflows. This indicates that higher taxes primarily discourage people from moving into certain counties, rather than encouraging them to move away from those counties. However, most coefficients across specifications are insignificant, suggesting that changes in taxes are not a primary driver in local-level US migration (Yushkov, 2025). On the other hand, Giesecke and Mateen (2022) do find that, after an average property tax increase of 12.3%, there was a significant decrease in township-level population by 0.35% (2.5 standard deviations) in Connecticut. In addition, they estimate that an increase in property tax by 10% leads to a significant increase in the propensity of individuals to migrate by 1.5%. However, these increases in net migration mainly consist of inter-state migration, as opposed to intra-state migration, which does not exhibit any significant increase or decrease. Also, these increases in property tax only came in response to rapid decreases in property values due to the Great Financial Crisis, i.e. a major shock to households, not a normal, marginal increase in tax rate like any other (Giesecke and Mateen, 2022).

Looking at taxation and migration in Switzerland, Liebig, Puhani and Sousa-Poza (2006) show that the level of the tax rate in Switzerland only significantly positively affects migration outflows for 21-35 year old Swiss citizens with a college degree, but even for this demographic group a 1 percentage point difference in local tax rate only increases migration outflow 0.1-0.9 percentage points between communities. The estimates for most other demographic groups among Swiss citizens are either insignificant or negative. However, when looking at changes in tax rate, the effect on migration outflow becomes positive and significant for almost all demographic groups among Swiss citizens, with 21–35-year-old college graduates now exhibiting a 2.8-3.3 percentage point increase in the migration probability in response to a 1 percentage point tax rate increase, when controlling for community expenditures. Note that in most models there is strong heterogeneity between the most mobile groups (young, highly educated people) and the least mobile groups (old,

less educated people), the latter group often only showing an effect between -0.2 and 1.2 percentage points when controlling for community expenditures (Liebig, Puhani and Sousa-Poza, 2007).

Having summarised some of the literature on the effect of taxation and tax changes on local migration, I conclude that there seems to be inconclusive evidence that migration in response to rising church taxes is a major limitation to my empirical strategy. Tax levels and tax changes do seem to play a role in determining migration, but the evidence for this suggests that any effects on migration outflows are mainly concentrated among certain demographic groups, such as high-income individuals, and major economic shocks. The lack of significant estimates of local migration effects for the broad population, in response to marginal tax changes, suggests that one could make the assumption that any such effect will not significantly bias estimated coefficients in my analysis. In addition, considering the difference in magnitude between an income tax change of 1 or more percentage points and a church tax change, whose average of the maximum annual increases during 2014-2023 was merely 0.27 percentage points in my dataset, it becomes apparent that any effect on migration by changes in the church tax will likely be small.

3.3 Empirical analysis

In recent years, several papers have been published highlighting the empirical issues arising from applying a normal two-way fixed effects difference-in-differences (DiD) model to a setting where treatment sets in at different points in time and some units experience multiple treatments. Since my dataset contains data on the church tax where some pastoral units experience a change in church tax at different points in time and several pastoral units actually experience multiple changes in church tax, the conclusions from these papers are important to bear in mind for my analysis.

Looking at the fundamental empirical issues in this case, de Chaisemartin & D'Haultfoeuille (2020) analyse in their paper what estimates traditional two-way fixed effects (TWFE) regressions (regressions controlling for group and time fixed effects) give when the treatment effect is heterogeneous. They find that, in such a case, the regression does not give an easily interpreted parameter, as it would in a normal case with homogenous treatment effects, such as an average treatment effect (ATE). Instead, they stipulate that, if the assumption of parallel trends holds, the treatment coefficient estimated corresponds to a weighted sum of several difference-in-differences de Chaisemartin & D'Haultfoeuille (2020).

When treatment happens at different points in time, the control group in one time period may have already been treated in a previous period, which can result in negative weights when already-treated groups are compared to other treated groups and treatment effects are subtracted away. Such negative weights can then lead to the treatment coefficient becoming negative, even though the treatment effect is positive in all cohorts and time periods. In other words, when treatment is heterogeneous, TWFE models are not always estimating the coefficients that you think they are. Hence, for TWFE models to work cleanly, treatment effects must be constant across all treatment cohorts and time periods de Chaisemartin & D'Haultfoeuille (2020). As a solution to this issue, de Chaisemartin & D'Haultfoeuille (2020) propose a new estimator, which estimates the ATE for "switching" units at the time when they switch from not being treated to being treated. This estimator avoids using already treated groups as control groups and is thus robust against heterogeneous treatment effects as long as the assumption of parallel trends (among others) holds. These results have important implications for many statistical studies, because de Chaisemartin & D'Haultfoeuille (2020) find that around 20% of papers published in the American Economic

Review (AER) between 2010 and 2012 use TWFE regressions where negative weights may occur (de Chaisemartin and D'Haultfœuille, 2020).

As a response to the highlighting of these issues, some more papers have been published trying to address the problems in question. In my opinion, the methodology best suited to address these issues in my particular setting is the one implemented in the R package did_multiplegt_dyn, which is based on de Chaisemartin and D'Haultfœuille (2024). The reason is that, among the many treatment designs it allows, my setting is one of them: It involves a staggered, non-binary treatment that can be increased or decreased multiple times. This is exactly what the package did_multiplegt_dyn allows for (Naqvi, 2024). I shall go into the details of the package in the methodology section 5.3.3.

The paper behind the package by de Chaisemartin and D'Haultfœuille (2024) provides the theoretical background for the package: To begin with, for this estimator to work out, two important assumptions need to be fulfilled: First, parallel trends must hold. This parallel trend assumption states that for two pastoral units which had the same initial treatment intensity, and which did not change it during the time period, their exit trends would have evolved similarly without treatment. This is actually a weaker assumption than the normal parallel trends assumption because it only requires parallel trends across groups with the same initial treatment level, not across all groups with different treatment levels. As figure 1 shows, in my setting, the correlation between church tax rate and the number of exits among pastoral units is -0.291, indicating that pastoral units with different treatment intensities may already differ systematically regarding outcomes, which would in turn imply that they may not be able to satisfy the parallel trends assumption. Therefore, to avoid having to impose stronger parallel trends assumptions, the paper restricts comparisons to groups with the same baseline treatment level. The other crucial assumption is that there is no anticipation of the future change in treatment. Hence, current outcomes must not depend on future treatment changes. In this particular setting, this means that the number of exits should not change in response to a church tax increase *before* the treatment change takes place (de Chaisemartin and D'Haultfœuille, 2024).

Having explained the assumptions, let us go into the model itself. In essence, the key idea (building on earlier work) of the estimator is based on comparing groups (pastoral units) experiencing a treatment change to groups whose treatment intensity has not yet changed. Technically, all groups may be considered as treated, because the estimator proposed by the paper characterises groups by their treatment path over time. Each pastoral unit always has some level of church tax and is thus always exposed to some level of treatment intensity. Therefore, instead of simply only estimating treated versus untreated units, the estimator estimates the effect of the respective group of units deviating from the own baseline treatment intensity. The eventual outcome under maintaining this initial treatment level is formally called the "status-quo" potential outcome, i.e. the outcome a group would have received if its treatment intensity had remained unchanged. Thus, the causal effect is defined relative to the so-called "counterfactual" in which the treatment level remains unchanged (de Chaisemartin and D'Haultfœuille, 2024).

In simple terms, the estimator by de Chaisemartin and D'Haultfœuille (2024) has the following mechanics: For each estimation, g denotes the group (pastoral unit in my case), F_g the time period in which treatment changes, $F_g - 1$ the last period before treatment changes and l denotes how many periods after the first treatment change the effect in questions occurs. For a group experiencing its first treatment change in period F_g, the proposed estimator compares the outcome evolution (in this setting the evolution of number of exits) from the period $F_g - 1$ to the period

$F_g - 1 + l$. If $l = 1$, this period is F_g, if $l = 2$ this period is $F_g + 1$ and so on. The difference-in-differences coefficients are then estimated by comparing these first differences to the differences of groups with the same initial treatment level for whom treatment has not yet changed by that same period. Hence, this is the same logic as with normal DiD estimators, except that the control group is assigned based on the baseline treatment level and on whether groups have already changed treatment intensity or not. The condition that treatment and control groups must be on the same baseline treatment level is crucial. The authors say that without this condition a causal identification would require even stronger parallel trend assumptions. In order to check whether the parallel trend assumption holds true, they propose placebo tests comparing groups whose treatment intensity changed and groups whose treatment intensity remained unchanged, with both having the same baseline treatment before the first treatment change took place. Under this assumption, the estimator identifies an effect representing an actual effect compared to the status-quo effect, i.e. the difference between the de-facto post-treatment outcome and the counterfactual outcome if the treatment change had not occurred (de Chaisemartin and D'Haultfœuille, 2024).

In the paper by de Chaisemartin and D'Haultfœuille (2024), the groups which experience a change in treatment are called "switchers" and the groups not experiencing a change in treatment "non-switchers". Switchers can both experience an increase in treatment (switching "in") and a decrease in treatment (switching "out"). "Switching in" implies an increase in treatment intensity, which in this case corresponds to increasing church tax, and "switching out" implies a decrease in treatment intensity, which in this case corresponds to decreasing church tax. The reason for this adjusted terminology is the afore mentioned focus on treatment intensity and the resulting avoidance of the treatment-control dichotomy, which may appear inappropriate in settings such as this one where treatment can increase, decrease and change multiple times (de Chaisemartin and D'Haultfœuille, 2024).

All of this means that, because the model only computes weighted averages of valid DiD comparisons, the problems mentioned above (already treated groups acting as controls and the thereby resulting negative weights) are effectively avoided.

3.4 Significance of this study

As outlined above, nobody has previously been able to find a significant positive effect of increases in church tax on the number of exits from the church on a pastoral-unit level. Therefore, establishing such a significant positive relationship would plug a considerable gap in previous empirical literature. It would imply that one could predict a change in membership following an increase in church tax. In consequence, such a result would become crucial with regards to policy decisions on a pastoral-unit level, particularly when contemplating church tax adjustments. Hence, this study contributes to the economic body of knowledge in two ways: First, it goes beyond the individual-level analyses estimating the propensity to leave the church by specific individuals to an aggregate-level analysis where the overall effects of microeconomic decisions are estimated in order to understand their influence on the wider economic situation of pastoral units and other ecclesiastical bodies. Secondly, this study provides empirical data and insights on how other associations, which are similar to pastoral units, and which also have a subscription fee for their members, can think about the economic consequences of adjusting their subscription fee.

4 Hypothesis

Considering the already published evidence on the relationship between church tax and membership, especially the analyses by Svärdhagen (2024) and Lyytikäinen & Santavirta (2013), which I have presented thus far, this thesis questions whether this relationship can be made even more explicit and evident by adjusting previous methodologies, such that they capture effects not previously captured, and applying modern econometric methodologies, such as de Chaisemartin and D'Haultfœuille (2024), which apparently have not yet been applied to this setting, even though they may be more appropriate.

The central hypothesis in this analysis is that an increase in church tax leads to a significant increase in the number of exits from a pastoral unit. The corresponding null hypothesis is that an increase in church tax has no significant effect on the number of exits.

5 Data and methodology

5.1 Data

5.1.1 Data selection

In my analysis I am using data on the number of members and exits (people leaving a pastoral unit) for each pastoral unit in the Church of Sweden and the church tax set by each individual pastoral unit during 2014-2023, the same period as Svärdhagen (2024) used, provided to me by the Church of Sweden. In addition, I use data on the diocesan church tax rates in the same time frame, also provided by the church. 2023 was the latest year that the church was able to provide data for. Before 2014 there was a major restructuring of pastoral units, which means that most pastoral units which existed independently before 2014 were later merged into larger ones, complicating the analysis. Therefore, I chose only to use the period of 2014-2023.

Data on the number of members aged 20-34 was not included, which is something Svärdhagen did to control for demographic shifts. Svärdhagen himself does not provide any evidence for an increase in the number of 20–34-year-olds during the event time window. He only mentions that there are high standard errors on the pastoral unit level (Svärdhagen, 2024). However, on a national level, the only increase in that number was 0.04 percentage points between 2014 and 2016, whereas between 2016 and 2023 it decreased by only 0.4 percentage points (SCB, 2026a). Under the assumption that these 20–34-year-olds create a higher revenue than cost for the church, one could make the assumption that an increase in the number 20–34-year-olds is approximately uncorrelated to the fact that a pastoral unit raises its church tax. The absence of a major national-level change in the number of 20-34-year-olds during 2014-2023 thus indicates that it is less likely that demographic shifts on a pastoral-unit level create an upward bias in the post-treatment event-study coefficients. Therefore, it is unclear whether including the number of 20-34-year-olds really makes a difference in reality.

For implementing the empirical analysis, I also used a few datasets from Statistics Sweden ("Statistiska centralbyrån", SCB). First of all, in order to know, among other information, which diocese each pastoral unit belongs to, I used data from SCB containing the name, code, diocese, deanery, parish(es), county (län) and municipality (kommun) of each pastoral unit during 2014-

2023, some of which is available online and some of which has to be ordered (SCB, 2026b). In order to know which of the pastoral units were merged during the period of 2014-2023, I used the biannually published documents "Indelningsändringar" ("division changes"), containing information on which pastoral units were subject to administrative changes, during the period of 2014-2023, most of which are available online (only the documents regarding administrative changes reported in 2014 and implemented in 2015 have to be ordered) (SCB, 2026c).

Finally, for creating the descriptive statistics plotting the municipal tax rates against the church tax rates below, I used data on the municipal tax rates during 2014-2023, which is available online in the SCB database (SCB, 2026d). In order to connect pastoral units to their respective municipalities, I utilised the document containing the municipal and county codes for 2026 (SCB, 2026e).

5.1.2 Data processing

For the pastoral units that merged during the time period, I created harmonised units, i.e. if two pastoral units were merged into one larger then I used the later larger one for all years and added together the number of exits for the other two old pastoral units and made a weighted church tax based on the number of members in each pastoral unit before the merger. The raw dataset contained 742 pastoral units, out of which 159 were collapsed into merged pastoral units, leaving me with 583 pastoral units, which were used as the basis for the harmonised pastoral units. On average, one harmonised pastoral unit contains 1.27 former pastoral units.

There are two special cases with respect to mergers: The first one is Göteborgs kyrkliga samfällighet, which is a pastoral unit that was divided up into several smaller pastoral units in 2018 (Lundberg, 2015). In this case, I applied the same logic as for the other mergers, except that I created one harmonised unit for the period 2018-2023. So, in essence, this case is about dividing pastoral units, not merging them, but the same logic applies for solving it. The other exception is Roslagens västra pastorat and Roslagens östra pastorat, which are two pastoral units created in 2018 out of several smaller pastoral units, some of which were divided up themselves at the same time and allocated to the two different new pastoral units (Sverke, 2015). This means that the divided pastoral units ended up in two different new pastoral units, which made it impossible to analyse those pastoral units consistently over the 10-year period. I solved this problem by simulating another merge between Roslagens östra pastorat and Roslagens västra pastorat into an even larger pastoral unit, which I henceforth called "Roslagens pastorat", thus putting all former pastoral units into one, so that they get treated consistently.

In addition to the usual church tax, there is an additional tax for all pastoral units in the diocese of Visby of 0.15%, which is aimed at preserving the cultural heritage of the many medieval churches within the diocese (Svenska kyrkan, 2026c). This additional church tax was simply added to all pastoral units within the diocese.

The dataset largely had no missing values. The only cases of missing values were the following: Kumla pastorat and Kumla, Tärna och Kila pastorat had missing values in 2014 and 2015 respectively for the number of exits. For the number of members, Kumla, Tärna och Kila pastorat had missing values in 2014 and 2015. These missing values were filled in using linear interpolation. The pastoral unit Västra Vingåkers och Österåkers pastorat had no values at all for the church tax. Hence, this pastoral unit was dropped from the dataset, thereby decreasing the number of pastoral units in the dataset from 583 to 582.

Like Svärdhagen, I exclude the diocesan church tax in the church tax values because it leads to many pastoral units being dropped when those with multiple increases are dropped, thereby substantially reducing sample size (Svärdhagen, 2024). This was important in order to replicate the method by Svärdhagen as closely as possible. Also, when the model by Svärdhagen (2024) is improved in section 5.3.2 by adding diocese-year fixed effects, it should not affect the results at all if you add the diocese-level church tax to the total church tax. The reason is that any variation in diocese-level church tax will automatically be absorbed by the diocese-year fixed effects, and thus not be able to independently affect estimation.

5.2 Descriptive statistics

Table 1 in the appendix displays basic summary statistics concerning my dataset. It is notable that there is large variation across pastoral units, e.g. the number of members ranges from 497 to 275230. Since the mean is far bigger than the median, it seems that there are a few very large pastoral units driving up the mean. The same is true for the number of exits where the mean number is almost twice the median number, indicating strong heterogeneity. Nevertheless, the relatively small standard deviation in the exit rate suggests that, despite large variation in exit levels, exit rates are relatively stable across pastoral units. Exit rate is here defined as the number of exits in year t divided by the number of members in year t. Regarding tax changes, one could note the fact that among 5820 observations, there are only 539 changes in tax (these changes are defined by a threshold of at least 0.005), out of which 76 were decreases in tax and 463 were increases in tax.

In table 2, you can see the mean members, mean church tax, mean exits and mean exit rate for each diocese. There appears to be moderate variation in church tax across dioceses, for example between the lowest average church tax in Stockholm of 0.794% and the highest average church tax in Visby of 1.487%. Interestingly, despite substantial variation in mean exits between a minimum of 51 exits on average and a maximum of 196 exits on average, the mean exit rate is fairly similar across dioceses, with the range being only between 0.008 and 0.011. The large size differences between dioceses highlight the possible need for fixed effects and other regional controls.

Figure 2 shows the distribution of church tax changes, i.e. the frequency of certain changes in church tax. Looking at the distribution, it is clear that most changes in church tax are closely concentrated around zero, with the majority of changes sitting in the range between -0.1 and + 0.1, which means that church tax changes are usually marginal and cautious rather than large and discrete. This plot also illustrates the strong asymmetry between increases and decreases mentioned above.

One issue that could be raised regarding the methodology presented in section 5.3 is that it does not control for the size of the pastoral unit. This could be problematic because pastoral units with many members are being compared to pastoral units with few members in the same way, but this may not be appropriate for the data. For example, it may theoretically be the case that large pastoral units have unproportionally more exits than small pastoral units. However, when plotting the average number of exits against the average number of members (see figure 3), it becomes clear that the relationship between exits and members is very linear, with an almost perfect correlation of 0.99. Therefore, it appears to be unnecessary to control for the number of members in the main model specification.

In figure 1 and 4, you can see the distribution of church tax vs exits and church tax vs exit rate, respectively. Starting with figure 1, you can see that there is a moderately negative relationship (correlation = -0.291) between the number of exits and the church tax rate. However, since the church tax rate and number of exits are recorded in the same year, this negative relationship is not causal in any way. The correlation may simply exist because small pastoral units both mechanically have fewer exits and a higher church tax, which you can see in figure 5, showing the negative correlation between the number of members and the church tax rate (correlation = -0.36).

Figure 4 on the other hand accounts for the size of the pastoral unit, by plotting church tax against the exit rate. When doing that instead we see that the correlation becomes essentially zero (-0.002), suggesting that there is no correlation between exits and church tax levels when accounting for the number of members. This is important because it means that church tax rates are not systematically associated with exit rates. In other words, whether the church tax for example is 0.8% or 0.9% does not really seems to matter in determining the propensity to leave the church.

Figure 6 plots the changes in church tax versus the number of exits. What you can see is a slight negative correlation (-0.1) between changes in church tax and the number of exits. Still, the data appears to be quite noisy and without a clear pattern. Looking at figure 7, which plots the change in church tax against the exit rate, the picture becomes a little bit clearer. Now we get a very small positive correlation (+0.03), which is essentially zero, representing a flat relationship between tax change and exit rates. This is a key descriptive result because it implies that changes in the church tax rate are not systematically associated with changes in the exit rate in the same year.

Figures 8 and 9 plot the total number of exits from the Church of Sweden and the exit rate, respectively, over the entire event time window. First of all, the fact that they both appear to have very similar patterns tells us that these patterns likely are determined by far more things than just the size of the Church of Sweden. Instead, it means that the propensity to leave actually changed over time and that the apparent increases in exits around 2016-2017 and 2021 likely were driven by exogenous demand shocks. This is something Svärdhagen (2024) talks about in his analysis as well, pointing out that negative publicity received by the Church of Sweden around 2016 and 2017 may have played a role in increasing the number of people leaving the church (Svärdhagen, 2024). The substantial variation in exits shown in these plots across the time period underlines the importance of using time-fixed effects in the regressions estimating effects on the number of exits.

At the same time as the rapid increase in exits around 2016-2017 the church also experienced an increase in the number of increases in the church tax, something which can be seen in figure 10 plotting the number of increases in church tax across the period of 2014-2023. The fact that the majority of the increases seem to be concentrated within the first three years of the time period puts into context the issues explained above regarding traditional TWFE estimators in cases such as this one where treatment is staggered. Because treatment is highly concentrated in a few years, some later years may only have a very limited number of "clean" control units. The high concentration of church tax increases in some years may also suggest that the increases may be responses to broader economic conditions, such as financial pressure or large-scale membership decline. Especially the latter point needs to be kept in mind in relation to issues surrounding reverse causality, which are outlined in section 5.3.4.

On average, the diocesan church tax during the period 2014-2023 was 0,056%. The lowest was 0,025% in the diocese of Stockholm and the highest was 0,11% in the diocese of Härnösand. The four dioceses that changed the diocesan church tax during the period are shown in figure 11 in the appendix.

In figure 12, I plotted the average municipal tax rate during 2014-2023 against the average church tax rate during 2014-2023. This plot was created by matching pastoral units to their respective municipalities, in which they are located. I only included pastoral units that contain only one unique municipality within their geographical area. Such pastoral units make up 97.13% of all pastoral units in the dataset. The trendline indicates a positive, albeit not particularly strong, relationship (correlation = 0.43) between the average municipal tax and the average church tax. However, the variation in municipal tax (more than 5 percentage points) is a lot larger than the variation in church tax (approximately 1 percentage point). This means that anyone willing to minimise their tax by migrating between administrative regions would profit a lot more from moving between municipalities than from moving between pastoral units. Therefore, these results strengthen the conclusions from above that changes in church tax will likely not lead to people moving from one pastoral unit to another.

5.3 Methodology

5.3.1 Replicating the methodology by Svärdhagen

First of all, I replicate the methodology by Svärdhagen on my dataset to see whether I can get similar results as he did. Since my dataset is not precisely the same as the one Svärdhagen used (for example regarding the merging logic, which is not mentioned in his analysis), the results cannot be expected to be exactly the same, but perhaps the conclusions, i.e. that pre-treatment trends make parallel trends seem unlikely (Svärdhagen, 2024).

Svärdhagen (2024) used a difference-in-differences (DiD) design, which is a statistical method employed to establish a causal relationship between a treatment and an effect. As the name implies, DiD is about establishing a difference of differences. This is being done by calculating the difference in outcome between before and after the treatment of some kind for the group that was treated (the treatment group) and then doing exactly the same thing for the group that did not receive the treatment (the control group). Finally, these two differences are taken, and the difference is calculated between them. Under the assumptions that there are no spillover effects and that the treatment group and control group would have had parallel trends with regards to the outcome after the treatment timing if the treatment had not happened, this difference-in-differences corresponds to a causal relationship (Statistics How To, 2026).

As mentioned above, in his analysis, Svärdhagen drops all pastoral units which increased the church tax multiple times or decreased it at any point. Svärdhagen uses the first year of implementation of the church tax change as the reference year, i.e. k = 0, because he assumes that people will not have had the time to leave the church before the decision was made or it became public. Thus, he implicitly assumes that there is no anticipation effect going on. Svärdhagen used nine years before the treatment and eight years after the treatment as the event time window. Standard errors were clustered at the pastoral unit level. The dataset used by Svärdhagen contained 476 pastoral units in total, out of which 176 raised the church tax during 2014-2023 (Svärdhagen, 2024).

The main model specification according to Svärdhagen is the following:

$$Y_{pt} = \alpha + \beta_0 TREAT_{pt} + \beta_1 TREAT_{pt} + \beta_2 X_{pt} + \phi_p + \gamma_t + \varepsilon_{pt}$$

Y_{pt}: The number of exits in pastoral unit p in time period t

$TREAT_{pt}$: Dummy variable for the time period relative to the treatment that is equal to 1 when it is (-9, -8, -7, -6, -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5, 6, 7, 8) periods away from the treatment

X_{pt}: The number of 20-34-year-olds in each pastoral unit at each time period

ϕ_p and γ_t: Dummy variables for time-specific and pastoral-unit specific effects constant across all time periods

ε_{pt}: Error term

However, in the paper Svärdhagen seems to be using an event study specification, partly in order to test for parallel trends and in order to show the table containing the estimated coefficients for each event time (years before the treatment and after the treatment) (Svärdhagen, 2024).

Regardless of what exactly the real main model specification by Svärdhagen looks like, the main goal with replicating his methodology is to see whether I can get the same results from his event study design checking parallel trends. Therefore, in my replication, I shall focus on the event study design:

$$Y_{pt} = (\sum_{\substack{k=-9 \\ k \neq 0}}^{8} \beta_t TREAT_{pt}^{k}) + \phi_p + \gamma_t + \varepsilon_{pt}$$

$k \in \{-9, -8, -7, -6, -5, -4, -3, -2, -1, 1, 2, 3, 4, 5, 6, 7, 8\}$

This model estimates coefficients in each year relative to the treatment in year 0. Apart from the control for the number of 20-34-year-olds, the model specification here follows the same intuition as with Svärdhagen.

Because the church tax values for the harmonised units of the merged pastoral units have many decimal places and therefore a lot of noise, I made sure to implement a threshold for treatment of 0.005 percentage points, which means that in order for a change in tax to registered as a treatment effect it must have a magnitude of at least 0.005 percentage points. Among the pastoral units which were never merged, none of them ever changed the tax by less than 0.005 percentage points. Therefore, this threshold only filters out merged pastoral units. The idea here is to filter out any tax increases that are economically meaningless and thus only distort the number of treatment/control pastoral units. I shall also apply this threshold to all other models henceforth, unless specified otherwise.

5.3.2 Improving the model by Svärdhagen

Having presented the replication of the methodology by Svärdhagen, I attempt to analyse whether his results can be improved by adjusting the model in various ways. Standard errors are always clustered at the level of the pastoral unit. The baseline model here remains largely the same, but I adjust it step by step. Here are the various adjustments to the baseline model that I undertake:

a. To start with, I am changing the reference period from k = 0 to k = -1, just in case the church tax change had already been anticipated in the previous year or reverse causality is at play. In addition, I also add diocese-year fixed effects so that I may be able to control for regional trends which you cannot recognise on a national level. Also, instead of dropping pastoral units that

increased the church tax multiple times, I decided to include them and restrict their treatment to the first increase.

b. In the next step, I divide the treatment up in increases and decreases of the church tax. Hence, instead of one, this gives two event studies, one for increases and one for decreases.

c. After this, I conduct placebo tests one year and two years prior to the actual treatment, which pretend that treatment happened earlier than it actually did, in order to check whether we still get an effect.

d. In another robustness check, I change the treatment threshold from 0.005 to 0.025 to see if it makes any difference to only include tax changes that are economically even more impactful on an individual level.

e. In the final robustness check, instead of using the number of exits as the treatment effect, I use the exit rate as the treatment effect, i.e. the number of exits per 1000 members. This may cause cyclical calculations (because the number of exits obviously affects the number of members), but it can still be interesting to see how such a specification affects the results.

5.3.3 Using did_multiplegt_dyn

As mentioned above, the R package did_multiplegt_dyn is ideal for addressing the various issues associated with DiD settings where treatment is heterogeneous. In the following, I explain the arguments that the package takes in the different specifications I use. In all the specifications, the reference period is the year before the treatment change was implemented. i.e. the last year with the old church tax before it was changed. However, unlike the previous models, when speaking of the results from did_multiplegt_dyn, this period will be referred to as k = 0. Also, standard errors will be clustered at the pastoral-unit level, and the threshold for a change in church tax to be classified as treatment change will still be 0.005. In this case, the treatment change in the main analysis is still defined as binary, not continuous. The reason is that the research question is still whether tax increases in general affect exits, i.e. not whether one-unit tax increases marginally affect exits. Also, the binary setup is preferable because it is easier to interpret causally and, as we have seen above, the treatment change sizes are highly heterogeneous.

a. Main baseline model

In my application, the baseline model takes the following arguments:

did_multiplegt_dyn (df = panel_dyn, outcome = "exits", group = "Name", time = "year", treatment = "tax_for_dyn", effects = 8, placebo = 8, switchers = "in", cluster = "Name", effects_equal = "all", design = c(0.8, "console"), date_first_switch = c("", "console"))

“df = panel_dyn” defines the dataframe used for the estimation, “outcome = “exits”” defines the dependent variable (in this case the number of exits), “group = “Name”” defines the group variable (in this case the pastoral unit). “time = “year”” defines the time variable, which in this case is annual and assumes even spacing between each period, which is the case here. “treatment =

“tax_for_dyn”” defines the treatment change variable, which in this case is the church tax adjusted for the treatment change threshold and excludes the diocesan tax. “effects = 8” stipulates that the function should estimate eight post-treatment event-study effects, which in this case is eight, in order for each treatment change to have at least one untreated period before the reference period. Similarly, “placebo = 8” stipulates that the function should estimate eight placebo (pre-treatment) event-study effects, in order to test whether treated units already had a different trend in exits before the treatment change (Quispe *et al.*, 2026).

However, as we will see, the package may only estimate less than eight placebo estimates, likely because the data does not support eight placebos, which happens quite frequently. “switchers = “in”” restricts the attention of the function to only units whose treatment increases, which in my setting implies that the function focuses on pastoral units experiencing an increase in church tax. Units whose treatment intensity decreased, on the other hand, are classified as “not-yet-treated” before their treatment change (and thus valid controls) and classified as invalid controls once they are treated. The opposite logic applies to “switchers = “out””, where units that increased the church tax are classified as valid controls as long as they are not treated themselves (Quispe *et al.*, 2026).

“cluster = “Name”” ensures that standard errors are clustered at the pastoral-unit level. In addition to these arguments concerning the mechanics of the model, there are some additional diagnostic arguments: “effects_equal = “all”” is an argument that tests whether all estimated effects are statistically equal to each other. If this is rejected, it implies that the treatment effect is dynamic rather than constant. “design = c(0.8, "console")” is a command that helps us understand what kinds of paths are being aggregated. It makes sure that treatment paths of at least 80% of relevant switchers are being shown in the console (directly in R). “date_first_switch = c("", "console")” prints out in the console when pastoral units experienced their first treatment change, by showing how many units experienced their first treatment change during each year. This may be helpful to see whether first increases in church tax are overrepresented in certain years or spread evenly across the panel (Quispe *et al.*, 2026).

b. Normalised effect per tax unit

In the next specification, I am looking into whether it makes any difference to rescale the treatment interpretation to a per-unit model. Unlike the main model outlined above, which simply estimates the treatment effect of being exposed to a higher tax path, this specification estimates something closer to the treatment effect per unit of treatment dose. This is being done with the help of the arguments “normalized = TRUE” and “normalized_weights = TRUE” (Quispe *et al.*, 2026).

c. Decreases in church tax

Having looked at the effect of increases in church tax on the number of exits, it would be interesting to see the treatment effect of decreases in church tax on the number of exits. This is being done by defining “switchers = “out”” instead of “switchers = “in””.

d. Stift-specific non-parametric trends

In order to find an analogous version of the diocese-year fixed effects in the TWFE model above, I use the argument "trends_nonparam = "stift_name"". This argument compares switchers only to controls within the same diocese-level trend group and thus controls for diocese-level time shocks (Quispe *et al.*, 2026).

e. Adding diocesan church tax

Up until now, the diocesan church tax was not included in the model specifications because it biases sample selection and may introduce measurement noise. However, in order to analyse the effect of *total tax burden* on the number of exits, the diocesan church tax would have to be included in the total church tax level. Hence, this robustness check rebuilds the dataframe for church taxes, such that the church tax values now include the diocesan church tax as well.

f. Only never-switchers as controls

Another thing that may alter the results is to adjust the control group. By using the argument "only_never_switchers = TRUE", the function restricts the control group to units that never change the church tax at all. Thus, unlike previous specifications, this excludes even units that have not yet changed the tax (but will do it later). This can create a cleaner control group, but it also may imply less data and worse precision (Quispe *et al.*, 2026).

g. Less conservative standard errors

Finally, the last robustness check includes the argument "less_conservative_se = TRUE", which alters the standard error calculation in such a way that the standard errors become less conservative. The reason to include this argument is that for units experiencing multiple treatment changes, the default standard errors calculated by the model may be too conservative (Quispe *et al.*, 2026).

5.3.4 Internal validity

As with any empirical analysis, there are some issues which may invalidate the identification strategy. Hence, I shall outline two of them here. First, one potential limitation is spillover effects, i.e. that an increase in church tax in one pastoral unit leads to a treatment effect in a different pastoral unit, thereby biasing post-treatment coefficients because the control group will be contaminated. However, I would argue that spillover effects in this setting are unlikely. One reason is that, as explained above, increases in church tax will likely not lead to people switching pastoral units because it would mean having to move residence.

In other words, the choice of pastoral unit is always linked to the choice of residence, thus making spillovers weaker in this setting compared to cases where switching between units is easier. Thus, increases in church tax only primarily *directly* affect members of the particular pastoral unit where it is increased.

Also, for spillovers to be an issue, members of pastoral units would have to be very aware of tax levels of pastoral units other than their own in order to notice tax changes in those other pastoral units. That seems highly unlikely because, as we have seen above, most of the church tax increases are fairly small and the church tax is probably not very visible to outside observers. For anyone who has that much awareness of other pastoral units, it seems likely that such a person should also be able to reliably anticipate a church tax increase (or lack thereof) in their own pastoral unit.

Hence, such active comparisons across pastoral units are possible in theory, but unlikely to widely affect behavioural responses. Another point is that, assuming individuals can receive information on church tax increases in advance, individuals have no concrete reason to leave the church based on what other pastoral units or individuals are doing. The own church tax burden is unaffected by increases in church tax elsewhere and completely unaffected by the number of exits. Therefore, despite not being able to definitively rule out the possibility of spillover effects, I would argue that they seem to be unlikely.

Another issue that could be raised regarding the internal validity of the analysis is reverse causality, i.e. the possibility that a decrease in membership leads to an increase in the church tax (in order to compensate for the membership loss). In this case, reverse causality is a credible concern, but I would argue that several aspects of the empirical strategy and institutional setting reduce the likelihood that it significantly determines the results.

If it is true that pastoral units anticipate a decline in membership before the increase in tax, then those pastoral units should look different with regard to the number of exits *before* the treatment. Whether this is the case can be tested using the pre-treatment event studies and placebo tests mentioned above. If the placebo tests do not show significant pre-treatment effects, this reduces (but does not eliminate) the concerns around reverse causality.

It is important to remember that the decision to raise the church tax is probably not a light-hearted decision taken as a quick response to a decline in membership. Instead, pastoral units may consider many different aspects when contemplating a tax raise, such as operating costs, investments, demographic changes, diocesan church tax developments, long-run financial planning etc., not just changes in membership. Therefore, it is plausible to assume that church tax changes are undertaken with consideration to many more things than just a membership decline.

Furthermore, my specifications are designed in a way that reduces the influence of reverse causality. For example, time fixed effects and diocese-year fixed effects can absorb national shocks to exits or regional shocks to exits, respectively.

Another problem with the possibility of reverse causality is that, if a pastoral unit decided to raise the church tax rate based on a decline in membership, it would have to predict the number of exits for a given year far in advance, thus making this prediction with high uncertainty. As explained above, a decision to raise the church tax has to be reported by the 14th of November each year, which means that a pastoral unit will not have enough time to reliably predict the number of exits the following year *based on the year of the decision*, especially considering the fact that the month with by far the most exits is October (Svärdhagen, 2024). Therefore, if a pastoral unit wanted to make a prediction on the number of exits in any given year based on exits in previous years, it would have to make such a prediction at least two years in advance. Such predictions are inherently uncertain and thus unlikely to be used as the *only* motivation for raising the church tax. Hence, this lag in decision-making renders the idea behind reverse causality less straightforward.

Even in case reverse causality is considerable, if anything it would most plausibly create an upward biased coefficient. If a pastoral unit expects a high number of exits and thus raises its church tax, and then the exits for whatever reason actually turn out to be high, the model might attribute the high number of exits to the church tax increase when, in reality, the true effect of the church tax increase is zero. In conclusion, I would say that the likelihood that reverse causality fully explains any estimated results is low, although I can never say that the possibility is eliminated.

6 Empirical results

6.1 Replicating Svärdhagen

6.1.1 Empirical results by Svärdhagen

In his analysis, Svärdhagen has 4308 observations. The results from his event study for checking parallel trends show that four out of nine pre-treatment coefficients are significantly different from zero. Svärdhagen concludes that these results are sufficient to doubt whether the parallel trends assumption actually holds, and that the lack of a clear break in trends at the time of treatment makes it improbable that the variation shown in the number of exits can be explained by increases in church tax. He also observes high standard errors for some periods, which he speculates to be result from a lack of observations in some periods (Svärdhagen, 2024). As you can see in figure 10, there is a large variation in the number of church tax increases across the time period. He notes that the variation in standard errors and significant values makes it difficult to build a consistent time series. Hence, he concludes that it is both difficult to confirm parallel trends and to find a clear causal treatment effect (Svärdhagen, 2024).

6.1.2 Results from replicating Svärdhagen

When replicating the methodology by Svärdhagen and applying it to my dataset, I received largely similar results, albeit with (as expected) not exactly the same numbers: In table 3 in the appendix, you can see a summary of the event-study model used to test the parallel trends assumption. I have 22 more observations than Svärdhagen and nine more treated units.

In table 4 in the appendix, you can find the various coefficients for all event times, except for the reference period ($k = 0$), their respective significance codes, standard errors and 95% confidence intervals (CI) and in figure 13, you can see the visualisation of the numbers in table 4. Although not identical, the results are similar in several respects: Firstly, there are two significant pre-treatment coefficients, which is not as many as Svärdhagen had, but still enough to cast doubt on the parallel trends assumption. Nevertheless, since the joint pre-treatment p-value (which checks whether all pre-treatment coefficients are jointly equal to zero) is 0.57, I cannot fully reject parallel trends, thus making it difficult to draw clear conclusions. Hence, despite seeing multiple significant and positive post-treatment coefficients (just as with Svärdhagen), it is uncertain whether these are actual causal estimates or merely results from other effects. Standard errors are high in my case as well, likely because I have a similar number of observations, and the same problem that the number of increases in church tax is not constant across all time periods.

6.2 Results from improving Svärdhagen's model

When attempting to improve the methodology by Svärdhagen in the various ways outlined above, I have received the following results:

a. Adding diocese-year fixed effects, changing the reference year to k = -1 and keeping pastoral units with multiple increases in church tax:

Upon implementing these first steps, we can immediately see an effect on the sample selection in table 5 in the appendix: When not dropping pastoral units that increased their church tax multiple times, and instead using their first increase as the treatment, the number of pastoral units in the regression jumps from 433 to 525 and the number of treated pastoral units jumps from 185 to 277.

In table 6 and figure 14 in the appendix, the effects become even more clear. Now, when including diocese-year fixed effects, setting the reference year to k = -1 and including pastoral units that increased the church tax multiple times, the significant pre-treatment coefficients disappear. Instead, we can only see insignificant coefficients, although some of them are almost significant and indicate some sort of increase in the number of exits right before and a decrease right after the treatment. In addition, the coefficients are still jointly insignificant (joint pre-treatment p-value = 0.28). There are no significant coefficients after the treatment. These results can be said to be a slight improvement from the method used by Svärdhagen. Because there are no significant coefficients, parallel trends appear slightly less unlikely, but it also means that we cannot confirm a causal relationship between increases in church tax and increase in exits. In addition, standard errors are still considerable.

b. Separate event studies for increases and decreases:

The results from doing separate event studies for increases and decreases, with the same adjustments as in step one, can be seen in tables 7-10 and in figures 15-16 in the appendix. In these specifications, no pastoral units were dropped, hence the high number of treated/control units.

The results for the increase-only event study in table 7, table 9 and figure 15 look very similar to previous specifications: All pre-treatment coefficients are both individually and jointly insignificant, and they follow a similar pattern. The most notable difference to previous specifications is that there is one significant post-treatment coefficient, the one in period 6, indicating weak evidence for significant treatment effects of tax increases. This means that, when including pastoral units that decreased the church tax in the control group, there seems to be some positive treatment effect, but it is slow and imprecise. In fact, when trying to estimate a single event study coefficient for this particular specification, it becomes insignificant.

The results for the decrease-only event study in table 8, table 10 and figure 16 paint an entirely different picture: Now there is strong evidence for pre-treatment trends. Two pre-treatment coefficients are significant and, crucially, the joint pre-treatment p-value is 0.0000, indicating a systematic pattern in exits in the periods before the treatment for pastoral units which are about to decrease the church tax. This is important because it means that pastoral units which decrease the church tax are systematically different from pastoral units which increase it in terms of how many exits they experience before the decrease in tax, which in turn could support the previous model specifications where those pastoral units were dropped. The post-treatment coefficients are highly significant and negative in three cases, but these are not credible from a causality standpoint.

c. Placebo tests one and two years before treatment:

The results from the placebo tests one and two years before the actual treatment timing are shown in table 11 in the appendix. These placebo tests are done for the specification used in step a. of adjustments. The fact that both the p-values for one year and two years are insignificant suggests that there is no detectable difference in exits before the tax change, thus indicating that there is no evidence of pre-treatment effects. In other words, treatment and control units behave similarly before the treatment, which strengthens the parallel trends assumption.

This result is important because, as mentioned above, the absence of significant pre-treatment placebo test coefficients weakens the concern that there might be reverse causality going on.

d. Changing the treatment threshold from 0.005 to 0.025:

When adjusting the model by changing the treatment threshold from 0.005 to 0.025, the pre-treatment coefficients are still all individually and jointly insignificant, indicating that the parallel trends assumption may hold. These results can be seen in table 12, table 13 and figure 17. However, instead of merely one significant post-treatment coefficient, we now get two, in periods 6 and 7. This indicates that large tax increases do matter, albeit after some time has passed, which may be due to delayed awareness of the tax change or gradual decision-making. Smaller changes in tax are thus likely too weak to trigger behavioural responses on the aggregate pastoral-unit scale.

e. Using the exit rate as the treatment effect:

When using the exit rate as the treatment effect, there is one significant pre-treatment coefficient, which may indicate that the parallel trends assumption is less likely to hold. However, it is almost insignificant (p-value = 4.5%), all other pre-treatment coefficients are insignificant, and they are all jointly insignificant. These results can be found in table 14, table 15 and figure 18 in the appendix. Among the post-treatment coefficients, two are significant, in period 1 and 4. The fact that the coefficient for year 1 (-0.4) is significantly negative is interesting because it implies that pastoral units that raise the church tax have a negative exit rate right after the increase. However, this may also be a mechanical effect or a timing/measurement issue, especially considering the fact that figure 7 paints a different picture. In addition, it is notable that all coefficients have a magnitude smaller than 1. Hence, even if there is a causal effect, it only means that an increased church tax rate leads to one additional exit per 1000 members. From the perspective of the pastoral unit, such an effect would only be marginal, because it would imply a decrease in revenue by on average only 0.2% (assuming that half of the members are taxpayers and the one additional individual leaving the pastoral unit is random with regards to the size of their tax base). These results also confirm the fact that members and exits are almost perfectly correlated.

6.3 Results from using did_multiplegt_dyn

a. Main baseline model:

The results from the main baseline model with the did_multiplegt_dyn function can be found in tables 16-19 and figure 19 in the appendix. From the pre-treatment placebo tests (table 18), whose coefficients are all individually and jointly insignificant (table 16), I conclude that there is little evidence of pre-treatment trends. Thus, unlike Svärdhagen (2024), parallel trends appear to be more credible. Looking at the post-treatment coefficients (table 17), we can find that tax increases are followed by negative, albeit insignificant, short-run effects on exits. In later periods, the coefficients become positive, and effect 7 is even significant, albeit only barely. It seems possible that an increase in church tax may take a few years for people to notice and then react to it. However, jointly, the coefficients are still insignificant at the 5% significance level (p-value = 6.02%).

The highly insignificant average cumulative effect (table 19) indicates that there is no clear evidence of an overall effect of increases in church tax on exits. This is likely due to imprecise effects or short-run negative effects being offset by long-run positive effects. Hence, this baseline model suggests that parallel trends are not rejected and increases in church tax may affect exits in some periods or the timing of exits, but not the total number of exits after the treatment.

Also, the fact that the dynamic pattern is so variable and the one significant coefficient is almost insignificant paints a picture where it is difficult to draw definitive conclusions regarding causality. One thing to be remembered is that later effects have less switchers, thereby possibly reducing statistical power for those coefficients, for example effect 1 has 228 switchers whereas effect 7 only has 153 switchers (table 17). This decrease in the number of switchers for later effects will remain for all other specifications as well, which has to be kept in mind because of the reduced statistical power for those later event times.

b. Normalised effect per tax unit:

The results from the model with a normalised effect per tax unit can be found in tables 20-23 and figure 20. Just as with the main baseline model, the pre-treatment placebo coefficients are jointly and individually insignificant (table 20 & 22).

The post-treatment coefficients also show a similar pattern, with short-run negative and insignificant coefficients, and long-run positive coefficients, with effect 7 being significant again (table 21). The joint p-value of post-treatment coefficients (table 20), as well as the average cumulative effect (table 23), are also insignificant again at the 5% significance level.

Therefore, qualitatively speaking, normalising the treatment effect does not strengthen the results, rather it rescales them. Note that the estimation produces a warning indicating that there may be a high degree of correlation between the estimated coefficients in certain periods, which means that coefficients and significance levels ought to be interpreted cautiously, thereby diminishing the strength of this specification even further.

The fact that the average cumulative effect is insignificant means that if you wanted to calculate the elasticity of demand for being a member of the church, you would not be able to do that because you would not know the overall sensitivity of members of the church against marginal

changes in church tax. But this is also not surprising considering the results above indicating that even increases in general, i.e. regardless of the relative increase, do not create an average cumulative effect on exits.

c. Decreases in church tax:

The results from the robustness check for decreases in church tax can be found in tables 24-27 and figure 21. These results paint a slightly different picture than the previous ones. First of all, it becomes clear that all coefficients are jointly and individually insignificant, both pre- and post-treatment (table 24 & 25), thus tax decreases do not appear to affect exits at all. However, the standard errors and confidence intervals are very large, especially for the placebo coefficients (table 26), indicating that the results are too noisy to be statistically informative. The same is true for the average cumulative effect, which shows an extremely large standard error and confidence interval (table 27). In addition, the small number of switchers (pastoral units that decreased the church tax) likely causes this specification to be critically underpowered, because at most there are only 23 switchers (table 25). All of this implies that this particular specification is not informative for causal inference.

d. Diocese-specific non-parametric trends:

The results from the robustness check controlling for diocese-specific non-parametric trends can be found in tables 28-31 and figure 22 in the appendix. Looking at the placebo effects in this specification, a finding that stands out immediately is the significant placebo effect 4, which possibly weakens the parallel trends assumption (table 30). However, the joint pre-treatment p-value remains insignificant, thus not necessitating a complete rejection of parallel trends (table 28). Regarding the post-treatment coefficients, there is still only one significant coefficient which is not part of a consistent pattern (table 29).

Hence, it is no surprise that the average cumulative effect also is insignificant (table 31). One interesting difference to the baseline model is the finding that the short-run negative coefficients largely disappear in this specification and that they instead become more positive and slightly unstable, indicating that the baseline results may be partly influenced by diocese-specific time trends (table 29).

This suggests that the original results may partly have been driven by such diocese-specific trends, rather than variation in church tax changes, which in consequence weakens those original results. Therefore, the evidence appears to be slightly more inconclusive than initially assumed. However, given the small number of switchers, which is between 18 and 93 (table 29), the findings are likely underpowered, suggesting that the results from this specification should be interpreted cautiously.

e. Adding diocesan church tax:

The results from the robustness check for including the diocesan church tax in the church tax variable can be found in tables 32-35 and figure 23 in the appendix. When including the diocesan church tax in the calculation and thus using the entire church tax rate faced by each individual, it is interesting to see that increases in total tax do not appear to significantly affect exits. All

coefficients, both pre- and post-treatment, both jointly and individually are insignificant, meaning that, in spite of a lack of evidence for pre-treatment trends, no meaningful conclusions can be made from the post-treatment coefficients, except that they resemble the baseline model (initially being negative and then progressively positive) (table 32, 33 & 34). Similarly, the average cumulative effect is highly insignificant (table 35). These results indicate that even when using total church tax burden, there is no significant effect on the number of exits from an increase in total tax. Considering that the average diocesan church tax increase during the time period of 2014-2023 merely was +0.021 percentage points, this result is not very surprising because on an individual level a 0.021 percentage point increase in tax burden may not be very impactful. Therefore, the original conclusion that there is no strong causal effect stays the same.

f. Using only never-switchers as controls:

The results from the robustness check using only never-switchers as controls can be found in tables 36-39 and figure 24 in the appendix. Now, in figure 24 we immediately see that for the first time we get a significant negative coefficient in the first post-treatment period, indicating a short-run decrease in exits after the tax increases. As with the baseline model, effect 7 is also significantly positive, whereas all remaining post-treatment coefficients are insignificant (table 37) as well as the joint p-value (table 36).

Thus, the pattern is very similar to the baseline results, except that we get a significant effect already in the first post-treatment period. As the placebo coefficients still are jointly and individually insignificant, parallel trends become credible, and the identification is still valid (table 38).

The average cumulative effect is still insignificant (table 39). While the short-run negative effects already observed in the baseline model now become slightly stronger, using never-switchers as controls still shows no persistent treatment impact pattern, thereby supporting the baseline model and leaving overall conclusions largely unchanged.

g. Less conservative standard errors:

The results from the last robustness check with less conservative standard errors can be found in tables 40-43 and figure 25. As with the previous robustness check, the post-treatment effects 1 and 7 become significant, effect 1 being significantly negative and effect 7 significantly positive (table 41).

However, the interesting thing now is that the joint post-treatment p-value is significant at the 5% significance level (p-value = 0.0397) (table 40). Hence, in this specification we reject that there is no effect and suggest that there are some statistically detectible effects. The qualitative pattern with short-run negative effects and long-run positive effects remains the same, except that there is a significant coefficient both short-term and long-term (figure 25).

Importantly, since placebo coefficients are individually and jointly insignificant, the identification remains valid (table 40 & 42). The cumulative effect still remains insignificant (table 43). It is very tempting to draw causal conclusions from this final robustness check, but because the point estimates are very similar to the baseline model and the qualitative pattern is identical, it seems that the increased significance comes from the less conservative inference, not from stronger effects. Therefore, even though this is a very interesting result, we need to be cautious and cannot not change the overall conclusions.

7 Discussion

7.1 Discussion of results

Looking at the collection of results presented above with regards to my original hypothesis yields somewhat mixed conclusions. The first TWFE event-study model, in which the methodology by Svärdhagen was replicated, suggests that the design was somewhat fragile. The existence of significant pre-treatment coefficients raised concerns about parallel trends and thus the validity of causal inference regarding the significant post-treatment coefficients. Therefore, these results confirm the conclusions by Svärdhagen that, due to identification issues, a simple TWFE model is insufficient as main causal evidence.

When improving the simple TWFE design by Svärdhagen in various ways, such as not dropping units with multiple increases, changing the reference year to k = -1 and adding diocese-year fixed effects, the evidence for pre-treatment trends largely disappears, thus making the design appear more credible.

However, the lack of significant post-treatment effects in the main improved specification suggests that the earlier apparent significant coefficients may have been driven by diocese-specific or time-varying effects, rather than actual variation in tax changes. When adding additional robustness checks, these baseline results experienced slight alterations, such as showing single significant post-treatment coefficients (increases only and increasing the treatment threshold), but these effects are often late (lagged), appearing only in year 7 or 8 after the treatment. This indicates that the post-treatment coefficients and significance levels are slightly sensitive to the specific model specification, but that even these are not stable nor systematic.

If anything, from figure 17 it seems that only large increases in church tax make any difference, and even then, effects are delayed and imprecise. Therefore, we can conclude that the improved TWFE specifications strengthen the credibility of the parallel trends assumption, thereby increasing the validity of causal estimates. However, because those very causal estimates are now largely insignificant or very sensitive to model specification choices, any real causality cannot be definitively confirmed.

Across the different robustness checks for the improvements there is no consistent pattern. Rather, we can see that some specifications show isolated significant post-treatment coefficients and some show no significant coefficients at all. Overall, these improvements to the TWFE specification by Svärdhagen (2024) show little robust evidence that increases in church tax led to increases in the number of exits. Still, considering the fact that this identification strategy is somewhat stronger than that by Svärdhagen (2024), the conclusion that there is no causal effect is strengthened.

Because of the strong suitability of the modern did_multiplegt_dyn estimator by de Chaisemartin and D'Haultfœuille (2024), the results from this model are likely to be the strongest in this study. Across the different specifications from this model, a clear pattern emerges: The placebo (pre-treatment) coefficients are all insignificant in most specifications, indicating that parallel trends cannot be rejected in those specifications and that the design is credible for causal inference. Despite this improved identification, like the pre-treatment coefficients, most of the post-treatment coefficients are also insignificant, with only isolated coefficients, such as effect 7, reaching significance.

The joint tests of all post-treatment coefficients are also almost always insignificant, and the average cumulative effects are consistently insignificant. That is the most crucial result from this model: Even when parallel trends cannot be rejected and the identification strategy thus is valid, there is no robust evidence that increases in church tax lead to an increased number of people leaving the church (exits). The specification which at first sight looks like showing the strongest causal evidence is the one with less conservative standard errors, because this is the only one where the post-treatment coefficients are jointly significant at the 5% significance level. However, when taking a closer look at the results and realising that the point estimates are very similar to previous specifications and that the average cumulative effect is still insignificant, it becomes clear that an increase in church tax does not lead to an overall increase in exits.

All remaining specifications show either no change in the results from the baseline model or a slight weakening of it (diocese-specific non-parametric trends). Regarding post-treatment coefficients, across the specifications there is a recurring, albeit unstable, pattern showing short-run decreases and long-run increases in the number of exits. Importantly, the fact that some of the short-run post-treatment coefficients are negative and even significant in some specifications, suggests that there is some evidence against strong reverse causality because, as mentioned above, if there was strong reverse causality, the short-run post-treatment coefficients should be positive. Note that, despite recurring patterns being visible, these short-run and long-run patterns are imprecise and not exactly identical across all specifications, which also is another argument against reverse causality.

This suggests that the increase in church tax may somewhat affect the timing of exits, but not the total number of exits for the full post-treatment window. Applying the modern DiD estimator based on de Chaisemartin and D'Haultfœuille (2024) to this data thus leads to the conclusion that when addressing the statistical difficulties related to this particular data setting, there is no significant, robust and consistent evidence that an increase in church tax leads to an increase in exits from a pastoral unit.

With regard to these empirical results and conclusions, we can say that, although there is some suggestion that in the short term and under some assumptions there may be a causal effect, the combined results lead us to the conclusion that there seems to be no significant and overall causal effect. In other words, going back to our original hypothesis, we fail to reject the null hypothesis that there is no causal effect.

In summary, we cannot say for sure that there is a causal relationship between an increase in church tax and an increase in exits. Simple TWFE specifications show some evidence for a causal relationship, but they also have substantial pre-treatment trend concerns. More robust TWFE specifications and especially the modern DiD estimator explained above largely eliminate these pre-treatment trends, but are unable to show significant, consistent and meaningful post-treatment effects. This means that even the most promising modern specifications are unable to display persistent patterns and show at most weak and isolated significant coefficients.

7.2 Limitations

My analysis is exposed to several important limitations, mainly concerning data, statistics and internal validity, some of which I shall explain in the following. First of all, looking at the data, it is notable that its aggregated setup may not be the best. As we saw in the previous literature, there appears to be a significant causal relationship between increases in church tax and exits on an *individual* level.

When aggregating up to the pastoral-unit level, it may be difficult to still retain that causal identification, possibly due to individual effects cancelling out and heterogeneity being hidden away. Hence, the use of aggregate data may mask individual-level responses and create the false impression that there is no effect, when, in fact, there may be many different effects.

Another issue in my particular dataset are the harmonised units due to mergers. The harmonisation method which was employed here may not be the most appropriate way to correct for mergers of different pastoral units. Instead, one could think of a method where, instead of using the number of members as the base for the weights determining the harmonised church tax, you would use the total tax base of members within the geographical area of the former pastoral unit. Using the total tax base instead of simply the number of members would eliminate the issue that some pastoral units may have disproportionally high tax bases relative to the number of members. For example, if two pastoral units that at some point in time were merged have the same number of members but a completely different total tax base, in my calculation the harmonised church tax would be a perfect weighted average of the two original church tax rates, whereas in reality the vast majority of church tax revenues come from the pastoral unit with the higher tax base. Hence, it would be fairer to allocate a larger weight to the pastoral unit with a higher total tax base, in order for the harmonised church tax rate to correspond the de facto revenues received from the members within the merged pastoral unit. Therefore, it is possible that some of the harmonised church tax rates in my processed dataset do not perfectly reflect the revenues received by the newly merged pastoral unit.

Another limitation related to the dataset in my analysis is simply the *lack of data*. Having a dataset containing only ten years of data naturally creates a limited time horizon. This is especially concerning, considering that there is a very limited amount long-run observations, with many increases in church tax rates only having a few post-treatment years to observe effects. It is fully conceivable that treated pastoral units may have different trends more than ten years before or after the treatment, something which I am unable to detect in my relatively short-run analysis. The limited number of observations likely also creates issues with statistical power, especially in the later post-treatment periods, with the number of switchers declining over the event time window.

The lack of statistical power is probably one of the biggest limitations in my analysis: In most specifications, standard errors are quite high and coefficients are imprecisely estimated. This in turn leads to considerable instability with regards to results across the different specifications. In consequence, the absence of significant causal estimates may partly be the result of limited statistical power, not the absence of a true causal effect, which would imply that my analysis fails to estimate the true effect. One reason for this may be that the treatment variation in my dataset is fairly small, with the median tax increase being only 0.05 (table 1). While these changes can have a large impact on the revenues of a pastoral unit, it is highly questionable whether a 0.05 percentage point increase in the tax rate makes an economically meaningful difference for individual taxpayers.

Hence, it is possible that any behavioural responses are too small to detect on an aggregate pastoral-unit level. In addition, it seems likely that other unobserved time-varying factors play a role in determining the number of exits from a pastoral unit. Such factors may be the number of 20-34-year-olds, which Svärdhagen included as a control in his analysis, or perhaps changes in municipal tax rates that happened simultaneously. If, for example, the municipal tax rate increased for some pastoral units at the same time as they increased their church tax rate, some individuals may have decided to leave the church in order to cut the costs that they are able to cut. The absence of such other demographic and economic controls is a clear limitation in my analysis.

Another point is that, although its existence is unlikely, it is difficult to completely rule out the possibility of reverse causality. I present several arguments against it in section 5.3.4, the placebo tests and pre-treatment event study coefficients do not indicate it and the short-run post-treatment coefficients are negative, instead of positive. Still, all of these arguments and results do not imply that the problem of reverse causality is completely eliminated.

Finally, it is important to note that this study only concerns itself with the church tax and membership in the Church of Sweden and is thus not fully applicable to other settings. The Church of Sweden has very specific institutional structures, church tax policies and regulations regarding membership and leaving the church. Therefore, any results and conclusions obtained in this analysis may not be easily generalised to other countries or settings with different institutional regulations and economic situations.

The same is true for the time period of 2014-2023, which may reflect specific trends which were going on at that time. Hence, we ought to be careful when applying conclusions from this time window to our current or future situation. Taken together, the limitations outlined here suggest that the absence of an overall significant relationship between church tax increases and the number of people leaving the church should not be interpreted as a representation of true behavioural responses. Rather, it should be interpreted cautiously, i.e. as a reflection of the specific dataset and identification constraints.

8 Conclusion and further research

8.1 Conclusion

Having discussed the main results of my empirical analysis and outlined their most important limitations, the deduction can be made that this study is unable to draw a definite conclusion on the relationship between changes in church tax and changes in the number of people leaving the church. On one hand, the improved empirical analysis, with for example more evidence supporting the parallel trends assumption, makes it more credible that there truly is no overall causal effect.

On the other hand, the big issues surrounding the lack of statistical power, the absence of significant cumulative effects, large standard errors, the sensitivity of results to different specifications, among other things, suggest that even these results should be taken with a grain of salt. On the whole, this means that we can accentuate that this analysis credibly improves the results brought forward by previous analyses, but the results are still not perfect, which, as mentioned above, leads us to fail to reject the null hypothesis that there is no causal effect.

As explained in the literature review, previous literature suggests that the church tax *does* play a role in people`s decision to leave the church, but it seems that the particular level of the church tax does

not really affect the decision that much. In other words, people leave the church because of the church tax, but whether the church tax rate is 1.0% or 1.05% does not make any particular difference, especially when summing up all individual behavioural responses to an aggregate pastoral-unit level. Putting this in economic terms, it means that demand is relatively inelastic towards the marginal changes in church tax apparent in my dataset.

From the perspective of pastoral units, these conclusions are very interesting with respect to policy decisions, such as raising or lowering the church tax. However, you need to be cautious not to make hasty decisions based on this evidence. You cannot say for sure that when increasing the church tax there will be no effect on membership at all. The only thing you can say is that there is no evidence supporting the claim that within the first eight years after the treatment there will be no significant overall increase in the number of exits. What happens after those eight years is very uncertain. Therefore, these results can be used as complementary data and to reassure policy makers, but since there is still so much uncertainty around them, they should probably not be used as a single justification to raise the church tax. In conclusion, the results from this study are very interesting and important to keep in mind for the church, but policymakers need to be aware of their limitations.

8.2 Further research

The various limitations of my analysis and the fact that it was unable to produce definite results suggest that there is still a lot of research to be done. First of all, there are plenty of possibilities to exploit the gaps surrounding the statistical limitations described in section 7.2, the most important of which likely is the lack of statistical power. Increasing the statistical power by for example expanding the event time window would be a natural first step. In addition, there is a case to be made that one could attempt to exploit heterogeneity in treatment effects across different demographic groups, such urban versus rural populations, high versus low-income individuals etc. Since my baseline models do not control for these factors, they might cancel out differences between these groups.

Furthermore, work still needs to be done on individual-level analyses, rather than aggregate pastoral-unit level analyses. Lyytikäinen & Santavirta (2013) found that it is totally possible to establish significant causal estimates for the elasticity of demand for being a member of the church in Finland (Lyytikäinen and Santavirta, 2013). It would be very interesting to see whether one could apply their methodological approach to the Swedish institutional context or even any other country.

Another point which may require more research is the methodology employed for solving the DiD identification issues. de Chaisemartin and D'Haultfœuille (2024) already present a very good solution, but it would be compelling to see whether their method could be improved even more by making it more robust against the various issues in my dataset. de Chaisemartin et al. (2026) published a working paper in April 2026 exploring the possibilities of employing DiD estimators where all units eventually receive some treatment dose, thus leaving no remaining never-treated control group at the end. They propose one solution where they separate treated units based on treatment intensity and stipulate that the control group consists of so-called "quasi-untreated" units, i.e. units only exposed to minimal treatment intensities. Under the assumption of parallel trends, they explain that causal identification can be established even when all units are treated at some point (de Chaisemartin et al., 2026). This setup could substantially facilitate an analysis of my dataset. Especially when expanding my event time window, you may have to face a situation where

the majority of pastoral units experienced a small increase in church tax, i.e. below 0.05 at some point in time. Hence, applying a setup where the control group becomes slightly larger may be able to ensure that statistical power remains high even when the vast majority of units receive some level of treatment dose during the event time window. Thus, this area of research ought to be monitored closely.

Finally, another conceptual issue worth considering is how cognitively important the church tax is to individuals. It may be the case that the church tax rates are not transparently communicated in every pastoral unit or that people simply are unaware of it because they are not actively choosing it. Hence, regarding the church tax rate, there may be an element of information asymmetry going on where individuals do not know what exactly they are paying for or when the tax has changed the last time. Therefore, the fact that I failed to reject the null hypothesis may not be due to an absence of economic incentives, but rather a low amount of salience regarding the church tax among the members of a pastoral unit, since it may be possible that members do not remember the exact church tax rate when preparing their annual tax return. It would be intriguing to conduct an analysis exploring whether the amount of salience acts as a mediator for the relationship between church tax increases and membership declines.

Bibliography

Akcigit, U. *et al.* (2018) *Taxation and Innovation in the 20th Century*. Cambridge, MA: National Bureau of Economic Research, p. w24982. Available at: https://doi.org/10.3386/w24982.

Akcigit, U., Baslandze, S. and Stantcheva, S. (2016) 'Taxation and the International Mobility of Inventors', *American Economic Review*, 106(10), pp. 2930–2981. Available at: https://doi.org/10.1257/aer.20150237.

Babajan, T. and Jonsson, P. (2021) 'Leaving a Folk Church: Patterns of Disaffiliation from the Church of Sweden', *Nordic Journal of Religion and Society*, 34(1), pp. 40–55. Available at: https://doi.org/10.18261/issn.1890-7008-2021-01-04.

Birkelbach, K. (1999) 'Die Entscheidung zum Kirchenaustritt zwischen Kirchenbindung und Kirchensteuer', *Zeitschrift für Soziologie*, 28(2), pp. 136–153.

Brundage, J.A. (2003) 'Tithe', *The Oxford encyclopedia of economic history*. Edited by M. Joel. Oxford, New York: Oxford University Press.

De Chaisemartin, C. and D'Haultfœuille, X. (2020) 'Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects', *American Economic Review*, 110(9), pp. 2964–2996. Available at: https://doi.org/10.1257/aer.20181169.

De Chaisemartin, C. and D'Haultfœuille, X. (2024) 'Difference-in-Differences Estimators of Intertemporal Treatment Effects'. Version 13. arXiv. Available at: https://doi.org/10.48550/arXiv.2007.04267

De Chaisemartin, C. *et al.* (2026) 'Difference-in-Differences Estimators When No Unit Remains Untreated'. Version 6. arXiv. Available at: https://doi.org/10.48550/ARXIV.2405.04465.

Feldstein, M. and Wrobel, M.V. (1998) 'Can state taxes redistribute income?', *Journal of Public Economics*, 68(3), pp. 369–396. Available at: https://doi.org/10.1016/S0047-2727(98)00015-2.

Giesecke, O. and Mateen, H. (2022) 'Local Governments' Response to Fiscal Shocks: Evidence from Connecticut', *SSRN Electronic Journal* [Preprint]. Available at: https://doi.org/10.2139/ssrn.4185090.

Hentze, T., 2023, Kirchensteuereinnahmen sinken kaufkraftbereinigt, Wirtschaftsdienst, 103(5), pp. 352-354

Jonsson, P. and Babajan, T. (2020) 'Effekter av höjda kyrkoavgifter på utträden', *Nyckeln till Svenska kyrkan: en skrift om organisation, verksamhet och ekonomi 2020*, Uppsala: Svenska kyrkan, pp. 72–79. Available at: https://www.svenskakyrkan.se/forskning/nyckeln-till-svenska-kyrkan-2020

Kleven, H. *et al.* (2020) 'Taxation and Migration: Evidence and Policy Implications', *Journal of Economic Perspectives*, 34(2), pp. 119–142. Available at: https://doi.org/10.1257/jep.34.2.119.

Kleven, H.J. *et al.* (2014) 'Migration and Wage Effects of Taxing Top Earners: Evidence from the Foreigners' Tax Scheme in Denmark *', *The Quarterly Journal of Economics*, 129(1), pp. 333–378. Available at: https://doi.org/10.1093/qje/qjt033.

Kleven, H.J., Landais, C. and Saez, E. (2013) 'Taxation and International Migration of Superstars: Evidence from the European Football Market', *American Economic Review*, 103(5), pp. 1892–1924. Available at: https://doi.org/10.1257/aer.103.5.1892.

Kyrkoordningen. Stockholm: Svenska kyrkan. Version as of 1 January 2026. Available at: https://www.svenskakyrkan.se/kyrkoordningen

Lag om Svenska kyrkan (1998) SFS 1998:1591. Available at: https://www.riksdagen.se/sv/dokument-och-lagar/dokument/svensk-forfattningssamling/lag-19981591-om-svenska-kyrkan_sfs-1998-1591/

Liebig, T., Puhani, P.A. and Sousa-Poza, A. (2007) 'Taxation and Internal Migration—Evidence from the Swiss Census Using Community-Level Variation in Income Tax Rates', *Journal of Regional Science*, 47(4), pp. 807–836.
Available at: https://doi.org/10.1111/j.1467-9787.2007.00529.x.

Lundberg, L. (2015) 'Delningen av Göteborgs samfällighet står fast', *Kyrkans Tidning*. Available at: https://www.kyrkanstidning.se/nyhet/delningen-av-goteborgs-samfallighet-star-fast/264844 (Accessed: 10 April 2026).

Lyytikäinen, T. and Santavirta, T. (2013) 'The effect of church tax on church membership', *Journal of Population Economics*, 26(3), pp. 1175–1193. Available at: https://doi.org/10.1007/s00148-012-0431-y.

Metzger, B.M. and Coogan, M.D. (eds) (1994) *Oxford Companion to the Bible*. New York, NY: Oxford University Press.

Moretti, E. and Wilson, D.J. (2017) 'The Effect of State Taxes on the Geographical Location of Top Earners: Evidence from Star Scientists', *American Economic Review*, 107(7), pp. 1858–1903. Available at: https://doi.org/10.1257/aer.20150508.

Naqvi, A. (2024) 'did_multiplegt_dyn (Chaisemartin and D'Haultfœuille 2024)'. Available at: https://asjadnaqvi.github.io/DiD/docs/code_r/07_did_multiplegt_dyn_r/ (Accessed: 27 April 2026).

Quispe, A. *et al.* (2026) 'DIDmultiplegtDYN: Estimation in Difference-in-Difference Designs with Multiple Groups and Periods'. Available at: https://cran.r-universe.dev/DIDmultiplegtDYN/doc/manual.html (Accessed: 27 April 2026).

Riegel, U. *et al.* (2019) 'Does Church Tax Matter?', *International Journal of Practical Theology*, 23(2), pp. 168–187. Available at: https://doi.org/10.1515/ijpt-2018-0028.

SCB (2026a) 'Folkmängden efter ålder och kön. År 1860–2024'. Available at: https://www.statistikdatabasen.scb.se/pxweb/sv/ssd/START__BE__BE0101__BE0101A/BefolkningR1860N/ (Accessed: 10 April 2026).

SCB (2026c) 'Indelningsändringar i Sverige (2015–2023, halvårsvisa rapporter)'. Available at: https://www.scb.se/hitta-statistik/regional-statistik-och-kartor/regionala-indelningar/abonnera-pa-indelningsandringar/ (Accessed: 20 April 2026).

SCB (2026e) 'Län och kommuner 2026 i kodnummerordning'. Available at: https://www.scb.se/hitta-statistik/regional-statistik-och-kartor/regionala-indelningar/lan-och-kommuner/lan-och-kommuner-i-kodnummerordning/ (Accessed: 20 April 2026).

SCB (2026d) 'Local tax rates. Year 2000–2026'. Available at: https://www.statistikdatabasen.scb.se/pxweb/en/ssd/START__OE__OE0101/Kommunalskatter2000/ (Accessed: 20 April 2026).

SCB (2026b) 'Pastorat med koppling till församlingar (2021–2023, XLS datasets)'. Available at: https://www.scb.se/hitta-statistik/regional-statistik-och-kartor/regionala-indelningar/kyrkliga-indelningar/ (Accessed: 20 April 2026).

Skatteverket (2026a) 'Avgift till andra trossamfund'. Available at: https://www.skatteverket.se/privat/skatter/arbeteochinkomst/askattsedelochskattetabeller/avgifttillandratrossamfund.4.18e1b10334ebe8bc80005629.html?q=Avgift+till+andra+trossamfund (Accessed: 8 April 2026).

Skatteverket (2026b) 'De svenska skatternas historia'. Available at: https://www.skatteverket.se/omoss/varverksamhet/skatternasochfolkbokforingenshistoria/desvenskaskatternashistoria.4.22501d9e166a8cb399f1f7d.html (Accessed: 8 April 2026).

Statistics How To (2026) 'Difference in Differences'. Available at: https://www.statisticshowto.com/difference-in-differences/ (Accessed: 8 May 2026).

Svärdhagen, G. (2024) 'Kyrkoavgiftsförändringar och utträden – finns det ett samband?', *Nyckeln till Svenska kyrkan: en skrift om organisation, verksamhet och ekonomi 2024*, chapter 2, Uppsala: Svenska kyrkan, pp. 24–34. Available at: https://www.svenskakyrkan.se/forskning/nyckeln-till-svenska-kyrkan-2024

Svenska kyrkan (2026d) 'Begravningsavgiften'. Available at: https://www.svenskakyrkan.se/begravning/begravningsavgiften (Accessed: 8 April 2026).

Svenska kyrkan (2026b) 'Kyrkoavgiften'. Available at: https://www.svenskakyrkan.se/medlem/kyrkoavgiften (Accessed: 24 April 2026).

Svenska kyrkan (2026a) 'Medlemmar i Svenska kyrkan 1972–2025'. Available at: https://www.svenskakyrkan.se/statistik (Accessed: 8 April 2026).

Svenska kyrkan (2026c) 'Samfälligheten Gotlands kyrkor'. Visby stift. Available at: https://www.svenskakyrkan.se/visbystift/sgk (Accessed: 10 April 2026).

Sverke, J. (2015) 'Kyrkan slår ihop pastorat', *Norrtelje Tidning*. Available at: https://www.norrteljetidning.se/norrtalje/kyrkan-slar-ihop-pastorat/ (Accessed: 10 April 2026).

The Bible. King James Version (1611). Available at: https://www.kingjamesbibleonline.org/Deuteronomy-Chapter-14.

Young, C. *et al.* (2016) 'Millionaire Migration and Taxation of the Elite: Evidence from Administrative Data', *American Sociological Review*, 81(3), pp. 421–446. Available at: https://doi.org/10.1177/0003122416639625.

Yushkov, A. (2025) 'Local Income Taxes and County-Level Migration: New Evidence from Indiana and Maryland', *National Tax Journal*, 78(4), pp. 1017–1036. Available at: https://doi.org/10.1086/738165.

Appendix

Figure 1: Church tax vs exits; Source: Rendering of Church of Sweden data by the author

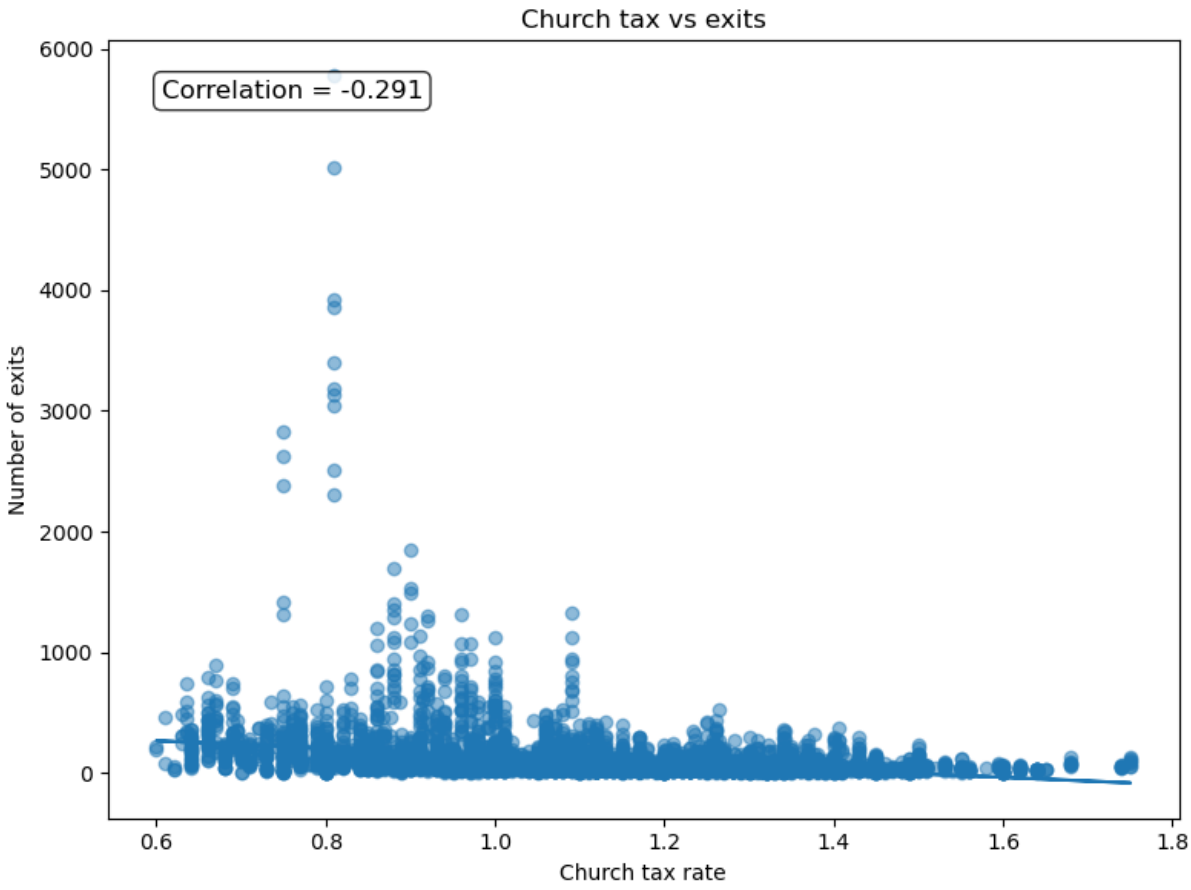

Table 1: Summary Statistics; Source: Rendering of Church of Sweden data by the author

Pastoral-unit-Year Panel, 2014–2023

Variable	Mean	Median	Std. dev.	Min	Max	N
Exits	108.6936	56.0000	215.5378	0.0000	5778.0000	5,820
Members	10081.1765	6176.5000	15426.2982	497.0000	275230.0000	5,820
Exit rate	0.0099	0.0094	0.0039	0.0000	0.0468	5,820
Church tax rate	1.1129	1.1200	0.2070	0.6000	1.7500	5,820
Tax change	0.0538	0.0500	0.0658	-0.1800	0.3750	539
Tax increase	0.0688	0.0500	0.0568	0.0050	0.3750	463
Tax decrease	-0.0377	-0.0223	0.0362	-0.1800	-0.0058	76

Note: Summary statistics are calculated with pastoral-unit-year observations for 2014–2023. A treatment threshold of 0.005 is enforced, only above which changes are classified as treatments. Exit rate is defined as exits divided by members.

Table 2: Diocese-level Summary Statistics; Source: Rendering of Church of Sweden data and Statistics Sweden (SCB, 2026b) data by the author

Averages across pastoral units within each diocese, 2014–2023

Stift code	Stift	Mean members	Mean tax rate	Mean exits	Mean exit rate
1	Uppsala stift	9,701	1.121	96	0.0092
2	Linköpings stift	11,850	1.237	130	0.0100
3	Skara stift	8,985	1.210	99	0.0102
4	Strängnäs stift	7,335	1.107	78	0.0102
5	Västerås stift	9,579	1.160	102	0.0108
6	Växjö stift	7,385	1.199	73	0.0091
7	Lunds stift	8,051	1.086	86	0.0095
8	Göteborgs stift	12,881	1.108	160	0.0112
9	Karlstads stift	8,438	1.152	74	0.0081
10	Härnösands stift	10,182	1.249	113	0.0105
11	Luleå stift	10,897	1.042	111	0.0089
12	Visby stift	3,831	1.487	51	0.0106
13	Stockholms stift	18,211	0.794	196	0.0106

Note: Exit rate is defined as exits divided by members.

Figure 2: Distribution of changes in church tax; Source: Rendering of Church of Sweden data by the author

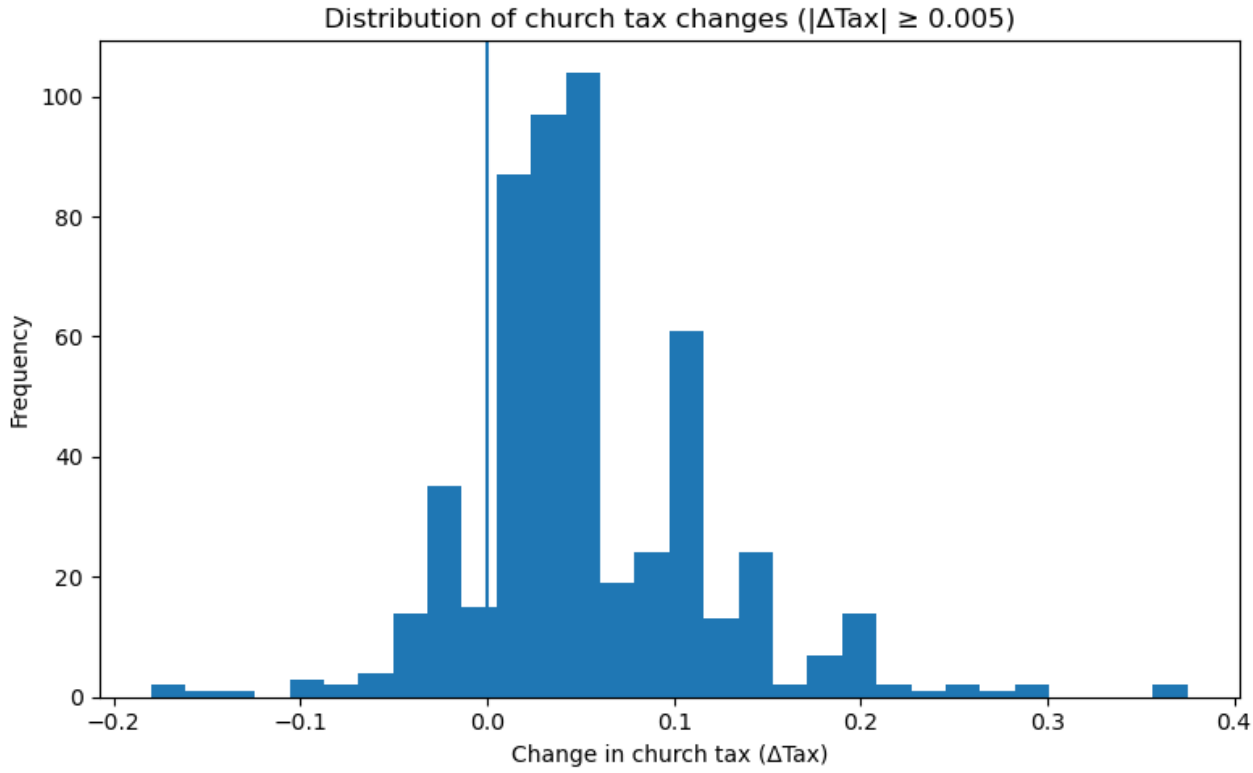

Figure 3: Members vs Exits; Source: Rendering of Church of Sweden data by the author

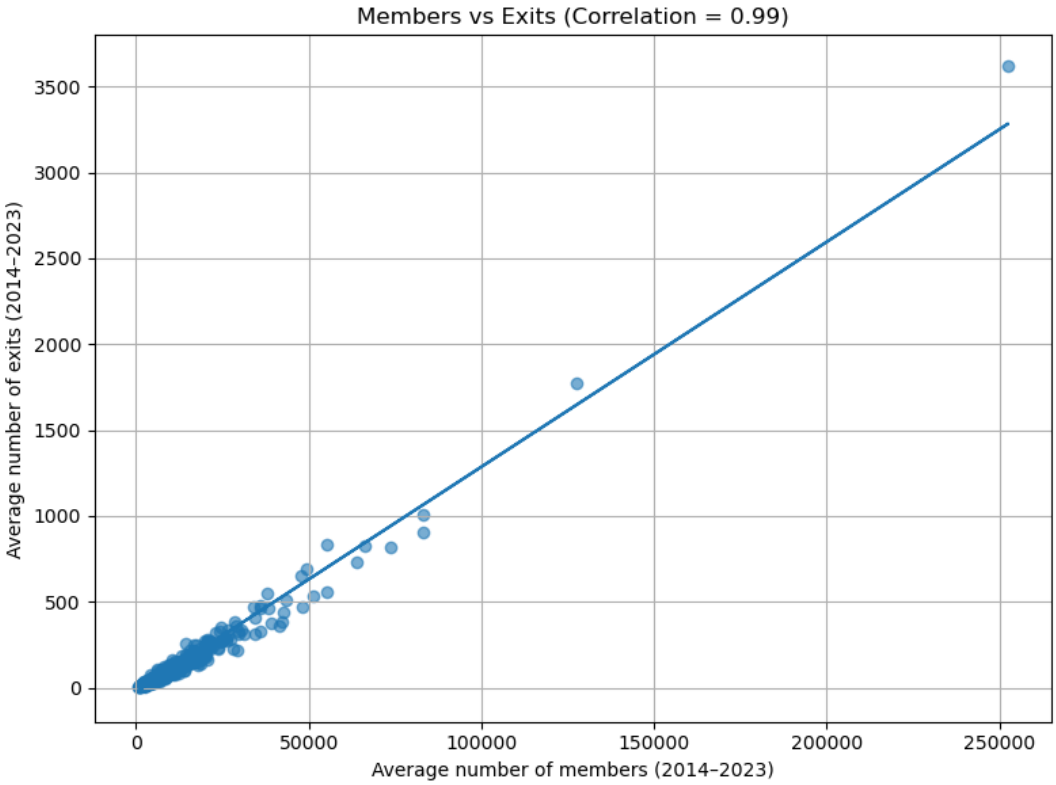

Figure 4: Church tax vs exit rate; Source: Rendering of Church of Sweden data by the author

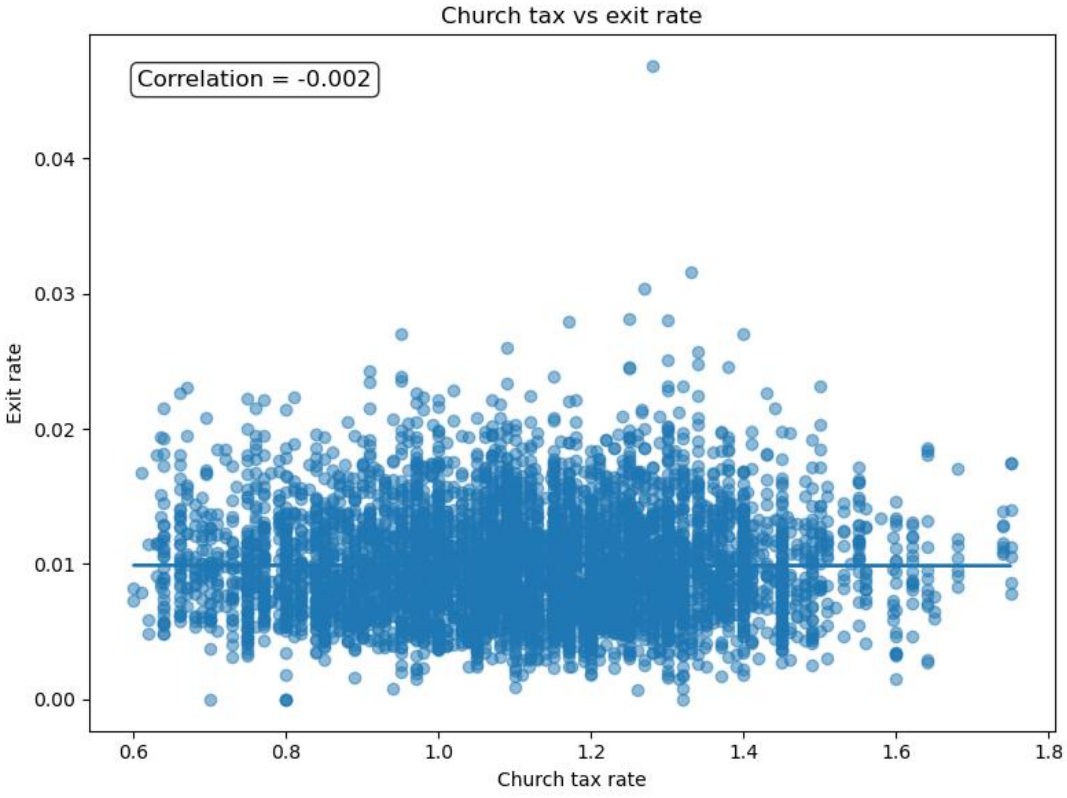

Figure 5: Church tax rate versus the number of members; Source: Rendering of Church of Sweden data by the author

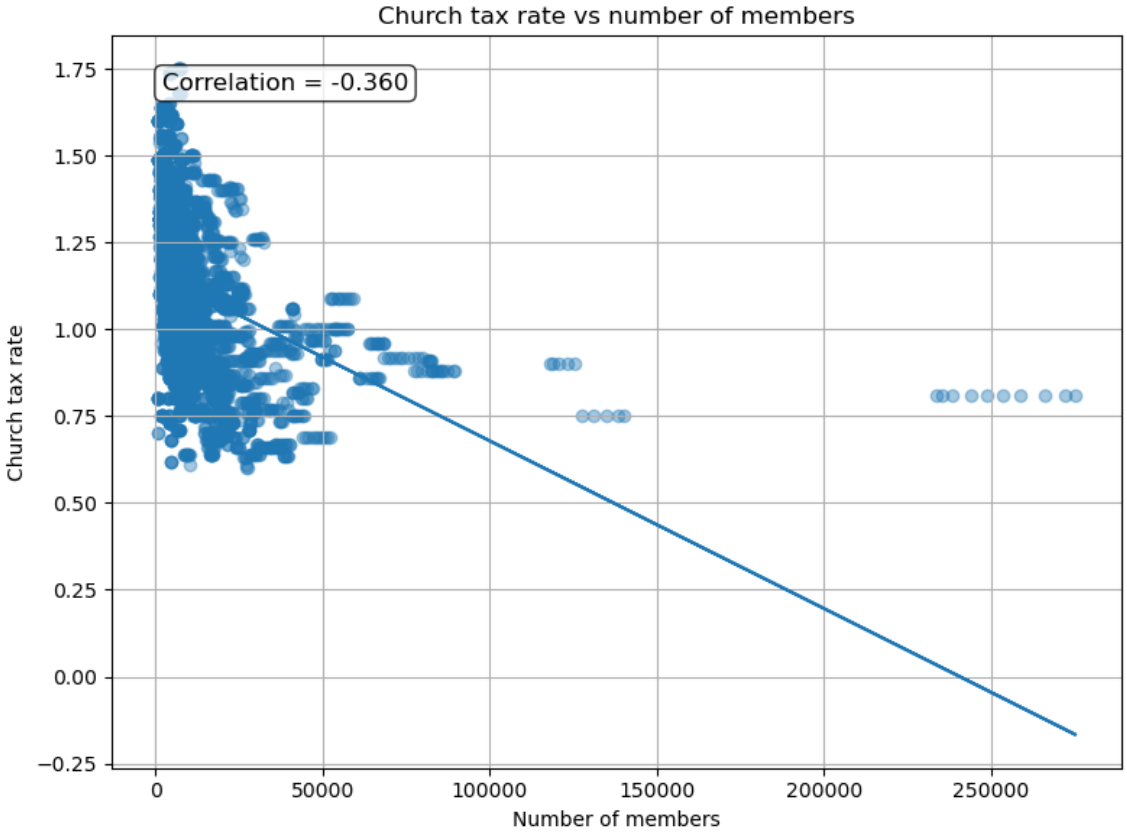

Figure 6: Change in church tax versus the number of exits; Source: Rendering of Church of Sweden data by the author

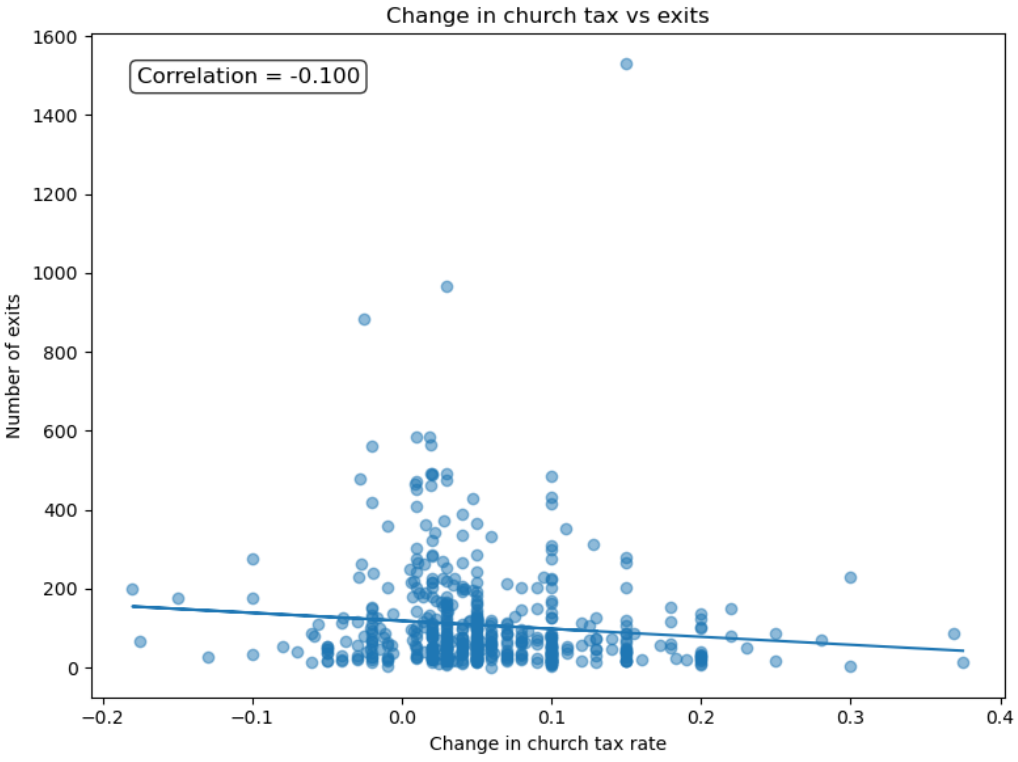

Figure 7: Change in church tax versus the exit rate; Source: Rendering of Church of Sweden data by the author

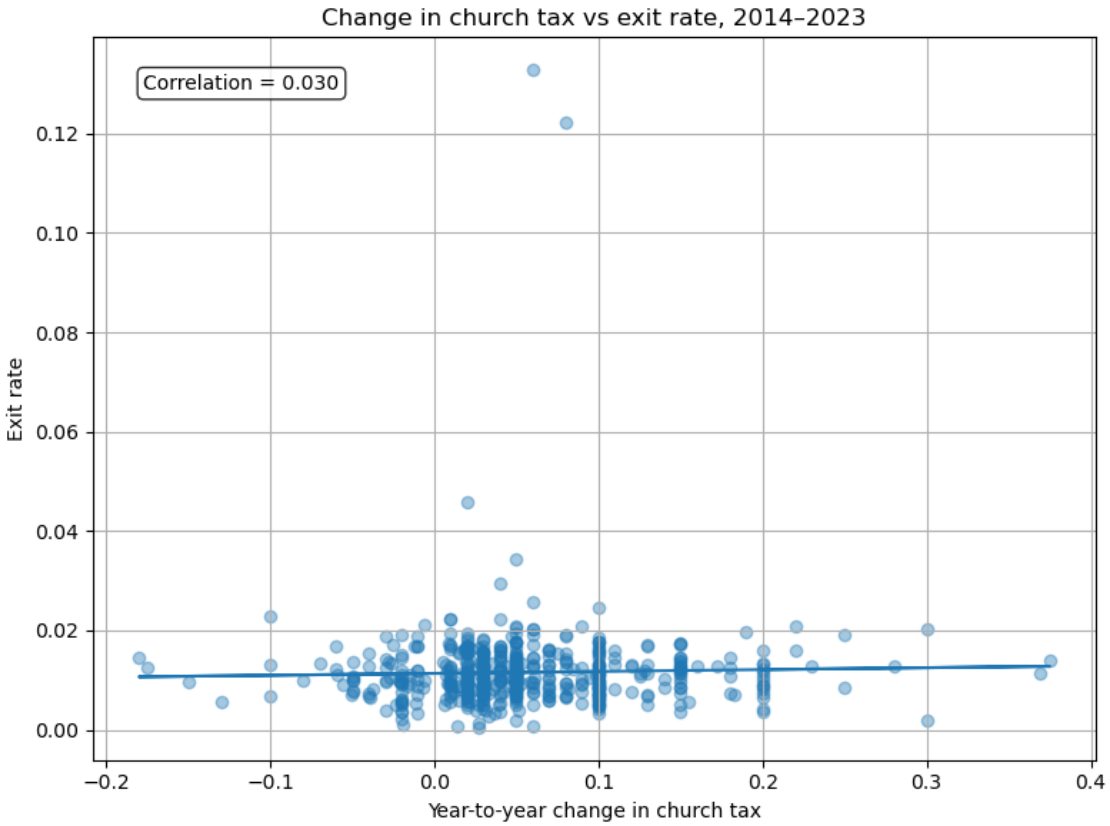

Figure 8: Total number of exits from the Church of Sweden during 2014-2023; Source: Rendering of Church of Sweden data by the author

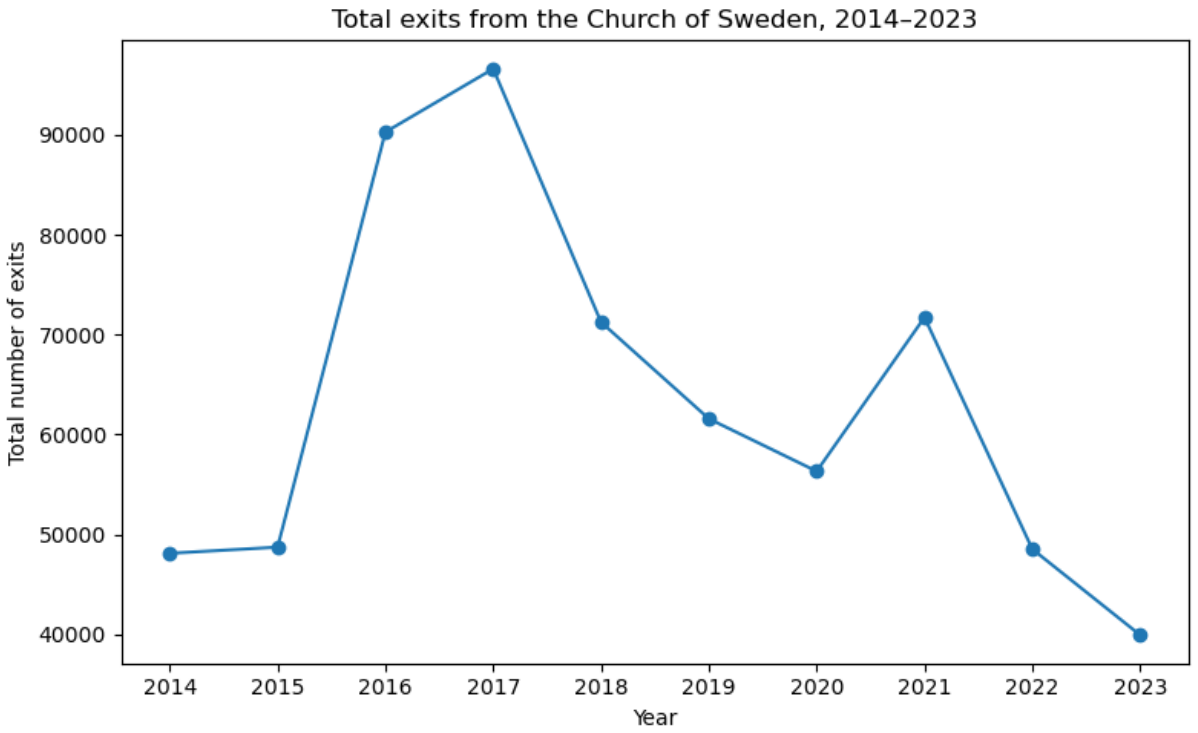

Figure 9: The exit rate from the Church of Sweden during 2014-2023; Source: Rendering of Church of Sweden data by the author

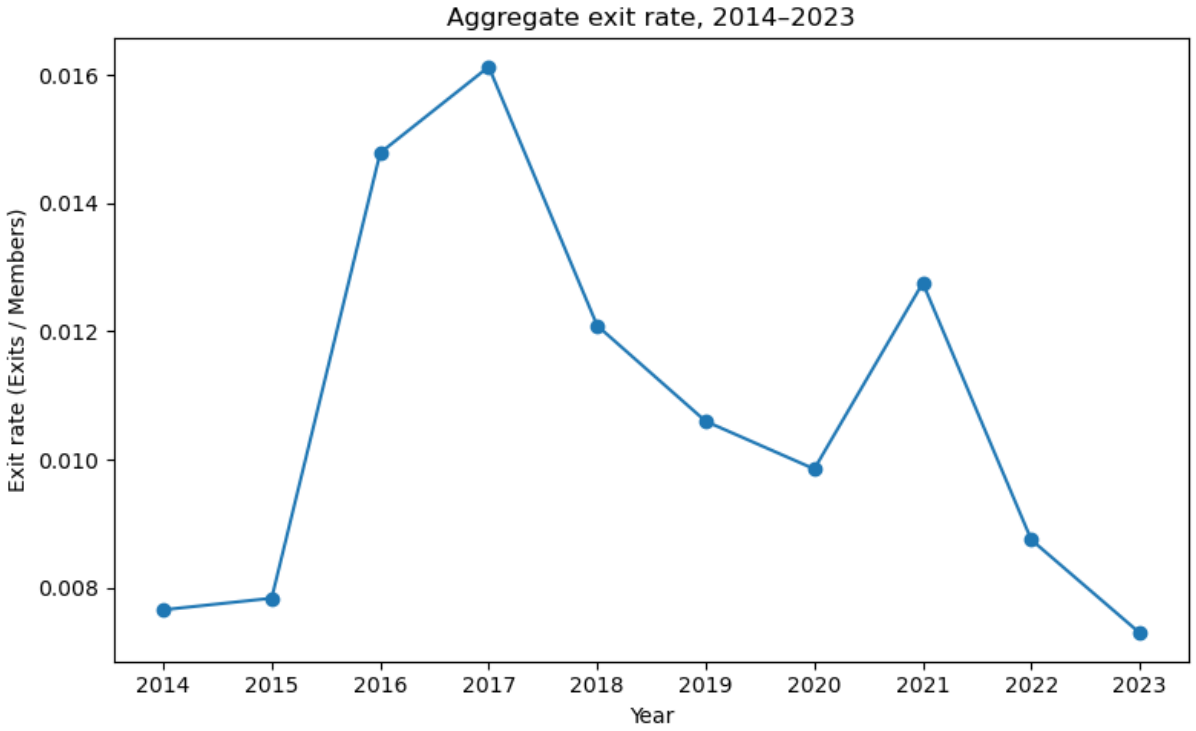

Figure 10: Number of church tax increases; Source: Rendering of Church of Sweden data by the author

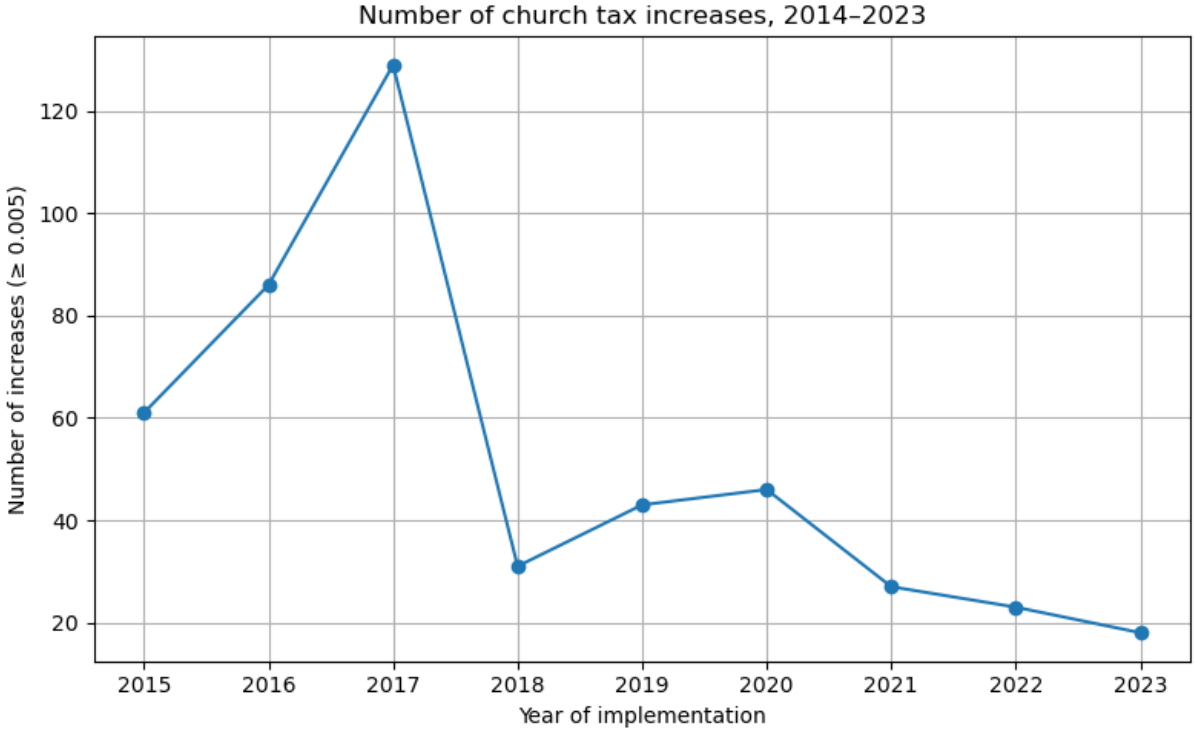

Figure 11: Diocesan church tax changes; Source: Rendering of Church of Sweden data by the author

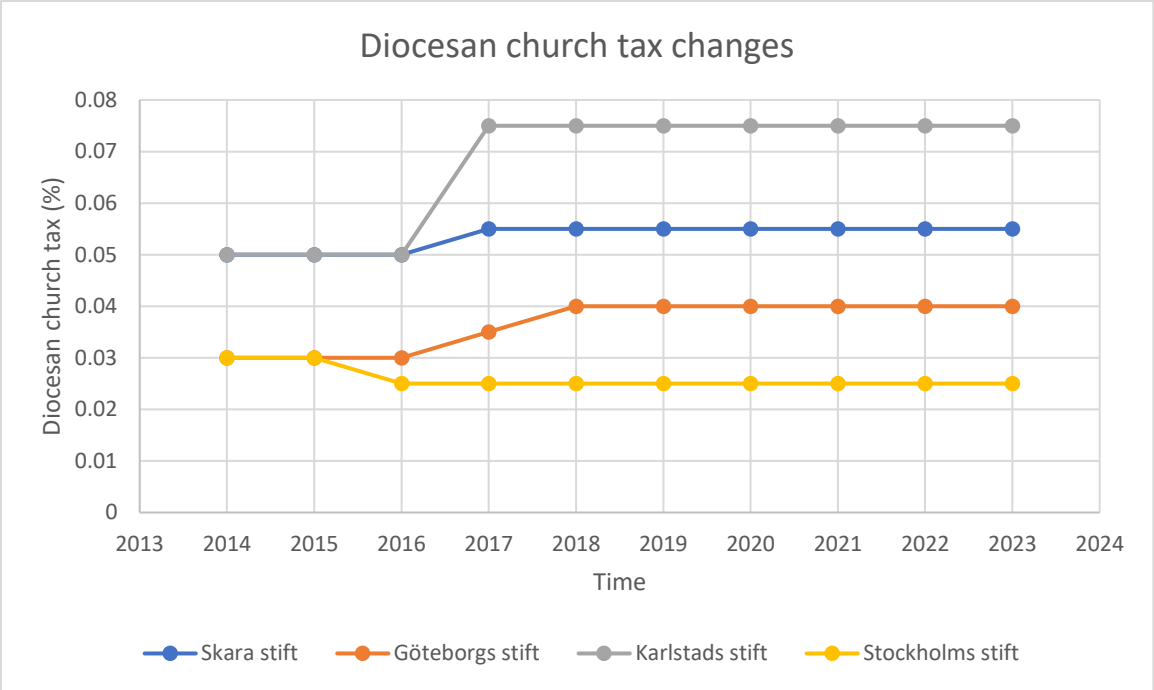

Figure 12: Municipal and church tax; Source: Rendering of Church of Sweden data and Statistics Sweden data (SCB, 2026d and SCB, 2026e) by the author

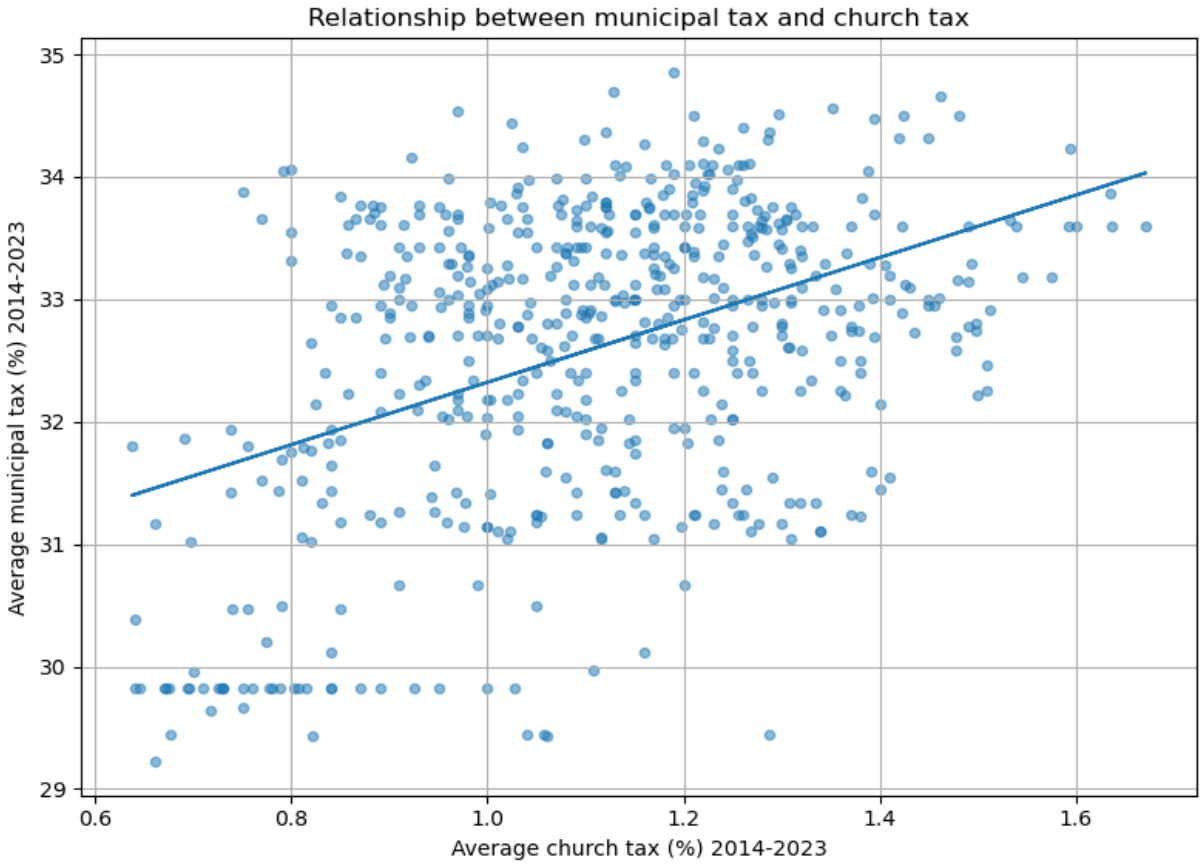

Table 3: Event-Study Model Summary; Source: Rendering of Church of Sweden data by the author

Effect of Church Tax Increases on Exits

Statistic	*Value*
Observations	4330
Pastoral units in regression	433
Treated pastoral units	185
Control pastoral units	248
Dropped pastoral units, total	149
Dropped due to decreases	57
Dropped due to multiple increases	92
R-squared	0.9199
Adjusted R-squared	0.9104
Joint pre-treatment p-value	0.5736
Pastoral unit FE	Yes
Year FE	Yes
Clustered SE	Pastoral unit
Reference period	k = 0

Note: The reference year k = 0 corresponds to the first year, in which the change in church tax started to be in effect.

Table 4: Event-Study Coefficients; Source: Rendering of Church of Sweden data by the author

Effect of Church Tax Increase on Exits

Event time (k)	*Coefficient*	*Standard Error*	*95% CI*
-9	15.9535	(8.9828)	[-1.6525, 33.5595]
-8	18.9872*	(7.7032)	[3.8891, 34.0852]
-7	6.3148	(8.4792)	[-10.3041, 22.9337]
-6	1.0768	(7.2429)	[-13.1191, 15.2727]
-5	-8.6073	(5.3900)	[-19.1715, 1.9570]
-4	-9.4085	(7.6862)	[-24.4732, 5.6563]
-3	17.5264	(9.3851)	[-0.8681, 35.9209]
-2	21.0380*	(10.1713)	[1.1026, 40.9734]
-1	14.4582	(8.6233)	[-2.4432, 31.3596]
1	7.2664*	(3.4073)	[0.5883, 13.9445]
2	16.0523***	(4.8038)	[6.6370, 25.4677]
3	13.3951*	(5.7692)	[2.0877, 24.7026]
4	10.8189	(6.5999)	[-2.1166, 23.7544]
5	21.0414**	(7.8746)	[5.6075, 36.4752]
6	27.9073**	(9.8148)	[8.6706, 47.1439]
7	25.1566*	(12.1472)	[1.3485, 48.9646]
8	26.4572*	(12.9839)	[1.0091, 51.9052]

Note: Coefficients are relative to reference period k = 0. Standard errors are clustered at the pastoral unit level. Confidence intervals are 95%. Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

Figure 13: Replication coefficients plot; Source: Rendering of Church of Sweden data by the author

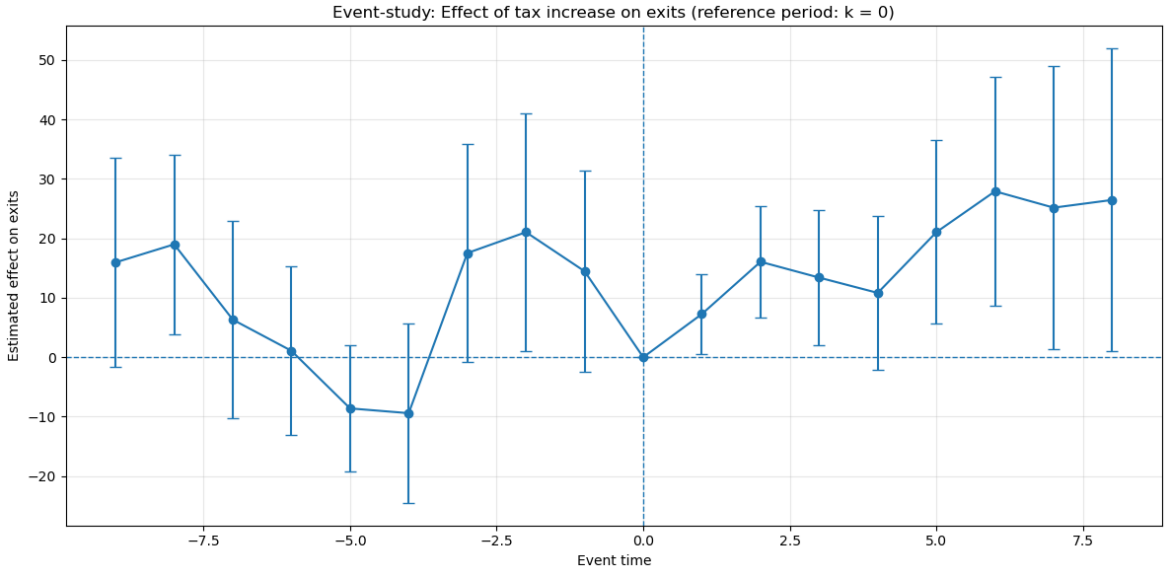

Table 5: Event-Study Coefficients; Source: Rendering of Church of Sweden data by the author

Effect of Church Tax Increases on Exits

Statistic	*Value*
Observations	5250
Pastoral units in regression	525
Treated pastoral units	277
Control pastoral units	248
Dropped pastoral units, total	57
Dropped due to decreases	57
R-squared	0.9224
Adjusted R-squared	0.9112
Joint pre-treatment p-value	0.2829
Pastoral unit FE	Yes
Diocese-year FE	Yes
Clustered SE	Pastoral unit
Reference period	k = -1

Note: The reference year k = -1 corresponds to the last year before the church tax change was implemented.

Table 6: Event-Study Coefficients; Source: Rendering of Church of Sweden data by the author

Effect of First Church Tax Increases on Exits

Event time (k)	*Coefficient*	*Standard Error*	*95% CI*
-9	5.1498	(5.0467)	[-4.7414, 15.0411]
-8	8.4906	(5.6466)	[-2.5766, 19.5579]
-7	-4.3599	(7.0535)	[-18.1845, 9.4648]
-6	-9.0804	(6.6826)	[-22.1780, 4.0172]
-5	-18.9443	(10.1058)	[-38.7514, 0.8628]
-4	-20.2746	(12.0103)	[-43.8143, 3.2652]
-3	3.8912	(5.1156)	[-6.1352, 13.9176]
-2	6.8016	(4.2105)	[-1.4509, 15.0541]
0	-10.5033	(6.2110)	[-22.6767, 1.6700]
1	-3.3549	(7.7460)	[-18.5368, 11.8270]
2	2.2279	(5.9361)	[-9.4065, 13.8624]
3	2.1122	(7.7518)	[-13.0811, 17.3054]
4	0.6055	(8.7276)	[-16.5002, 17.7113]
5	8.8444	(6.2782)	[-3.4607, 21.1495]
6	13.1308	(6.8553)	[-0.3053, 26.5670]
7	11.8400	(9.4304)	[-6.6432, 30.3231]
8	10.6670	(12.5123)	[-13.8567, 35.1907]

Note: Coefficients are relative to reference period k = -1. Standard errors are clustered at the pastoral unit level. Confidence intervals are 95%. Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

Figure 14: First improvement of TWFE analysis; Source: Rendering of Church of Sweden data by the author

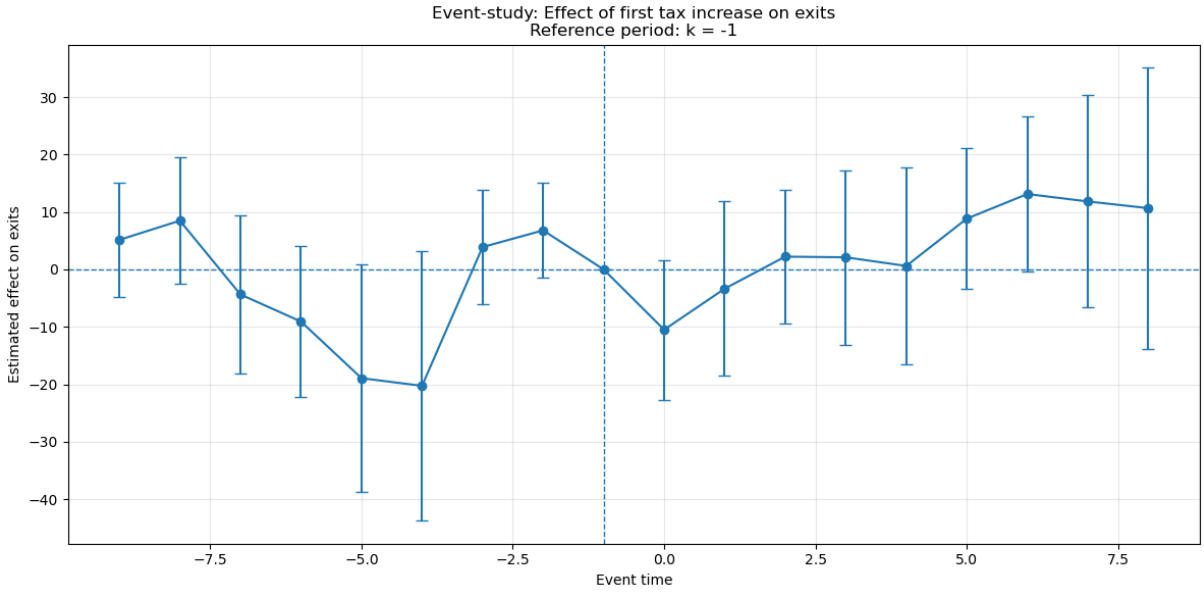

Table 7: Event-Study Summary: Tax Increases; Source: Rendering of Church of Sweden data by the author

Statistic	*Value*
Observations	5820
Pastoral units in regression	582
Treated pastoral units	312
Control pastoral units	270
R-squared	0.9217
Adjusted R-squared	0.9107
Joint pre-treatment p-value	0.1626
Pastorat FE	Yes
Stift × Year FE	Yes
Clustered SE	Pastorat
Reference period	k = -1

Table 8: Event-Study Summary: Tax Decreases; Source: Rendering of Church of Sweden data by the author

Statistic	*Value*
Observations	5820
Pastoral units in regression	582
Treated pastoral units	57
Control pastoral units	525
R-squared	0.9215
Adjusted R-squared	0.9105
Joint pre-treatment p-value	0.0000
Pastorat FE	Yes
Stift × Year FE	Yes
Clustered SE	Pastorat
Reference period	k = -1

Table 9: Event-Study (Tax Increases); Source: Rendering of Church of Sweden data by the author

Event time (k)	*Coefficient*	*Standard Error*	*95% CI*
-9	6.7568	(4.6556)	[-2.3680, 15.8816]
-8	9.9028	(5.1833)	[-0.2563, 20.0619]
-7	-3.8755	(6.6645)	[-16.9377, 9.1867]
-6	-7.5112	(5.9890)	[-19.2494, 4.2270]
-5	-16.4946	(9.2002)	[-34.5267, 1.5376]
-4	-17.5316	(10.9546)	[-39.0022, 3.9391]
-3	4.3031	(4.6655)	[-4.8412, 13.4474]
-2	7.5158	(3.8628)	[-0.0551, 15.0867]
0	-9.5120	(5.5893)	[-20.4669, 1.4429]
1	-1.6210	(6.9546)	[-15.2517, 12.0097]
2	3.7482	(5.3891)	[-6.8142, 14.3107]
3	3.4171	(7.0846)	[-10.4685, 17.3027]
4	2.1706	(7.8602)	[-13.2352, 17.5763]
5	8.7056	(5.9185)	[-2.8944, 20.3056]
6	12.9165*	(6.4308)	[0.3124, 25.5206]
7	10.5783	(8.2885)	[-5.6668, 26.8234]
8	8.2231	(10.7778)	[-12.9009, 29.3472]

Note: Coefficients are relative to k = -1. Standard errors clustered at the pastoral unit level. COnfidence intervals are 95%. Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

Table 10: Event-Study (Tax Decreases); Source: Rendering of Church of Sweden data by the author

Event time (k)	*Coefficient*	*Standard Error*	*95% CI*
-9	-48.5575***	(8.1194)	[-64.4713, -32.6437]
-8	-43.8003*	(19.9337)	[-82.8696, -4.7311]
-7	-33.8076	(20.8458)	[-74.6646, 7.0495]
-6	18.0024	(10.5439)	[-2.6632, 38.6680]
-5	2.9172	(10.9895)	[-18.6218, 24.4561]
-4	-19.9273	(11.3417)	[-42.1566, 2.3020]
-3	-12.3430	(14.2190)	[-40.2116, 15.5257]
-2	-10.4614	(11.0272)	[-32.0743, 11.1515]
0	-7.1100	(4.0534)	[-15.0544, 0.8345]
1	-26.5511***	(7.6859)	[-41.6152, -11.4869]
2	-26.0017**	(9.4177)	[-44.4602, -7.5433]
3	-18.9031*	(8.7329)	[-36.0192, -1.7870]
4	-12.5995	(8.2073)	[-28.6856, 3.4865]
5	-19.4448	(12.8786)	[-44.6865, 5.7969]
6	-20.2587	(11.9462)	[-43.6728, 3.1555]
7	3.7665	(9.0789)	[-14.0278, 21.5607]
8	11.2380	(9.3710)	[-7.1288, 29.6049]

Note: Coefficients are relative to k = -1. Standard errors clustered at the pastoral unit level. COnfidence intervals are 95%. Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

Figure 15: Increases only; Source: Rendering of Church of Sweden data by the author

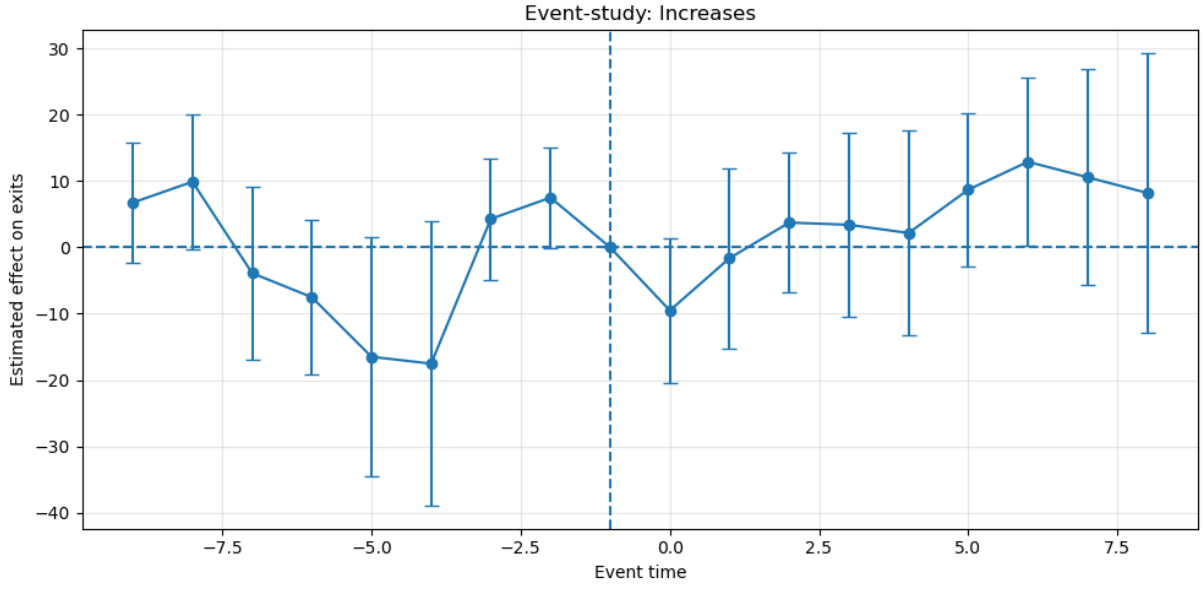

Figure 16: Decreases only; Source: Rendering of Church of Sweden data by the author

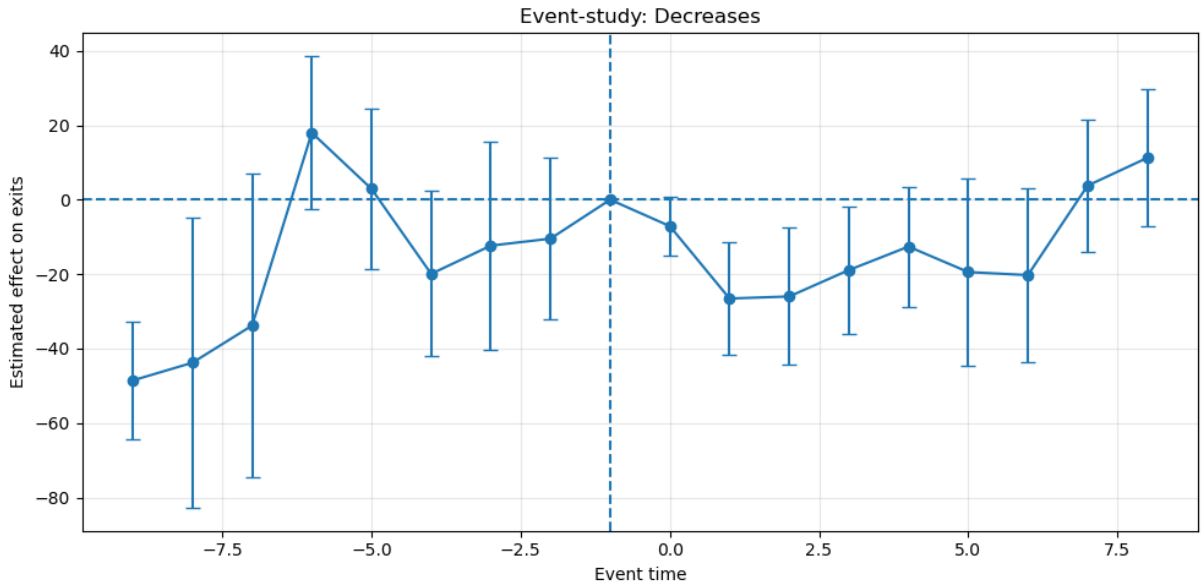

Table 11: Placebo Tests; Source: Rendering of Church of Sweden data by the author

Placebo	*Coefficient*	*Standard Error*	*p-value*	*95% CI*
1 year early	-3.0535	(3.9587)	0.4405	[-10.8125, 4.7054]
2 years early	4.6522	(3.2587)	0.1534	[-1.7347, 11.0392]

Note: Standard errors clustered at pastoral unit level. Significance levels:* p<0.05, ** p<0.01, *** p<0.001.

Table 12: Event-Study Summary; Source: Rendering of Church of Sweden data by the author

Only Larger Tax Increases: Tax increase ≥ 0.025

Statistic	*Value*
Observations	5250
Pastoral units in regression	525
Treated pastoral units	250
Control pastoral units	275
R-squared	0.9225
Adjusted R-squared	0.9113
Joint pre-treatment p-value	0.4142
Pastorat FE	Yes
Stift × Year FE	Yes
Clustered SE	Pastorat
Reference period	k = -1
Treatment threshold	Tax increase ≥ 0.025

Table 13: Event-Study Coefficients; Source: Rendering of Church of Sweden data by the author

Only Larger Tax Increases: Tax increase ≥ 0.025

Event time (k)	*Coefficient*	*Standard Error*	*p-value*	*95% CI*
-9	-0.3336	(6.2661)	0.9575	[-12.6150, 11.9478]
-8	1.4117	(6.8692)	0.8372	[-12.0516, 14.8750]
-7	-9.0636	(7.2649)	0.2122	[-23.3025, 5.1754]
-6	-9.8978	(6.5361)	0.1299	[-22.7083, 2.9127]
-5	-17.3812	(10.2134)	0.0888	[-37.3991, 2.6367]
-4	-22.1823	(12.3214)	0.0718	[-46.3319, 1.9673]
-3	2.4005	(5.1712)	0.6425	[-7.7348, 12.5358]
-2	6.7046	(4.4585)	0.1326	[-2.0339, 15.4432]
0	-12.3360	(6.7273)	0.0667	[-25.5212, 0.8492]
1	-6.8736	(7.9994)	0.3902	[-22.5522, 8.8050]
2	2.1284	(6.0507)	0.7250	[-9.7307, 13.9875]
3	1.5815	(8.3918)	0.8505	[-14.8661, 18.0291]
4	-1.0936	(9.6336)	0.9096	[-19.9751, 17.7878]
5	8.8257	(6.4718)	0.1727	[-3.8589, 21.5103]
6	14.6457*	(7.0396)	0.0375	[0.8484, 28.4431]
7	18.3574*	(9.2971)	0.0483	[0.1355, 36.5794]
8	17.8337	(12.4651)	0.1525	[-6.5974, 42.2648]

Note: Coefficients are relative to k = -1. Standard errors clustered at the pastoral unit level. Confidence intervals are 95%. Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

Figure 17: Increased treatment threshold; Source: Rendering of Church of Sweden data by the author

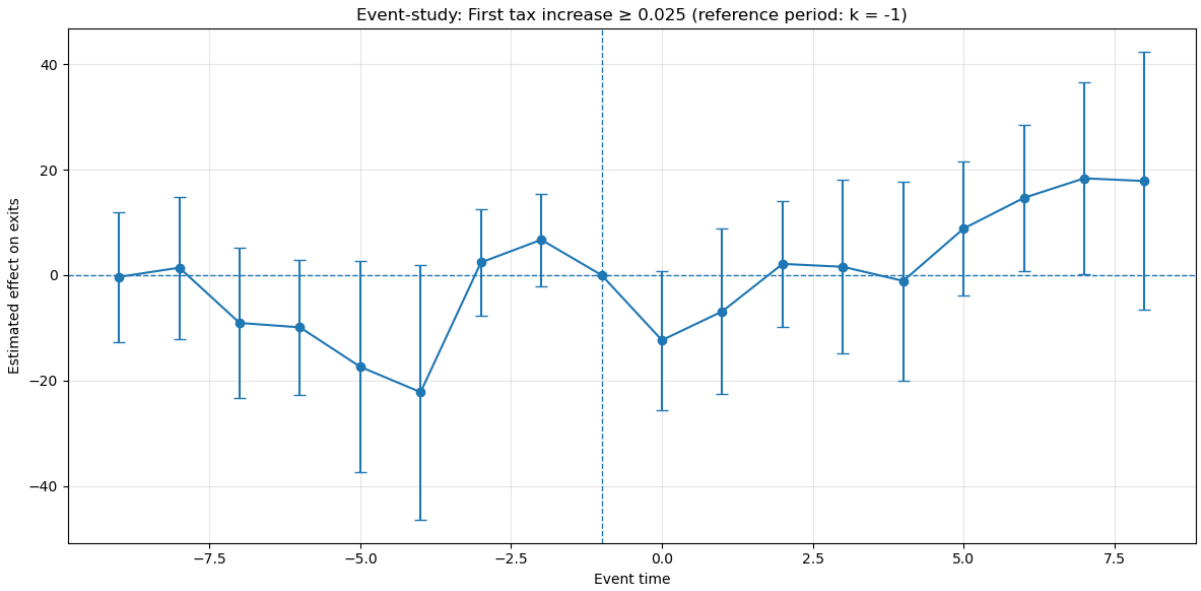

Table 14: Event-Study Summary; Source: Rendering of Church of Sweden data by the author

Robustness: Exit Rate as Outcome

Statistic	*Value*
Observations	5250
Pastoral units in regression	525
Treated pastoral units	277
Control pastoral units	248
R-squared	0.7660
Adjusted R-squared	0.7324
Joint pre-treatment p-value	0.4974
Pastorat FE	Yes
Stift × Year FE	Yes
Clustered SE	Pastoral unnit
Reference period	k = -1
Outcome	Exits per 1000 members
Treatment threshold	Tax increase ≥ 0.005

Table 15: Event-Study Coefficients; Source: Rendering of Church of Sweden data by the author

Robustness: Exit Rate as Outcome

Event time (k)	*Coefficient*	*Standard Error*	*p-value*	*95% CI*
-9	-0.8912	(0.5627)	0.1132	[-1.9940, 0.2116]
-8	-0.3924	(0.6715)	0.5590	[-1.7084, 0.9237]
-7	-0.5426	(0.4198)	0.1962	[-1.3654, 0.2802]
-6	-0.5150	(0.3492)	0.1402	[-1.1993, 0.1694]
-5	-0.3278	(0.3410)	0.3364	[-0.9961, 0.3405]
-4	-0.5862*	(0.2920)	0.0447	[-1.1586, -0.0138]
-3	-0.4088	(0.3095)	0.1865	[-1.0155, 0.1978]
-2	-0.2926	(0.2350)	0.2130	[-0.7531, 0.1679]
0	-0.4014*	(0.1845)	0.0295	[-0.7630, -0.0399]
1	0.1107	(0.2381)	0.6419	[-0.3559, 0.5773]
2	0.1698	(0.2485)	0.4945	[-0.3172, 0.6568]
3	0.3252	(0.2640)	0.2179	[-0.1921, 0.8426]
4	0.6403*	(0.2797)	0.0221	[0.0921, 1.1886]
5	0.4480	(0.3007)	0.1362	[-0.1413, 1.0373]
6	0.1898	(0.2961)	0.5215	[-0.3906, 0.7703]
7	0.3625	(0.3378)	0.2832	[-0.2995, 1.0246]
8	0.4856	(0.3991)	0.2237	[-0.2967, 1.2679]

Note: Coefficients are relative to k = -1. Standard errors clustered at the pastoral unit level. Confidence intervals are 95%. Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

Figure 18: Effect on exit rate; Source: Rendering of Church of Sweden data by the author

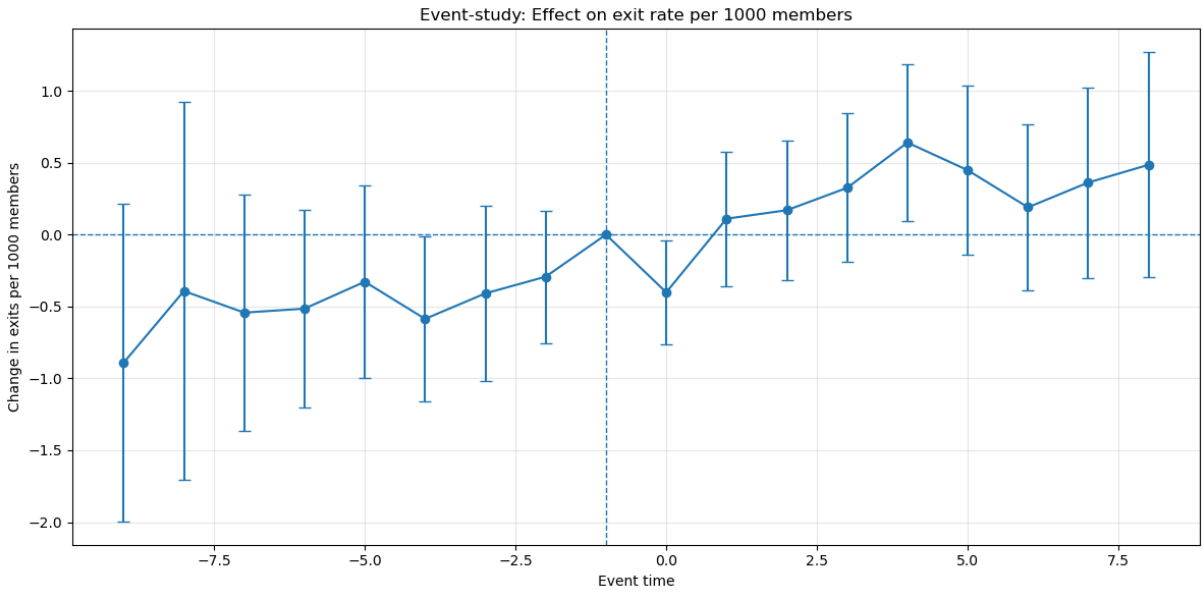

Table 16: Main baseline model summary; Source: Rendering of Church of Sweden data by the author

Statistic	Value
Number of effects	8
Number of placebos	4
Switchers in first effect	228
Joint p-value, effects	0.0602
Equality p-value, effects	0.0736
Joint p-value, placebos	0.2337
Switchers type	Switchers-in
Clustered SE	Pastoral unit

Table 17: Main model event-study coefficients; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Effect_1	-13.4089	(7.1874)	0.0621	[-27.4960, 0.6782]	1,061	228
Effect_2	-14.3074	(9.5658)	0.1347	[-33.0561, 4.4413]	950	220
Effect_3	-2.6961	(7.0584)	0.7025	[-16.5304, 11.1382]	862	213
Effect_4	3.7399	(9.1896)	0.6840	[-14.2714, 21.7513]	787	206
Effect_5	0.4622	(9.4972)	0.9612	[-18.1519, 19.0763]	658	180
Effect_6	7.9375	(5.8183)	0.1725	[-3.4661, 19.3411]	562	162
Effect_7	14.1435*	(7.0201)	0.0439	[0.3843, 27.9027]	525	153
Effect_8	8.2422	(5.7946)	0.1549	[-3.1151, 19.5994]	313	87

Table 18: Main model placebo effects; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Placebo_1	-0.2728	(4.7577)	0.9543	[-9.5976, 9.0521]	830	191
Placebo_2	-7.5984	(8.2314)	0.3560	[-23.7316, 8.5347]	522	131
Placebo_3	-40.6561	(26.5255)	0.1253	[-92.6452, 11.3329]	251	58
Placebo_4	-18.9481	(29.7918)	0.5248	[-77.3391, 39.4428]	182	43

Table 19: Main model average cumulative effect; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Average cumulative effect	-13.4657	(66.9810)	0.8407	[-144.7460, 117.8145]	3,591	1,449

Figure 19: Main model event-study coefficients plot; Source: Rendering of Church of Sweden data by the author

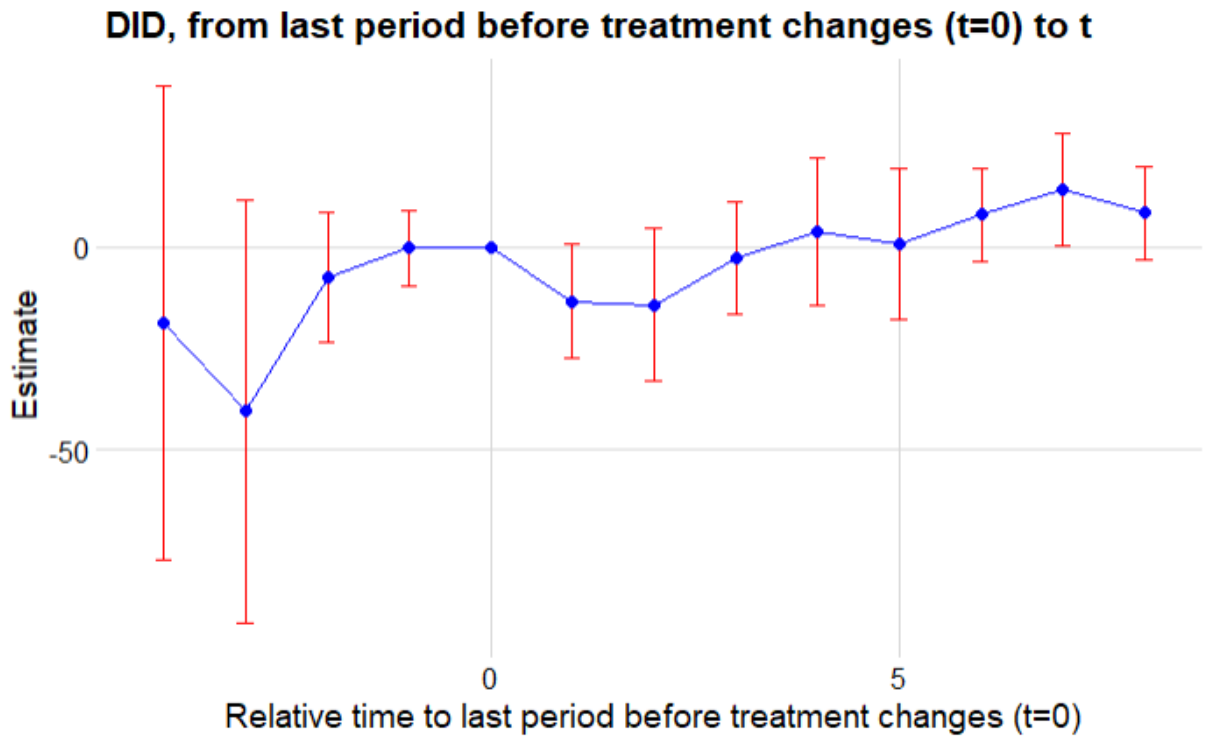

Table 20: Normalised effect model summary; Source: Rendering of Church of Sweden data by the author

Statistic	Value
Specification	Normalized effects per tax unit
Number of effects	8
Number of placebos	4
Switchers in first effect	228
Joint p-value, effects	0.0602
Equality p-value, effects	NA
Joint p-value, placebos	0.2337
Switchers type	Switchers-in
Normalized	Yes
Normalized weights	Yes
Clustered SE	Pastoral unit

Table 21: Normalised effect event-study coefficients; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Effect_1	-174.7643	(93.6769)	0.0621	[-358.3677, 8.8391]	1,061	228
Effect_2	-91.1643	(60.9519)	0.1347	[-210.6278, 28.2993]	950	220
Effect_3	-11.0129	(28.8318)	0.7025	[-67.5221, 45.4964]	862	213
Effect_4	11.2530	(27.6506)	0.6840	[-42.9412, 65.4471]	787	206
Effect_5	1.1015	(22.6340)	0.9612	[-43.2604, 45.4634]	658	180
Effect_6	15.1065	(11.0732)	0.1725	[-6.5966, 36.8096]	562	162
Effect_7	22.1537*	(10.9960)	0.0439	[0.6020, 43.7054]	525	153
Effect_8	12.0193	(8.4501)	0.1549	[-4.5426, 28.5811]	313	87

Table 22: Normalised effect placebo coefficients; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Placebo_1	-3.3603	(58.6138)	0.9543	[-118.2412, 111.5206]	830	191
Placebo_2	-43.2274	(46.8281)	0.3560	[-135.0087, 48.5539]	522	131
Placebo_3	-163.8676	(106.9131)	0.1253	[-373.4134, 45.6781]	251	58
Placebo_4	-57.3781	(90.2147)	0.5248	[-234.1958, 119.4395]	182	43

Table 23: Normalised effect average cumulative effect; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Average cumulative effect	-13.4657	(66.9810)	0.8407	[-144.7460, 117.8145]	3,591	1,449

Figure 20: Normalised event-study coefficients plot; Source: Rendering of Church of Sweden data by the author

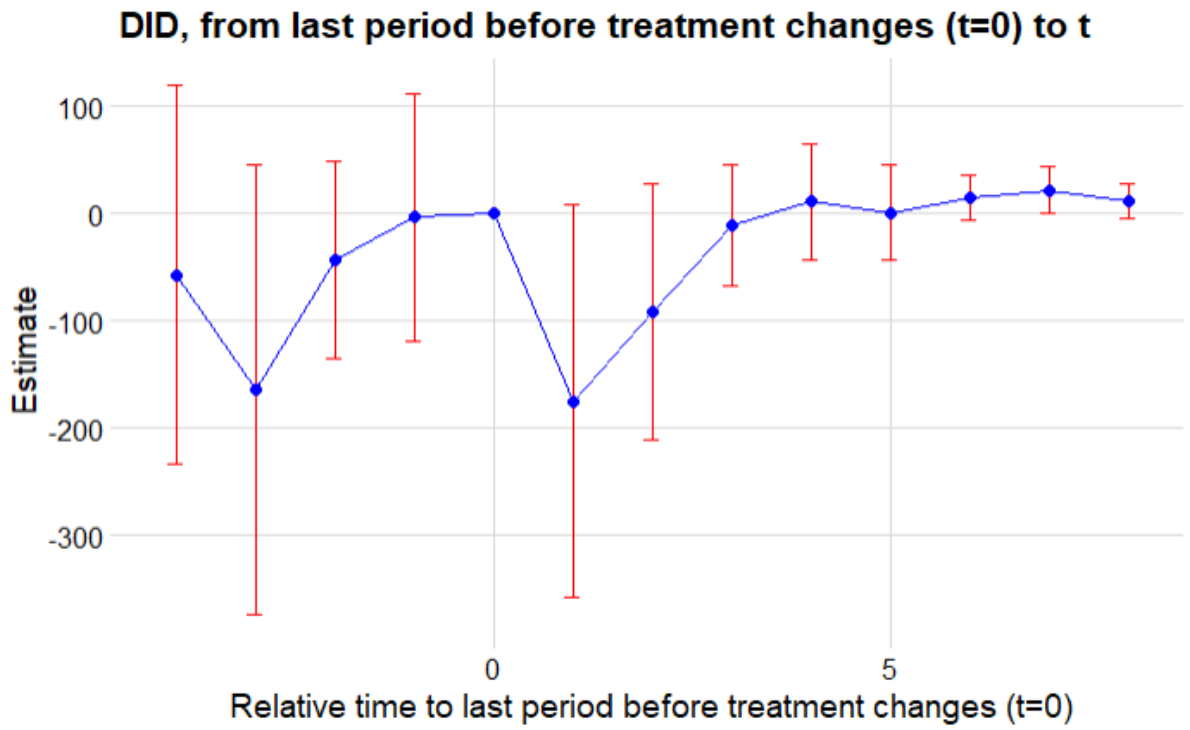

Table 24: Decreases in church tax model summary; Source: Rendering of Church of Sweden data by the author

Statistic	Value
Specification	Tax decreases (switchers-out)
Number of effects	8
Number of placebos	4
Switchers in first effect	23
Joint p-value, effects	0.8136
Equality p-value, effects	0.7276
Joint p-value, placebos	0.4376
Switchers type	Switchers-out
Clustered SE	Pastoral unit

Table 25: Decreases in church tax event-study coefficients; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Effect_1	4.1619	(8.6947)	0.6322	[-12.8794, 21.2032]	195	23
Effect_2	5.3046	(16.4549)	0.7472	[-26.9464, 37.5556]	155	20
Effect_3	6.7553	(18.6317)	0.7169	[-29.7621, 43.2727]	132	18
Effect_4	3.0996	(13.4032)	0.8171	[-23.1701, 29.3693]	127	18
Effect_5	0.2991	(19.0493)	0.9875	[-37.0368, 37.6350]	120	18
Effect_6	4.3438	(24.8763)	0.8614	[-44.4130, 53.1005]	110	17
Effect_7	6.1853	(17.8527)	0.7290	[-28.8054, 41.1760]	97	15
Effect_8	-5.1462	(5.9809)	0.3895	[-16.8685, 6.5761]	64	10

Table 26: Decreases in church tax placebo effects; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Placebo_1	14.6544	(13.7388)	0.2861	[-12.2732, 41.5820]	0	13
Placebo_2	11.5513	(39.7255)	0.7712	[-66.3092, 89.4119]	0	8
Placebo_3	16.0167	(84.1799)	0.8491	[-148.9729, 181.0063]	0	3
Placebo_4	-9.0000	(14.0819)	0.5227	[-36.6000, 18.6000]	0	1

Table 27: Decreases in church tax average cumulative effect; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Average cumulative effect	82.3309	(308.0372)	0.7893	[-521.4108, 686.0727]	822	139

Figure 21: Event-study coefficients plot for decreases in church tax; Source: Rendering of Church of Sweden data by the author

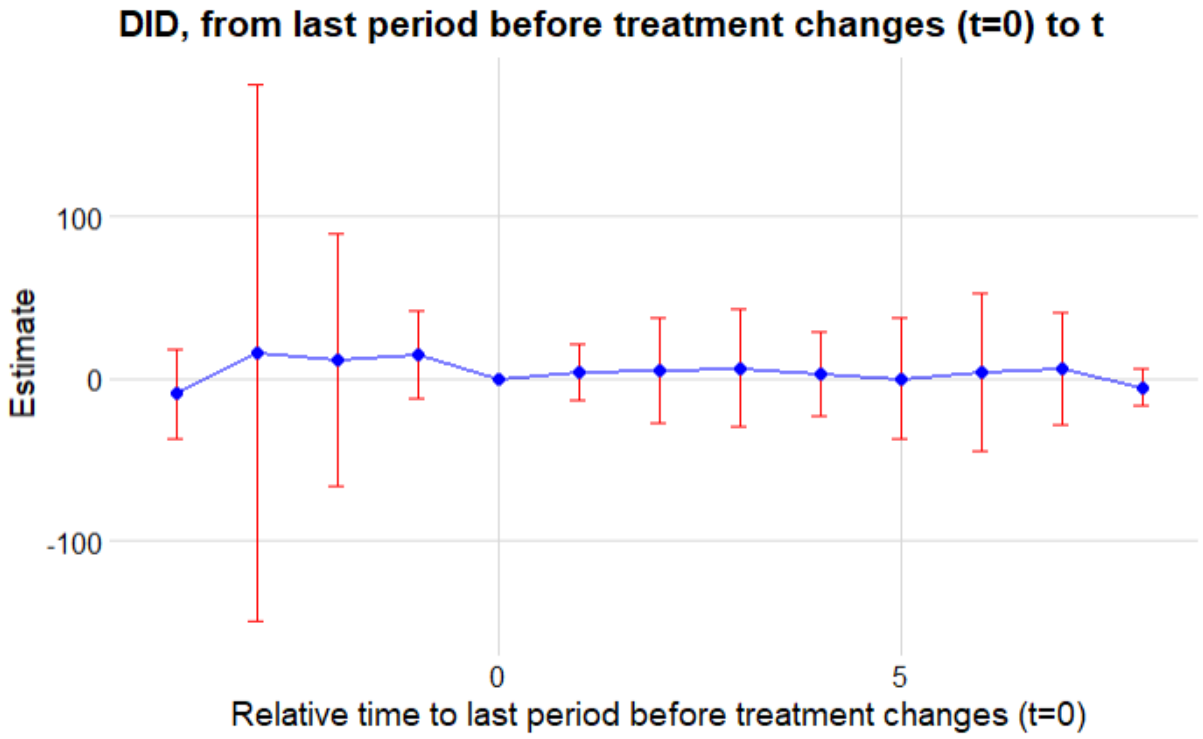

Table 28: Diocese-specific trends model summary; Source: Rendering of Church of Sweden data by the author

Statistic	Value
Specification	Tax increases with stift-specific nonparametric trends
Number of effects	8
Number of placebos	4
Switchers in first effect	93
Joint p-value, effects	0.2536
Equality p-value, effects	0.3367
Joint p-value, placebos	0.1073
Switchers type	Switchers-in
Nonparametric trends	Stift
Clustered SE	Pastoral unit

Table 29: Diocese-specific trends event-study coefficients; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Effect_1	-1.8181	(3.9857)	0.6483	[-9.6300, 5.9938]	221	93
Effect_2	4.6957	(7.3479)	0.5228	[-9.7059, 19.0973]	199	86
Effect_3	4.8689	(8.1844)	0.5519	[-11.1722, 20.9100]	172	75
Effect_4	18.1754*	(9.0627)	0.0449	[0.4128, 35.9380]	150	67
Effect_5	11.9660	(9.8503)	0.2244	[-7.3402, 31.2723]	118	54
Effect_6	1.8571	(7.3290)	0.8000	[-12.5074, 16.2216]	95	42
Effect_7	2.7883	(10.2654)	0.7859	[-17.3315, 22.9080]	81	37
Effect_8	-6.6296	(6.2174)	0.2863	[-18.8155, 5.5563]	39	18

Table 30: Diocese-specific trends placebo effects; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Placebo_1	-7.9566	(6.8552)	0.2458	[-21.3925, 5.4792]	175	73
Placebo_2	-15.4631	(13.7288)	0.2600	[-42.3711, 11.4449]	124	52
Placebo_3	-2.9762	(16.5795)	0.8575	[-35.4715, 29.5191]	68	28
Placebo_4	35.0167*	(17.5070)	0.0455	[0.7036, 69.3298]	44	20

Table 31: Diocese-specific trends average cumulative effect; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Average cumulative effect	65.0275	(68.1198)	0.3398	[-68.4849, 198.5399]	1,001	472

Figure 22: Event-study coefficients plot for diocese-specific non-parametric trends; Source: Rendering of Church of Sweden data by the author

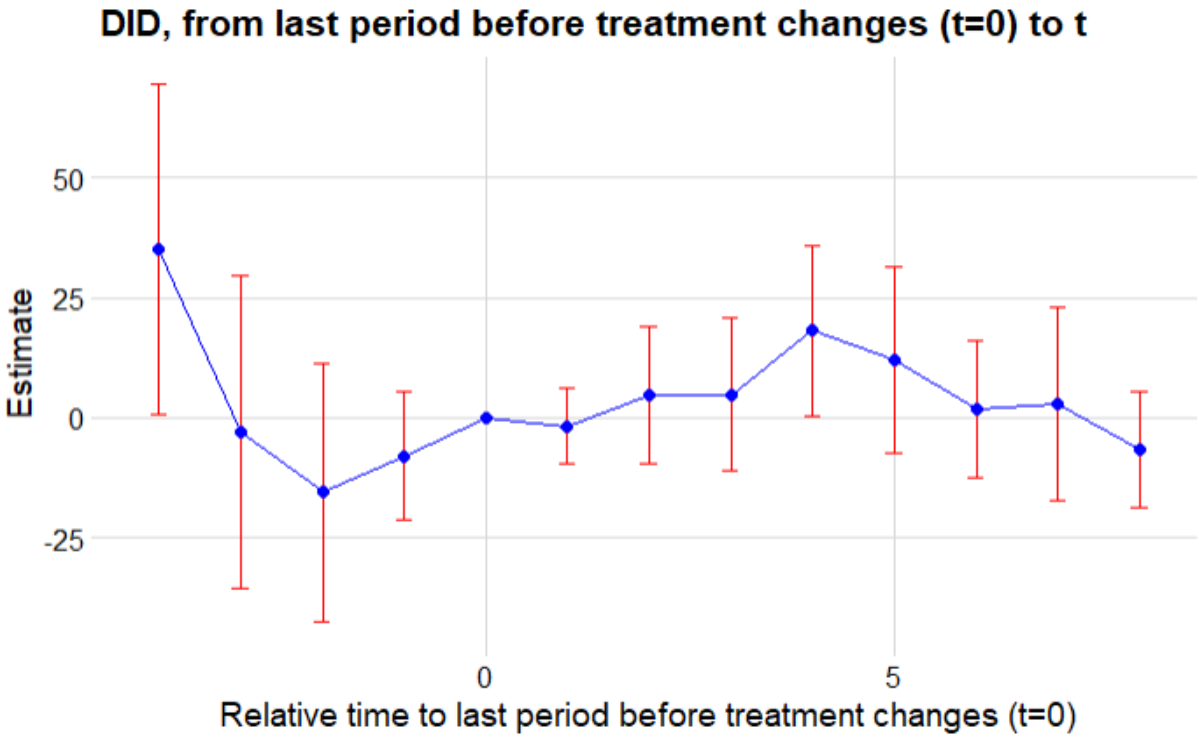

Table 32: Adding diocesan church tax model summary; Source: Rendering of Church of Sweden data by the author

Statistic	Value
Specification	Total tax including stiftsavgift
Treatment variable	Church tax + stiftsavgift
Number of effects	8
Number of placebos	4
Switchers in first effect	229
Joint p-value, effects	0.7414
Equality p-value, effects	0.7212
Joint p-value, placebos	0.8438
Switchers type	Switchers-in
Clustered SE	Pastoral unit

Table 33: Adding diocesan church tax event-study coefficients; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Effect_1	-9.3982	(7.9875)	0.2394	[-25.0535, 6.2571]	965	229
Effect_2	-6.9661	(8.6165)	0.4188	[-23.8542, 9.9220]	837	221
Effect_3	-3.9943	(8.1510)	0.6241	[-19.9700, 11.9815]	732	207
Effect_4	4.7550	(5.5013)	0.3874	[-6.0273, 15.5372]	663	196
Effect_5	5.4664	(6.1307)	0.3726	[-6.5495, 17.4824]	580	178
Effect_6	6.4429	(7.6952)	0.4024	[-8.6395, 21.5253]	537	169
Effect_7	10.1288	(7.8267)	0.1956	[-5.2113, 25.4689]	420	139
Effect_8	2.0267	(4.7585)	0.6702	[-7.2998, 11.3531]	216	70

Table 34: Adding diocesan church tax placebo effects; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Placebo_1	-2.7367	(10.9786)	0.8031	[-24.2543, 18.7810]	778	193
Placebo_2	-4.2201	(16.6952)	0.8004	[-36.9421, 28.5018]	491	136
Placebo_3	6.2497	(12.6010)	0.6199	[-18.4478, 30.9472]	216	56
Placebo_4	-4.3621	(15.7960)	0.7824	[-35.3216, 26.5974]	84	22

Table 35: Adding diocesan church tax average cumulative effect; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Average cumulative effect	0.2253	(64.3929)	0.9972	[-125.9824, 126.4331]	3,014	1,409

Figure 23: Event-study coefficients plot for total church tax; Source: Rendering of Church of Sweden data by the author

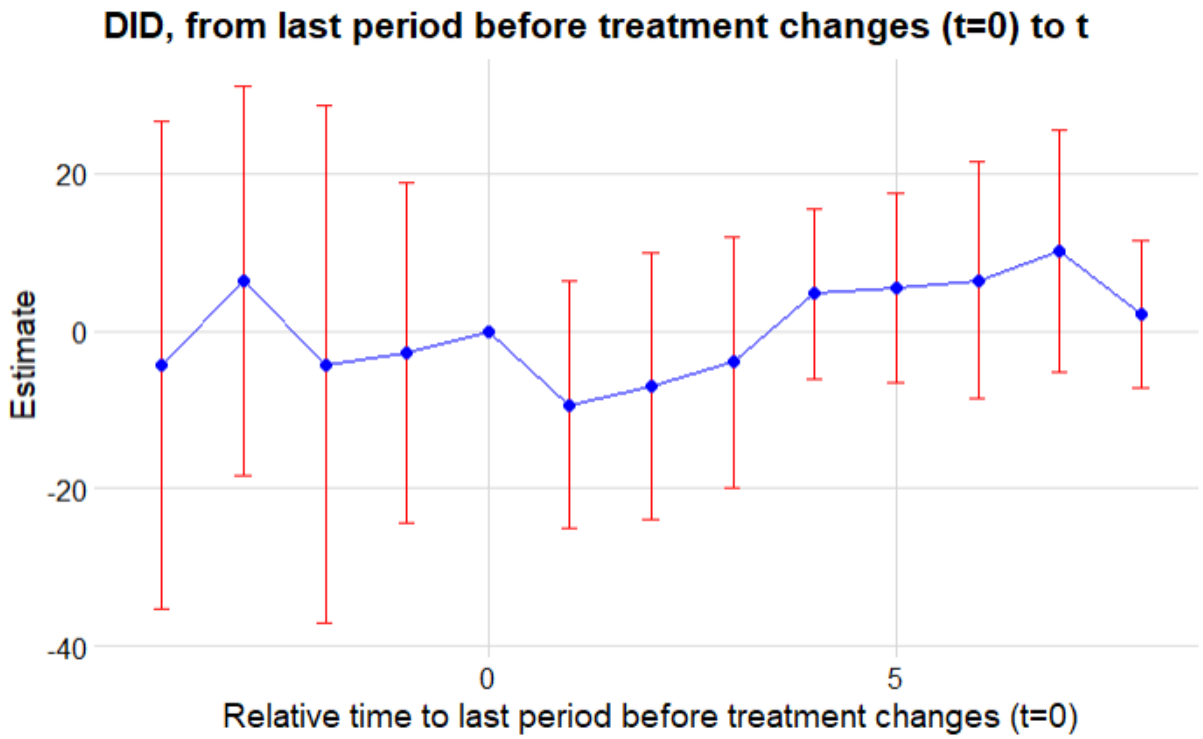

Table 36: Only never-switchers model summary; Source: Rendering of Church of Sweden data by the author

Statistic	Value
Specification	Only never switchers as controls
Controls	Never switchers only
Number of effects	8
Number of placebos	4
Switchers in first effect	226
Joint p-value, effects	0.0950
Equality p-value, effects	NA
Joint p-value, placebos	0.2932
Switchers type	Switchers-in
Clustered SE	Pastoral unit

Table 37: Only never-switchers event-study coefficients; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Effect_1	-15.8917*	(7.2851)	0.0292	[-30.1703, -1.6131]	814	226
Effect_2	-15.4633	(9.5140)	0.1041	[-34.1105, 3.1838]	791	219
Effect_3	-2.0008	(6.3439)	0.7525	[-14.4346, 10.4329]	759	212
Effect_4	3.8286	(8.3350)	0.6460	[-12.5077, 20.1648]	721	205
Effect_5	0.1256	(9.4688)	0.9894	[-18.4329, 18.6841]	618	180
Effect_6	7.6540	(5.8810)	0.1931	[-3.8725, 19.1805]	541	162
Effect_7	13.9608*	(7.0349)	0.0472	[0.1727, 27.7489]	512	153
Effect_8	8.2558	(5.7953)	0.1543	[-3.1028, 19.6144]	309	87

Table 38: Only never-switchers placebo effects; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Placebo_1	4.1837	(5.2128)	0.4222	[-6.0332, 14.4007]	679	189
Placebo_2	-2.4452	(7.7062)	0.7510	[-17.5491, 12.6587]	481	131
Placebo_3	-39.0601	(26.5161)	0.1407	[-91.0307, 12.9105]	246	58
Placebo_4	-18.3861	(29.7998)	0.5372	[-76.7927, 40.0205]	180	43

Table 39: Only never-switchers average cumulative effect; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Average cumulative effect	-19.2407	(61.4787)	0.7543	[-139.7367, 101.2553]	3,170	1,444

Figure 24: Event-study coefficients plot for only never-switchers; Source: Rendering of Church of Sweden data by the author

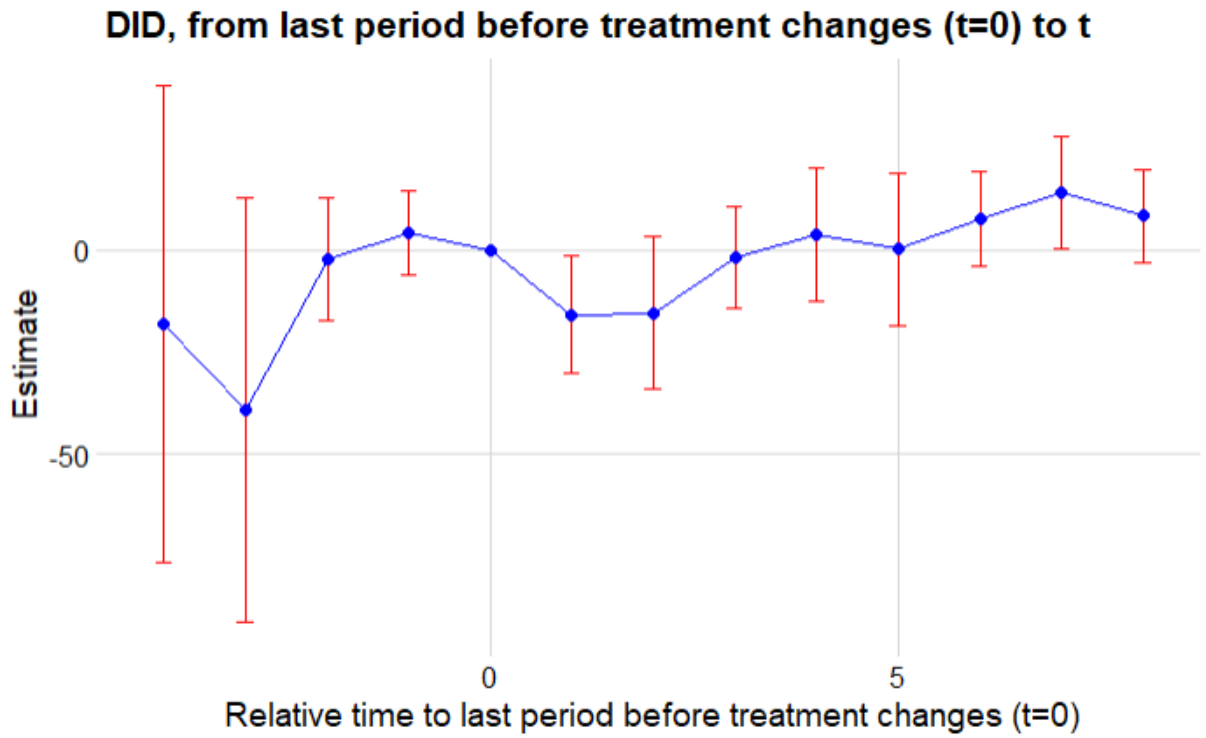

Table 40: Less conservative SEs model summary; Source: Rendering of Church of Sweden data by the author

Statistic	Value
Specification	Sensitivity with less conservative standard errors
Number of effects	8
Number of placebos	4
Switchers in first effect	228
Joint p-value, effects	0.0397
Equality p-value, effects	NA
Joint p-value, placebos	0.2337
Switchers type	Switchers-in
Less conservative SE	Yes
Clustered SE	Pastoral unit

Table 41: Less conservative SEs event-study coefficients; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Effect_1	-13.4089*	(6.7390)	0.0466	[-26.6171, -0.2007]	1,061	228
Effect_2	-14.3074	(9.2594)	0.1223	[-32.4556, 3.8408]	950	220
Effect_3	-2.6961	(6.6752)	0.6863	[-15.7793, 10.3871]	862	213
Effect_4	3.7399	(8.6512)	0.6655	[-13.2161, 20.6959]	787	206
Effect_5	0.4622	(8.7802)	0.9580	[-16.7467, 17.6711]	658	180
Effect_6	7.9375	(5.7324)	0.1662	[-3.2979, 19.1729]	562	162
Effect_7	14.1435*	(6.8954)	0.0403	[0.6287, 27.6583]	525	153
Effect_8	8.2422	(5.7743)	0.1535	[-3.0751, 19.5595]	313	87

Table 42: Less conservative SEs placebo effects; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Placebo_1	-0.2728	(4.7577)	0.9543	[-9.5976, 9.0521]	830	191
Placebo_2	-7.5984	(8.2314)	0.3560	[-23.7316, 8.5347]	522	131
Placebo_3	-40.6561	(26.5255)	0.1253	[-92.6452, 11.3329]	251	58
Placebo_4	-18.9481	(29.7918)	0.5248	[-77.3391, 39.4428]	182	43

Table 43: Less conservative SEs average cumulative effect; Source: Rendering of Church of Sweden data by the author

Term	Coefficient	Std. Error	p-value	95% CI	N	Switchers
Average cumulative effect	-13.4657	(63.1217)	0.8311	[-137.1820, 110.2506]	3,591	1,449

Figure 25: Event-study coefficients plot for less conservative standard errors; Source: Rendering of Church of Sweden data by the author

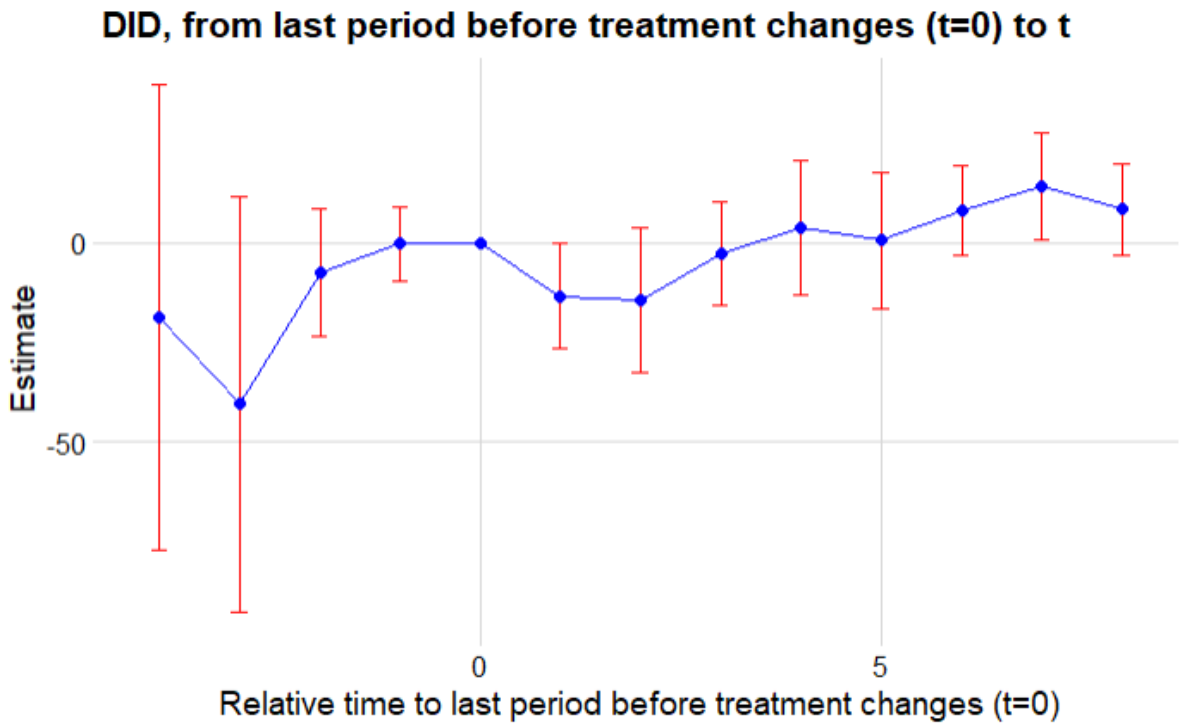

AI Disclosure

In the process of writing my bachelor's thesis in economics, as approved by my supervisor Sampreet Goraya, I have applied the use of generative artificial intelligence (AI), in order to facilitate the writing process. In general, the AI models used have increased the quality of the thesis. However, in my opinion, they have done that mostly indirectly by relieving me of a lot of tedious work that I would otherwise have had to do myself. If I had not had access to generative AI, I would have had to spend days, if not weeks, doing the tasks which AI was able to do within a few minutes. Therefore, because AI gave much more time to work with, I was able to focus more on the substantive parts of the thesis and thereby improve their quality. I shall give more concrete examples of quality improvements below.

I have used AI for two purposes:

1. Finding appropriate scientific articles connected to the area of research that my thesis concerns itself with. For this purpose, I used the AI tool on scispace.com, which is an AI application specialised in academia. It helped me to find relevant scientific articles, whose contents I then included in my theory and my literature review.

   Scispace has contributed to my thesis in the way that I was able to find some of the most appropriate articles for my theory and literature review – articles which I may not have been able to find on my own. A risk with using Scispace was, however, that it could pick up keywords in an article and then claim that article is about a certain topic when, in reality, it was about a completely different topic. I tried to mitigate this risk by carefully checking every article that Scispace gave me, in order to see whether the article actually is about the topic that it is purported to be.

2. The second purpose of generative AI in my thesis was writing the codes for my empirical analysis. My empirical analysis involved using long scripts of code in Python and R, which formed the building blocks for the regressions, descriptive statistics and event studies described in my thesis. In order to facilitate the writing of these scripts, I used the generative AI models available on chatgpt.com. Specifically, I have mostly been using the model ChatGPT based on the current version GPT-5.5. Most of the tables and figures in my thesis were created by chunks of code that I instructed ChatGPT to write.

   ChatGPT has contributed to my thesis in the way that I was able to acquire the needed code scripts extremely quickly and efficiently, something that, as mentioned above, saved me many hours of work. In addition, given the quality of the codes provided by ChatGPT, I believe that they are somewhat more "clean" and empirically coherent than the codes that I would have written myself, even though I might have gotten the same results from my own coding. One central risk associated with using ChatGPT for coding is the fact that, despite its excellent ability to write codes, it is not always great at understanding my prompts, i.e. what I want from it. In other words, it is almost perfect at doing what it believes it is supposed to do, but not always good at understanding what I want it to do. In order to mitigate this problem, on one hand, I was very careful in writing good prompts, and on the other hand, I tried to make sure to cross-check the resulting tables etc. with manual calculations.

I gained various insights from utilising AI in my thesis. Most importantly, I realised that generative AI is not always as good as you think it is. It makes mistakes frequently, especially when it comes to complex econometric topics, and you can never completely rely on it, hence highlighting the importance of always taking its answers and results with a grain of salt. Perhaps even more crucially, I have observed that generative AI always is completely convinced of its own excellence and thus never initially considers the possibility of being wrong (before being proven wrong). I experienced this first hand when I was debugging errors in my code and ChatGPT basically tried gaslighting me into believing that the code works when, in fact, it did not work at all. Such behaviour where AI tools give you information with complete confidence when it, in reality, is nonsense, can be disastrous for a thesis, especially in case you are dependent on e.g. the coding knowledge of an AI tool. Therefore, in my thesis writing process, I have tried to be aware of this issue and always be slightly sceptical about answers given to me by AI tools.