The London School of Economics and Political Science

Essays in Public Economics

Mohammad Vesal

A thesis submitted to the Department of Economics of the London School of Economics and Political Science for the degree of Doctor of Philosophy, London, July 2014.

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

The following disclimar applies to chapters 1 and 2 of the thesis that are based on HMRC data: “This work contains statistical data from HMRC which is Crown Copyright. The research datasets used may not exactly reproduce HMRC aggregates. The use of HMRC statistical data in this work does not imply the endorsement of HMRC in relation to the interpretation or analysis of the information.”

I declare that my thesis consists of 44206 words in total.

Mohammad Vesal
July 2014

Abstract

I present three essays in this thesis. The first essay investigates the decision of small businesses with respect to an optional Flat Rate Scheme (FRS) in the UK. FRS replaces VAT with a turnover tax providing some traders with a tax saving opportunity. Using the universe of VAT returns between 2004-05 and 2010-11, I find 26 percent of eligible traders have non-negative tax gains from FRS. I show gains are highly persistent and not so small, yet only 3 percent of gainers join the scheme after one year. Temporal and spatial correlations point to information frictions and learning as potential explanatory factors. Results show traders registering after introduction of FRS and those registering in high FRS density areas are more likely to join the scheme. The second essay estimates stimulus effect of the temporary reduction in the standard VAT rate in the UK. From 1 December 2008 to 31 December 2009, the standard-rate was reduced from 17.5 to 15 percent. I use the universe of VAT returns submitted to HMRC between 2002q1 and 2010q4 and compare changes in sales growth of standard-rated traders during the cut to that of zero-rated traders (difference-in-differences). To control for heterogeneous recession effects, I first rely solely on post-recession observations and utilize the fact that the cut and the recession don't fully overlap. Second, I allow for sector specific recession impacts. Both strategies show a small insignificant impact on gross sales and purchases which suggest a proportionate increase in quantity demanded in response to the tax induced price cut. The third essay estimates the impact of Iran Iraq war on educational attainment of children. I use a two percent sample of 2006 Iran Population Census, and compare exposed cohorts in war provinces to unexposed cohorts (difference-in-differences). The estimates suggest probability of finishing high school is respectively reduced by 4.8 and 1.9 percentage points for cohorts exposed to war in early childhood and those exposed during schooling (former significant at 10 percent, latter insignificant). Interestingly, the war impact on early childhood cohorts is robust to controlling for differential linear trends while the impact on school cohorts is not.

Acknowledgments

I am extremely grateful to my supervisors Tim Besley, Henrik Kleven, and Johannes Spinnewijn for all their support and encouragement through my PhD. Without their help and guidance the chapters of this thesis would not have been possible.

Chapters of this thesis have benefited from presentations at HMRC, RES Annual Conference 2014, PEUK Residential Conference 2014, Oxford CBT doctoral meeting 2013, LSE Public Economics, and Development Economics work in progress seminar series. I have also benefited from numerous discussions with colleagues at LSE and elsewhere. I especially would like to thank Michael Best, Florian Blum, Steve Bond, Shawn Chen, Frank Cowell, Michael Devereux, Jason Garred, Camille Landais, Li Liu, Ben Lockwood, Daniel Osorio Rodriguez, and Mazhar Waseem for great comments on various parts of the thesis.

I am also indebted to the HMRC datalab staff, especially Lucy Nicholson, Daniele Bega, Chioma Anaba, and John Haynes, for providing the data for chapters 1 and 2 and dealing with my numerous data requests.

Finally, I am incredibly grateful to my parents and my beloved wife Sabrieh for all their love and support through these years.

Contents

List of Figures

List of Tables

Chapter 1

Optimization Frictions in the Choice of the UK Flat Rate Scheme of VAT

1.1 Introduction

There is growing evidence in public economics that optimization frictions play an important role in shaping individual behavior. Whether small businesses are subject to similar frictions has not received much attention. An individual owner-manager is often responsible for business decision making but theoretically, one cannot generalize the individual-based evidence to small businesses. Business owners have shown particular skills (e.g. started a business) that might reduce the effect of frictions. Understanding role of optimization frictions in the business environment is important from two perspectives. Conceptually, it affects the way economists think about profit maximization. From a policy perspective, it is important to understand frictions in business decision making to design effective support schemes.

In this chapter, I study the decision of VAT registered traders with respect to the Flat Rate Scheme of VAT for small businesses (FRS). I use HM Revenue and Customs' (HMRC) VAT returns data to calculate FRS tax gains for eligible traders. This is the first paper that analyzes FRS using tax return data. FRS is an optional scheme introduced in 2002 to alleviate compliance burden of VAT on small businesses. Normally, VAT liability is the difference between VAT on sales and purchases. HMRC requires record keeping of business transactions showing separation of zero, reduced, and standard-rated sales and purchases. FRS liability[1] is, however, calculated as a

[1] I refer to VAT liability under FRS as FRS liability, but once traders join FRS this is their VAT

percentage of gross sales, relieving traders of the need to account for various rates separately. In order to compensate for the inability of FRS traders to reclaim purchases VAT, HMRC sets sector specific flat rates so that on average FRS and VAT liabilities are equalized.

In order to join the scheme, traders need to fill out a one-page form telling HMRC of their main activity (and hence flat rate) and declaring their eligibility. In the absence of optimization frictions, eligible traders should join FRS when expected net benefits are positive. While the scheme could potentially benefit traders via reduced tax payments and lower compliance cost, I focus on pure tax savings for two reasons. First, anecdotal evidence suggests tax savings play a key role in the FRS joining decision. For example, an HMRC study of compliance cost of VAT conducted by KPMG reports "the predominant theme ... is that [traders] enter into the FRS to save them money in terms of the amount of VAT paid to HMRC" (KPMG (2006)). Second, returns data does not provide any information on the amount of time businesses spend on preparing their VAT returns or whether they use tax preparators.

I define FRS gainers as eligible VAT traders with observed FRS liability less than or equal to the reported VAT liability. I show that between 2004-05 and 2010-11, 26 percent of eligible traders are FRS gainers. Following FRS gainers over time reveals little responsiveness. The estimated probability of joining within one year of gaining is 3 percent and increases to 10 percent after six years. This is despite the fact that gains are persistent and not very small. On average 70 percent of FRS gainers in a given year remain a gainer in the following year and the median FRS gainer would save about 12 percent on VAT payments upon joining the scheme.

Since FRS joining decision is made ex ante, inaction of gainers is not necessarily a sign of sub-optimal choices. Risk neutral traders would join the scheme when expected benefits are positive. Presence of uncertainty could result in observed gains even if expected gains are negative. Two pieces of evidence, however, go against this explanation. First, I show the probability of joining FRS rises sharply as traders get slightly positive gains. This suggests that at least for a sub-sample of traders, observed gains could be interpreted as expected gains[2]. The caveat here is that the sub-sample of responsive traders might have different risk preferences or face a

liability from HMRC's perspective. Similarly I refer to tax liability under normal VAT accounting as VAT liability.

[2]This requires the assumption that traders joining the scheme are not making a mistake themselves.

different level of uncertainty.

The second piece of evidence against uncertainty is the fact that FRS gains are highly persistent. Even after controlling for sector and year dummies, last year gainers are on average 62 percentage points more likely to gain in the following year. Furthermore, the probability of gaining in future rises very sharply right at zero past gains and goes beyond 80 percent for traders with gains above £1000 during last year. The distribution of current FRS gains conditional on gaining in the last year shows a median tax saving of 10 percent of VAT liability and a mean of just above zero for large enough traders[3].

After discussing that uncertainty cannot fully explain inaction of FRS gainers, I move to characterize the frictions that prevent traders from joining using temporal and spatial correlations. Here, the FRS joining patterns support a combination of broadly defined information frictions and learning as key drivers of inaction. I define information frictions to include both lack of knowledge about FRS rules and unawareness of its existence. I use learning to refer to a case where traders know about the scheme but are not certain about its benefits. This could be a result of uncertainty or a consequence of incorrect prior beliefs about suitability of FRS.

First, I conjecture that VAT registration is a period of intense learning about VAT rules. Therefore the chance of coming across FRS is the highest during this time. I split the sample into three groups based on the date of VAT registration: a) pre-FRS traders who registered before introduction of FRS, b) early-FRS traders who registered after introduction of FRS but before major reforms in 2004, and c) late-FRS traders who registered after favorable FRS reforms in 2004[4]. Late-FRS traders could learn about the reformed FRS and are expected to have the highest chances of joining the scheme. On the other hand, pre-FRS traders registered when FRS was not in place and should have least awareness of the scheme. Consistent with this reasoning, non-parametric estimates of joining probabilities are always significantly higher for late-FRS compared to early-FRS traders. Similarly early-FRS traders show higher joining probabilities relative to pre-FRS traders. Restricting the sample to FRS gainers confirms a similar pattern: late-FRS gainers are significantly more likely to join FRS with early and pre-FRS groups lagging behind.

[3]With risk averse preferences, positive expected FRS gains may not justify optimality of uptake. In section 1.5 I discuss some features of the scheme to argue that even gainers with risk averse preferences might benefit from the scheme.

[4]In 2004 FRS rates were reduced and a temporary 1 percentage point discount was applied to traders joining the scheme during first year of VAT registration.

Second, I argue that traders registering in postcode districts (outcodes) with a higher density of FRS traders are expected to have higher FRS awareness (e.g. through peer groups). I look at joining probabilities for traders registering in high and low FRS density outcodes. The non-parametric estimates show, traders registering in the highest decile of FRS density are significantly more likely to join the scheme compared to those in the lowest decile. Furthermore, FRS gainers registered in outcodes with higher FRS densities are significantly more likely to join the scheme later on.

For both temporal and spatial correlations, I observe that joining probabilities increase over time. In other words, it seems that some FRS gainers realize that they could gain from FRS and join the scheme later on. While this pattern could be consistent with inertia (sluggish responsiveness), learning, or gradual spread of information about the scheme, I argue that the spatial correlations are not fully consistent with inertia. For example, inertia cannot explain the higher joining probabilities for high FRS density outcodes unless a disproportionate number of more active traders are registered in these places.

To look at the relative importance of these explanations and to rule out inertia I estimate Cox proportional hazard (CPH) models. After controlling for 5-digit sectors and FRS density deciles (stratified CPH), I still find traders registering later are more likely to join the scheme. Furthermore, I find support for learning. An additional year of gaining leads to higher likelihood of joining even after controlling for period of registration. Including a continuous variable for FRS density (instead of stratification on decile dummies) shows traders in outcodes with higher FRS densities are more likely to join the scheme.

The conclusion that small traders are susceptible to optimization frictions resonates with the results of Devereux et al. (2014) who find small incorporated businesses are not completely shifting their incomes to the corporate base while in a frictionless world it is optimal to do so. Their preferred explanation is illiquidity of corporate profits and the need for having a stable flow of income (e.g. in the form of personal income). In this paper, however, I argued for presence of information frictions which implies gainers would join FRS if they get the right information. My results suggest small businesses might be subject to optimization frictions similar to those observed in the context of individual decision making. Accepting this view in the case of FRS, calls for a more effective role of the government in publicizing the scheme.

The results are also consistent with the large empirical literature on the importance of frictions in the process of individual decision making. Chetty et al. (2011) find that presence of search costs and hours constraints imply individuals re-optimize only when the tax gains are sufficiently high. This is consistent with an observed positive correlation between estimated labor supply elasticities and size of tax variations in Denmark. Kleven and Waseem (2013) find a significant mass of individual tax filers in Pakistan locate in strictly dominated regions above tax notches. They provide evidence that 90% of wage earners and 50-80% of self-employed in these areas are not responsive to tax incentives potentially due to frictions. Jones (2012) provides evidence that inertia could explain why so many income tax filers receive a tax refund although it might be optimal to adjust tax payments and not pay the money in the first place.

Bhargava and Manoli (2013), Chetty et al. (2013), Liebman and Luttmer (2011), Saez (2009) find direct evidence that provision of information changes individual decisions. Bhargava and Manoli (2013) design a randomized experiment to understand high non take-up of EITC benefits. They find re-sending a reminder letter for potential EITC benefits is most effective in increasing take-up when the information is simplified and the size of potential benefits is displayed. Chetty et al. (2013) show neighborhoods with higher EITC information are more responsive to the incentives created by the program and households moving into high information areas start to optimize their EITC soon after. In the context of social security Liebman and Luttmer (2011) find an information brochure and an invitation for a web based tutorial increases labor force participation by 4 percentage points one year later. Saez (2009) shows both explaining incentives and presentation details matter for take-up of retirement savings subsidies.

Some other studies however find a minimal role for information indirectly pointing to significance of other frictions. Chetty and Saez (2013) show there is a limited effect of providing information on take-up of EITC in a randomized setting. Jones (2010) finds providing information about advance EITC, an add-on feature paying interim installments, does not change take-up of the program significantly. Investigating retirement saving decisions Choi et al. (2011) find providing information to 401(k) participants with strictly dominated contribution rates does not change their behavior significantly. They conjecture presence of biased preferences might be responsible for unresponsiveness.

In the next section, I give a detailed account of the rules around FRS. In the third section I describe the data. Section four establishes the fact that a significant number of VAT traders benefit from FRS but fail to join the scheme. In section five I discuss why uncertainty cannot fully explain inaction of FRS gainers. Section six presents temporal and spatial correlations that suggest information frictions and learning are potential explanations for low uptake. The last section concludes.

1.2 Flat Rate Scheme

HMRC first announced the Flat Rate Scheme of VAT for small businesses (FRS) with a consultation in June 2001. The scheme came to force from 24 April 2002 as part of the Finance Act 2002 with the stated purpose of reducing compliance burden of VAT on small businesses. Businesses in the UK must register for VAT when their annual turnover goes beyond a registration threshold (£67,000 during 2008). VAT features three different rates (standard, reduced, and zero) and a set of exempt activities. Normal VAT liability is the difference between VAT on sales and purchases while VAT liability under FRS is the multiplication of a sector specific tax rate and total turnover. As a result FRS requires businesses to keep track of total turnover rather than separate record of transactions under each of the various VAT rates and therefore it is thought to simplify compliance. Effectively VAT is a tax on value added while FRS liability is a tax on gross sales as shown below:

$$T_V = \tau_V v S_g \tag{1.1}$$

$$T_F = \tau_F S_g \tag{1.2}$$

where T_V and T_F respectively represent VAT and FRS liability, S_g is gross sales, v is share of value added (defined as $\frac{S_g - P_g}{S_g}$, with P_g being gross purchases), τ_V is effective VAT rate (defined as $\frac{T_S - T_P}{vS_g}$, with T_s and T_p respectively showing sales and purchases VAT), and τ_F is the flat rate percentage. Eligible traders decide ex ante to be liable either for T_V or T_F over an accounting period. HMRC sets flat rates by sector so the average traders within sectors are indifferent between FRS and VAT: "We calculate the flat rate percentages from the net tax paid by all the businesses that are currently registered for VAT and eligible for the scheme. The net tax paid varies with different trade sectors and so there are a variety of flat rate percentages"[5]. Nevertheless traders with lower than average purchases VAT would get substantial gains from FRS. For example, a management consultant with no purchases VAT could save 16 percent

[5]HMRC, Notice 733: Flat rate scheme for small businesses, February 2004.

on VAT payment by joining FRS during 2004-7[6]. There are around 16 distinct flat rates ranging from 2 to 14.5 percent (appendix A). On January 2004, HMRC lowered all but one flat rate, increased eligibility thresholds, and incentivized new VAT registrations to join FRS by offering a 1 percentage point discount on flat rates within the first 12 months of registration. To maintain the attractiveness of FRS when standard VAT rate changed, HMRC revised the flat rates on 1 December 2008, 1 January 2010, and 4 January 2011.

While FRS is advertised as a compliance cost saving scheme[7], anecdotal evidence suggests most businesses view the scheme as a tax saving opportunity. An HMRC study of VAT compliance cost reports that "the predominant theme ... is that [traders] enter into the FRS to save them money in terms of the amount of VAT paid to HMRC" (KPMG (2006)). Same study states that businesses spend resources to determine whether FRS is suitable for them, which suggests information about FRS gains is not readily available. In addition, in the initial FRS consultation, accountancy firms argued the scheme would not generate any of the intended savings and opposed the scheme as undermining VAT accounting discipline (HM Customs and Excise (2002)). Presence of any compliance cost savings would strengthen the evidence on the sub-optimality of the inaction of FRS gainers. But I ignore compliance cost savings in what follows because returns data does not provide any information on the amount of time businesses spend on preparing their VAT returns or whether they use tax preparators[8].

Eligible VAT traders could easily and quickly join or leave FRS. Traders wishing to join, fill in a one-page application form declaring main activity from the list in appendix A, the corresponding flat rate, and sign that they are eligible. FRS start

[6] τ_F for management consultants is 12.5 percent. With a standard-rate of VAT equal to 17.5 percent, the VAT rate on gross sales is $\tau_V = \frac{0.175}{1+0.175} = 14.9$ percent. Therefore, when the trader does not use any tax-refundable inputs (i.e. $v = 1$) the FRS gain as a percentage of current VAT liability is $1 - \tau_F/\tau_V = 1 - 12.5/14.9 = 16.1$ percent.

[7] Initially FRS was claimed to save on average about £750 (HM Customs and Excise (2002)) but later an impact assessment puts the average compliance savings at £45 (HMRC (2009)). The first estimate is based on saving 45 minutes of clerical time at an hourly wage of £16 over the course of 52 weeks plus £100 saving on accountants' fees. The second estimate uses a "Standard Cost Model" but details of calculations are not disclosed.

[8] There is some evidence that a move to FRS might actually increase compliance costs. Accounting software seemed to have lacked FRS capability until recently. For example SAGE 50 Accounts introduced FRS capability in the 2011 upgrade (GfK Business (2008), an HMRC sponsored study, shows from the 58 percent of businesses using accounting software for VAT, 61 percent use SAGE.). Furthermore, there is anecdotal evidence that FRS traders calculate both VAT and FRS liabilities not to lose money on FRS. The mental cost of worrying about losing money and the time cost of calculating two tax liabilities are likely to increase FRS compliance costs. This could be a competing story for the frictions I study in section 1.6.

date is normally the beginning of next VAT period (a quarter for most of traders) and backdating is not normally allowed. Businesses wishing to leave the scheme write to HMRC of their decision and normally stop FRS at the end of current VAT period. Again retrospective departure is usually not allowed. There is no statutory minimum term for being on FRS but once left FRS, the trader cannot rejoin within the following 12 months. As a measure of revenue protection HMRC reserves the right to withdraw the scheme (even back date the withdrawal) in fraudulent cases.

FRS eligibility is based on turnover and non-turnover criteria. Table 1.1 shows turnover eligibility rules. Joining eligibility is based on two tests. Expected taxable turnover should be below a threshold (£150,000 during 2004-10) and expected total turnover should be less than a second threshold (£187,500 until December 2010). Once on the scheme, traders remain eligible until their FRS turnover crosses the continuation threshold (£225,000 during 2004-10). The joining tests are based on forecasts of turnover. Instead, I use actual turnover to determine eligibility. This should do no harm because HMRC suggests traders could use last year turnover as a benchmark for their forecasts and also there is no penalty for falling above the joining threshold once on the scheme. Furthermore, during my sample, a small fraction of eligible traders become ineligible in the following year (8 and 10 percent of FRS gainers and losers respectively).

There are five mostly unobservable non-turnover eligibility criteria that apply at all times[9]. Since the main claim in this paper is that some eligible traders are missing out on tax saving opportunities, it is important to rule out unobserved ineligibility of gainers as a potential explanation. First, traders who were on FRS during the past 12 months cannot rejoin the scheme. Second, firms registered or eligible to be registered as a VAT group in the past 24 months are ineligible. While I observe traders registered as groups during the sample, I do not have information on those eligible for group treatment or prior group registrations. It is, however, encouraging to note that only 0.3 percent of VAT traders below FRS continuation threshold are registered as a group.

Third, FRS cannot be combined with certain VAT schemes (capital goods[10], cash accounting, retail, tour operators, margin and auctioneer's schemes). I do not have reliable information on take-up of these schemes but several observations justify

[9]Unfortunately, official data on the number of ineligible traders or applications ruled out as ineligible is not available.

[10]Traders purchasing property or doing refurbishment with a value greater than £250,000 or acquire computer and related equipment with value greater than £50,000 must use the capital goods scheme.

Table 1.1: FRS turnover eligibility criteria

Dates	Joining eligibility		Continuation eligibility FRS turnover (incl. VAT)
	Test 1 Taxable turnover (excl. VAT)	Test 2 Total turnover (excl. VAT)	
April 02 - December 03	<100k	<125k	<150k
January 04 - February 07	<150k	<187.5k	<225k
March 07 - December 10		<187.5k	
January 11 - now		-	<230k

Notes: Taxable turnover (test 1) is the sum of zero, reduced and standard rated supplies excluding any VAT. It excludes exempt supplies and non-business income like charitable or educational activities. Total turnover (test 2) is taxable turnover plus exempt supplies, and non business income such as charitable and educational activities. During March 2007 until December 2010, total turnover for test 2 includes VAT. FRS turnover is VAT inclusive total turnover (e.g. includes exempt, zero, reduced, and standard rated supplies plus any VAT but exclude non business income). Non turnover eligibility criteria are the same across the years. Sources: FRS notices dated February 2004, March 2007, January 2010, April 2011, August 2011, October 2012.

ignoring them. FRS provides an alternative to cash accounting and retail schemes. Furthermore, it is unlikely that traders on margin and tour operator schemes benefit from FRS because of the high level of VAT refunds they receive with these schemes. Therefore, remaining on another scheme is unlikely to be an important factor in analysis of FRS gains.

Fourth, any VAT conviction or dishonesty in the past 12 months disqualifies the firm. Data on VAT dishonesties and convictions is not available. It is, however, unlikely that a big part of FRS gainers fall in this category. National Audit Office reports that out of 196,000 investigations during 2002-03 financial year around 30% of cases had VAT under-declaration but only 4% received a penalty (National Audit Office (2004)). Furthermore, traders with negative VAT liability are under greater scrutiny and a disproportionate number of them are caught in fraudulent activities (National Audit Office (2006)). But traders receiving a net VAT refund would not gain from FRS since my calculated FRS liability is always positive.

Fifth, businesses associated with others[11] are ineligible. This measure was put in place to stop artificial splitting of activities into different entities for tax benefits. For example, a trader with several businesses could concentrate standard rated sales

[11]HMRC clarifies that this is based on commercial reality not legal form and applies to cases where a company has the right to give directions to another or complies with directions of another.

under FRS running entity but report purchases under the one using normal VAT. While HMRC collects data on connections to other businesses from VAT registration form, this data is not available for the current paper. Given the large number of gainers and the small size of traders involved it seems unlikely this criterion creates a major problem.

1.3 Data

Data used in this paper is the annualized version of all VAT returns submitted to HMRC between 2004-5 and 2010-11 financial years. This data has become available recently and this is the first paper analyzing FRS using this data. VAT returns include information on sales, purchases, and corresponding VAT on each but does not provide separate account of transactions under each VAT rate. The returns data is merged with part of HMRC's trader characteristics dataset which provides information on date of registration, date of deregistration, date of joining/leaving FRS, sector of activity, frequency of submitting returns, ownership form, and a few other variables. I refer to this dataset as returns-level data as it includes all returns submitted by traders. From this, I also construct a trader-level dataset which has one observation per trader and records the date of certain events of interest (e.g. VAT registration, joining FRS, etc.). The trader-level dataset only contains traders who are observed to be eligible at least once during the sample (includes FRS traders as well).

Table 1.2 shows the total number of available observations before and after cleaning, and the number of returns submitted by VAT and FRS traders during each financial year. There are around 2 million VAT registered traders in each year (column (1)). Dropping inactive traders, returns reporting zero sales, and other anomalies (see table notes and appendix C for more detail) result in around 1.5 million returns per year (column (2)) . This constitutes the working sample for the analysis in the paper. Based on observable eligibility criteria (see section 1.2) on average 54 percent of VAT traders are FRS eligible (column (4)). Column (5) reports the number of returns submitted by FRS traders which is a relatively small fraction of total returns (column (6)). The fraction of FRS returns increases from 9 to 21 percent of all eligible traders between 2004 and 2010 (column (6))[12]. The increase in share of FRS traders during

[12]Eligible traders is used to refer to VAT traders who are eligible for FRS. All eligible traders include eligible VAT traders and FRS traders.

Table 1.2: Number of VAT and FRS traders

Financial year	All observations (1)	Workable Sample (2)	VAT traders (3)	% FRS eligible (4)	FRS traders (5)	FRS % of eligible (6)
2004-5	1,894,281	1,472,918	1,398,324	56%	74,594	9%
2005-6	2,177,146	1,512,156	1,413,470	57%	98,686	11%
2006-7	2,221,095	1,529,537	1,404,911	54%	124,626	14%
2007-8	2,118,562	1,575,018	1,420,959	54%	154,059	17%
2008-9	2,173,977	1,422,206	1,256,822	51%	165,384	21%
2009-10	2,123,413	1,448,423	1,280,881	52%	167,542	20%
2010-11	2,120,552	1,499,923	1,320,226	52%	179,697	21%
Total	14,829,026	10,460,181	9,495,593	54%	964,588	16%

Notes: Column (1) is number of all available returns. Column (2) shows the cleaned data used for all subsequent analysis and restricts the sample to a) live traders (not reported to be deregistered and identified as live trader at the end of fiscal year by HMRC), b) observations with positive and non missing sales, c) observations with outputs and inputs less than the 99th percentile of the respective distributions, d) observations implying an effective output and input tax rate less than the standard rate plus half a percentage point, e) firms listed as sole proprietors, partnerships, and incorporations, and f) traders with monthly or quarterly VAT returns. Column (3) shows number of VAT returns on normal VAT accounting. Column (4) demonstrates the fraction of VAT traders eligible for FRS based on all observable eligibility criteria (see text for details). Column (5) shows the number of FRS traders and column (6) present FRS traders as a fraction of all eligible traders (actual FRS and FRS eligible traders).

the sample period suggests FRS awareness is increasing but this pattern could be a result of sluggish responsiveness (inertia) or experimenting with VAT (learning).

Many of the traders joining FRS are doing so right at the time of VAT registration. Figure 1.1 shows Kaplan-Meier nonparametric estimate of probability of joining FRS over time[13]. The analysis time reflects the months FRS option was available to the trader. 9 percent of traders join FRS as soon as they have the option to do so. While in principle this jump could be a result of existing VAT traders joining when FRS was introduced, evidence shows this is due to a large number of new traders joining FRS at the time of VAT registration (figure 1.14). After the initial jump, the joining probability continues to rise and by the end of 9 years of exposure to FRS it reaches 18 percent[14].

[13] See section 1.6 for a discussion of Kaplan-Meier method.

[14] The end point estimate of probability of joining FRS is smaller than the fraction of FRS traders as of April 2011 (reported in column (6) of table 1.2) for two reasons. First, the analysis here is based on once eligible traders which includes traders eligible for FRS in 2011 but also those who were eligible earlier and are not eligible at this time. Therefore the number of FRS traders is divided by a larger denominator. Second, figure 1.1 is based on trader rather than return level data and uses Kaplan-Meier estimate of survival function which is not necessarily equivalent to

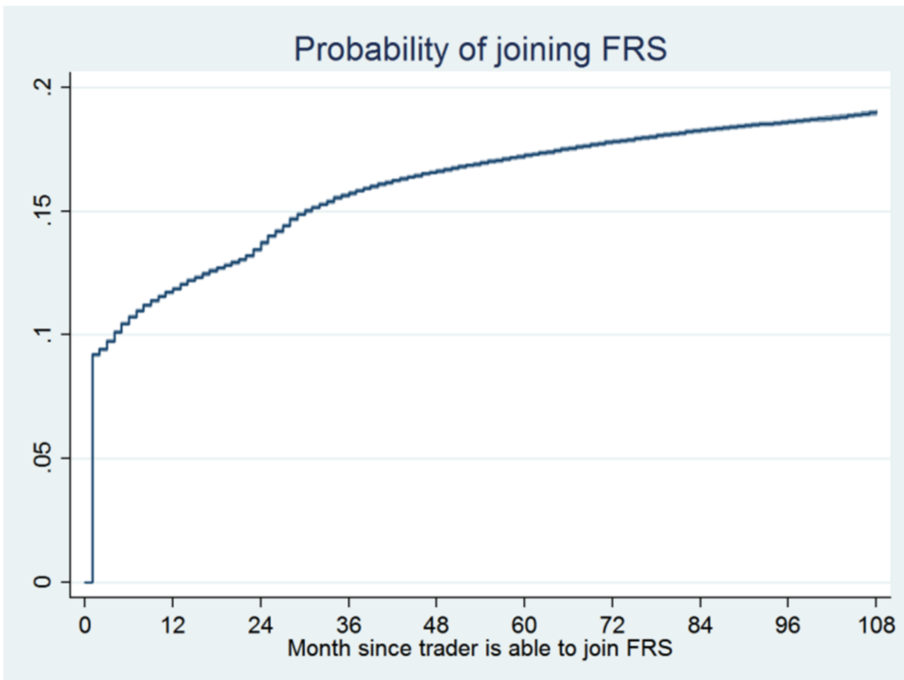

Figure 1.1: Probability of joining FRS on or before analysis time

Notes: Figure shows Kaplan-Meier nonparametric estimate of probability of joining FRS on or before analysis time. Analysis time measures the time since traders had the option of joining FRS. The zero corresponds to date of VAT registration for traders registering after April 2002, when FRS is available, but is fixed at April 2002 for those already registered when FRS was introduced. Traders who were VAT registered at the time of FRS introduction in April 2002 had the option of joining FRS for 109 months at the end of sample on April 2011. Figure uses trader-level dataset with 1,803,179 traders. 165,967 join FRS as soon as they have the option to do so ($t = 1$) and 129,318 join after this time until the end of analysis time. Data includes all traders who were observed to be eligible for FRS or were on FRS at least once during the sample.

Figure 1.2: Composition of FRS inflow and outflow

Notes: Figure uses returns-level dataset and follows traders overtime. The inflow figures are based on last year status of traders observed on FRS during 2005-2010 financial years (148,332 average number of traders on FRS in this period). The outflow figures are based on what happens to traders on FRS during 2004-2009 financial years in the next year (130,815 is the average number of FRS traders during this time). New VAT registrations are traders within the first twelve months of VAT registration.

Figure 1.2 shows composition of traders joining and leaving FRS. On average 81 percent of current FRS traders remain on FRS and only 3 percent revert to normal VAT in the next year. 16 percent of current FRS traders also exit data which seems normal given the small size of eligible traders. On the inflow side, new VAT registrations comprise a significant addition to FRS. While 71 percent of current FRS traders were on FRS in the last year, 23 percent are coming from new registrations as opposed to 6 percent from existing VAT traders. In summary, figure 1.2 shows FRS is close to an absorbing state and most of the additions are from newly registered traders.

Table 1.3 shows summary statistics for three sub-samples: a) VAT traders below FRS continuation threshold of £225,000, b) FRS traders, and c) eligible VAT traders with gains from FRS (next section). The top panel lists tax variables while the bottom panel shows indicator variables. Average FRS trader has a similar turnover to average eligible gainer but they are smaller than average VAT trader. FRS traders pay higher net VAT compared to VAT traders but slightly less than eligible gainers. Eligible gainers also have much lower average inputs and input VAT compared to VAT traders. This is consistent with the intuition that FRS is beneficial for firms using fewer inputs. FRS traders report inputs only if they purchase capital goods with a value greater than £2000 or under special circumstances. This pulls down average inputs and input VAT for FRS traders.

Incorporated businesses, with a share of 70%, dominate the population of FRS traders. They have a more balanced share among VAT traders and FRS gainers (43 and 48 percent respectively). Both sole proprietors and partnerships are under-represented in FRS. This suggests that sole proprietors and partnerships are less

cross-sectional estimates of fraction on FRS.

Table 1.3: Summary statistics

Variables	A. VAT traders (sales≤225k)			B. FRS traders			C. eligible FRS gainers		
	Mean	S. Dev.	Median	Mean	S. Dev.	Median	Mean	S. Dev.	Median
Gross Outputs	82,543	61,268	71,711	76,197	82,671	68,393	75,548	45,913	70,916
Output VAT	9,463	8,715	7,306	8,758	9,592	7,809	10,903	6,679	10,211
Gross inputs	62,746	161,909	37,836	4,805	32,542	0	25,068	46,783	12,967
Input VAT	6,335	18,303	3,464	360	2,559	0	2,161	2,889	1,119
Net VAT	3,190	18,837	2,818	8,407	9,323	7,545	8,821	5,672	8,045
% sole proprietor	37.8			23.8			35.7		
% incorporated	43.4			69.8			48.1		
% partnership	18.9			6.4			16.2		
% EC Trader	21.7			9.6			20.6		
%Group registrations	0.3			0			0		
% Partial Exempt	1.4			0.2			0.91		

Notes: Based on 2004-10 data and the working sample shown in 1.2. The number of observations are 5,822,956 for VAT traders, 964,588 for FRS traders, and 1,049,218 for eligible gainers. 255,215 of FRS returns show non zero input and input VAT but some of these relate to traders who are submitting a mix of FRS and VAT return. There are 720,856 pure FRS returns (12 months on FRS) and 85,476 of these report a non-zero input VAT (12 percent) with an average input VAT of £2,125. EC Trader counts both former and present traders with EU transactions. Partial exempt counts all traders with some form of partially exempt supplies. Group registration shows fraction of divisional and representative registration.

likely to utilize FRS opportunity[15]. The last two rows show the fraction of group registrations and partially exempt traders are very small among VAT businesses. Group registrations are ineligible for FRS and hence the zeros under panel B and C. It is also less likely that partially exempt traders benefit from FRS justifying smaller numbers under panel B and C.

[15] One likely reason for this could be the fact that a higher proportion of incorporated businesses use tax preparators and hence are more likely to get tax saving recommendations from their specialized agents. National Audit Office (2010) reports that 78 percent of corporation tax returns and 43 percent of VAT returns are submitted through tax agents. Incorporated businesses submit both corporation tax and VAT returns while sole proprietors and partnerships do not submit corporation tax returns.

1.4 FRS gainers

1.4.1 Calculation of FRS gains

In order to assess whether traders are choosing the minimum tax scheme I need to calculate tax liability under the alternative scenario. VAT traders report VAT liability (T_V in (1.1)). In order to calculate counterfactual FRS liability (T_F in (1.2)), I use traders' reported Standard Industry Classification 2007 (SIC2007) codes to determine the appropriate flat rate (τ_F) which is then multiplied by the sum of reported net sales and corresponding VAT. FRS gains are defined to be $T_V - T_F$. Similarly an eligible VAT trader is an FRS gainer if $T_V - T_F \geq 0$.

I give a brief overview of determination of flat rates and leave further discussions to appendix B where I also explain some complications in calculation of FRS gains. HMRC publishes applicable flat rates for 56 "categories of business" together with the list of associated "trade names". I match "trade names" to SIC2007 code descriptions from the Office of National Statistics (ONS) to form a mapping between reported SIC2007 codes and published flat rates. For example, ONS describes SIC2007 code of 70229 as "management consultancy activities (other than financial management)". This description matches with the FRS category for "management consultancy" with $\tau_F = 12.5$ percent during 2004-07. Using this manual matching, I assign flat rates to 78 percent of eligible traders. The largest sectors left out are construction and some retail sectors because reported SIC2007 codes map to several flat rates.

FRS traders make an active decision when joining FRS; therefore it is unlikely that they lose out from the scheme. Comparing FRS and VAT liabilities for FRS traders could shed light on importance of other issues (e.g. compliance cost savings) that might influence the joining decision. For example, observing some traders remain on FRS despite having a lower VAT liability suggests that they get compliance cost reductions under FRS. Unfortunately, FRS traders only report gross sales (S_g), and corresponding FRS liability (T_F), making it impossible to *calculate* counterfactual VAT liability (T_V)[16]. I must *estimate* VAT liability for FRS traders which requires estimation of τ_V and v in (1.1). Absence of enough observable characteristics renders regression based estimation of gains ineffective and therefore, I exclude FRS traders. Table 1.4 summarizes the focus of this paper. FRS traders are left out but VAT

[16] To be more precise FRS traders report FRS turnover which in some cases might differ from gross sales (see appendix B). Also notice that the less demanding reporting requirement is the main source of compliance cost saving under FRS.

Table 1.4: FRS gainers studied

	FRS gainer	FRS loser
FRS traders	$\hat{T}_V - T_F \geq 0$ Left out	$\hat{T}_V - T_F < 0$ Left out
VAT traders	$T_V - T_F \geq 0$ Focus of paper	$T_V - T_F < 0$ Analyzed

traders are analyzed. The main message of the paper is, however, about the group of VAT traders who are observed to gain from FRS.

1.4.2 FRS gainers characteristics

Table 1.5 shows aggregate number of FRS gainers. Column (1) reports the number of eligible VAT traders under investigation (assigned τ_F). On average 26 percent of 573,347 eligible traders are FRS gainers but the percentage of gainers drops from 28 to 23 percent during the sample (column (2))[17]. Columns (4) shows percentage of FRS gainers who join FRS in the following year. On average only 3 percent of FRS gainers join the scheme in the following year and there does not seem to be a clear time trend. However, 70 percent of gainers remaining on VAT (do not exit or join FRS) still gain from the scheme in a consecutive year (column (5)). Column (6) checks the robustness of fraction of gainers by setting τ_F to the maximum applicable rate in each financial year. Even using this conservative approach 12 percent of eligible traders are observed to gain from FRS. This, to some extent, alleviates concerns about errors in assignment of flat rates. Therefore, FRS gains seem to be persistent but majority of gainers are not responsive and remain on normal VAT.

To compare size of gainers and current FRS traders figure 1.3 plots sales distribution (frequency) for the two groups. Both distributions are right-skewed suggesting FRS is suitable for small businesses and is inline with HMRC's design of the scheme as a small business program. The number of FRS gainers is almost similar to FRS traders for low levels of sales, but the ratio of gainers to FRS traders increases after £100,000 annual sales. Around the joining threshold (first vertical line) there are three gainers for each FRS trader. Figure 1.3 also sheds light on gainers beyond

[17]The decline in the fraction of FRS gainers could be a result of information diffusion over time (in 2004 the scheme was in place only for two years). The flip side of this decline is a secular increase in fraction of traders on FRS which is reported in column (6) of table 1.2.

Table 1.5: FRS gainers among eligible VAT traders

year	FRS eligible (assigned τ_F)	% FRS gainer	# FRS gainer	% Joined FRS	% FRS gainer next year	% gainer (max τ_F)
	(1)	(2)	(3)	(4)	(5)	(6)
2004	618,810	28%	172,421	3.5%	72.0%	14%
2005	635,295	27%	174,639	3.9%	69.0%	14%
2006	596,803	27%	161,942	2.8%	71.0%	14%
2007	602,626	27%	165,170	3.6%	69.9%	12%
2008	503,013	25%	125,155	1.9%	68.0%	11%
2009	523,772	24%	124,967	2.8%	68.5%	7%
2010	533,107	23%	124,924	-	-	9%
Average	573,347	26%	149,888	3.1%	69.7%	12%

Notes: Column (1) shows number of VAT registered traders who are eligible for FRS and whom I was able to assign a flat rate to and calculate counterfactual FRS liability. Column (2) shows the percentage of FRS gainers out of column (1) traders, i.e. VAT traders with FRS liability equal or smaller than reported VAT liability. Column (3) is the number of gainers, i.e. column (2) multiplied by column (1). Column (4) follows the population of FRS gainers to the next period and reports the fraction joining FRS. Column (5) reports the fraction of FRS gainers gaining in the following year. This fraction is calculated as the number of second year gainers divided by all first year gainers who remain on normal VAT, i.e. do not exit and do not join FRS. Column (6) uses the maximum applicable flat rate (not the ones I have assigned) and reports the fraction of VAT traders with non-negative tax gains from joining FRS.

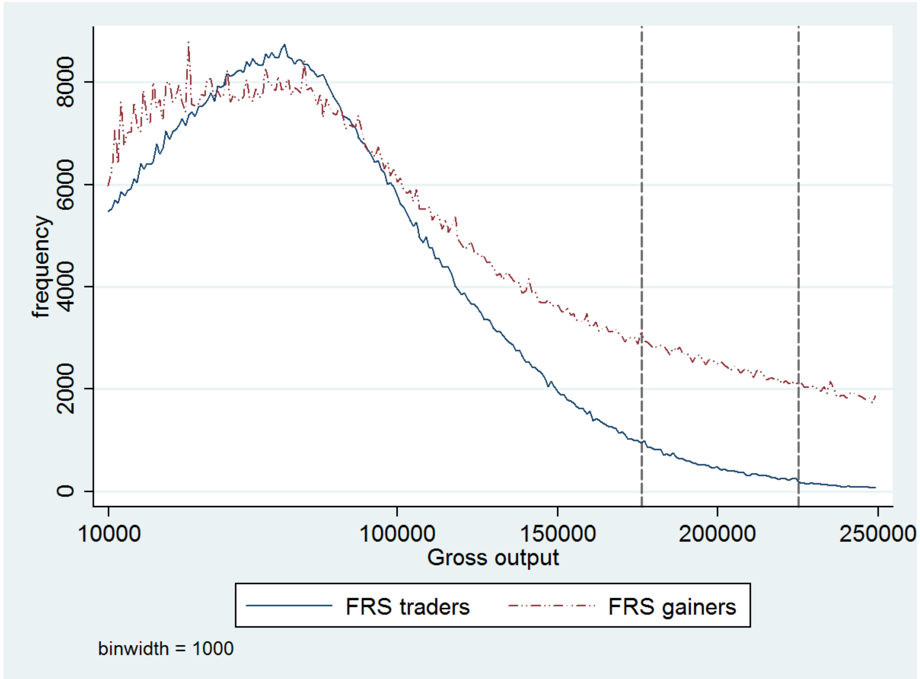

Figure 1.3: Sales distribution for FRS traders and FRS gainers

Notes: Figure shows number of traders within bins of gross output for FRS gainers and FRS traders. The sample is the returns-level dataset and includes all VAT returns submitted while traders are observed on FRS and all returns for FRS gainers during 2004 - 2010 financial years. The sample here is bigger than the one reported in the tables because it includes traders above the FRS eligibility thresholds depicted by the vertical lines. I, however, exclude traders who are ineligible based on observable non-turnover criteria. The first vertical line shows FRS joining eligibility threshold $(150,000\times(1+0.175) = £176,250$ during 2004-2010) while the second vertical line shows FRS continuation eligibility threshold ($£225,000$ during January 2004 until January 2011).

the joining eligibility. As we have seen in section 1.2 the joining threshold is not binding and traders above this threshold could in effect join the scheme. I ignore this possibility in table 1.5 but figure 1.3 shows there is a significant mass of traders who could potentially gain in this region.

In the remaining part of this section I establish four empirical facts about the population of FRS gainers:

Fact 1 Very few FRS gainers join FRS over time. 3 percent join in the following year and the estimated joining probability 6 years after gaining is 10 percent.

Fact 2 Gains are persistent. Gaining in the last period increases the probability of gaining by 62 percentage points after controlling for SIC2007 and year dummies. 34 percent of gainers are observed to gain (or join FRS) during all years they show up in the data.

Fact 3 Size of FRS gains are not small. Median gainer could save 12 percent on VAT liability by joining FRS. 92 percent of gainers have a gain of £100 or more and 46 percent gain £1000 or more.

Fact 4 Gainers are concentrated in a few services sectors (consultancy and personal services)[18].

Fact 1: Few gainers join the scheme

Figure 1.4 plots Kaplan-Meier non-parametric estimate of (cumulative) probability of joining FRS on or before the indicated number of months since traders are first observed to gain. Similar to table 1.5, 12 months after gaining, probability of joining is about 3 percent. Interestingly, the likelihood of joining FRS shows a very gentle increase over time and reaches 10 percent after 6 years (72 months). The gradual increase in uptake of FRS suggests a potential role for learning and inertia which I discuss in more detail in section 1.6.

Figure 1.5 looks at the percentage of gainers eventually joining FRS. X-axis shows the number of years traders are observed to gain. Figure 1.5a considers all eligible traders and plots the fraction of traders in each x-axis category that are observed on FRS at any time during the sample. 13 percent of one-year gainers and 12 percent of two year gainers are ever observed on FRS while only 8 percent of traders gaining for more than two years join the scheme. Interestingly, 4 percent of traders who never gain join the scheme. While this is one third of the fraction of two year gainers who join the scheme, it suggests my calculations are unable to uncover gains for these traders.

Splitting the data into traders with different lifespans[19] in figure 1.5b confirms the same pattern but also shows the percentage of gainers joining FRS is the highest among traders who are present in the full 7 years of my sample: almost 20 percent of one and two year gainers join FRS. In contrast, around 15 percent of one and two year gainers from 5 and 6-year traders join the scheme. The patterns observed in this figure could be consistent with inertia (sluggish responsiveness) and learning.

[18] I believe unobserved ineligibility is unlikely to overturn any of these facts. As discussed under section 1.2, some of the unobserved eligibility criteria are likely to be more binding for FRS losers and therefore would strengthen my results (e.g. past VAT convictions or uptake of alternative VAT accounting schemes). The only unobserved criterion that might pose a challenge is being associated with another business. I have no available information on business associations and assume the share of associated businesses is not disproportionately high among FRS gainers.

[19] This is defined as the number of years traders show up in my data.

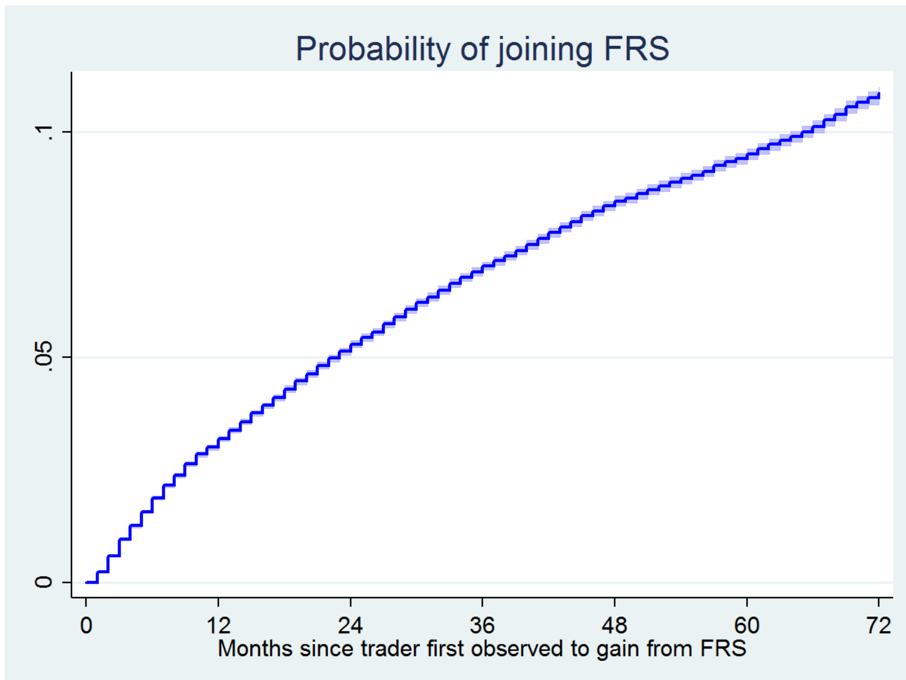

Figure 1.4: Probability of joining FRS versus months since first gained

Notes: Figure shows Kaplan-Meier non-parametric estimates of the probability of joining FRS on or before analysis time. The zero of analysis time (x-axis) corresponds to end of first financial year traders observed to gain from FRS. Data used here is the trader-level dataset and includes all traders who were observed to be eligible for FRS and gained at least once during the sample period. Traders exiting the data before joining FRS are censored after exit. Figure uses the trader-level dataset and estimates joining probability from the sub-sample of 457,297 traders who gain at least once during their lifetime.

Observing one and two year gainers for longer (higher lifespan traders) increases the joining probability. Gaining for second years rather than one year also increases joining probability for 7-year traders (but not for traders with shorter lifespans).

Fact 2: Gains are persistent

Figure 1.6 looks at the persistence of FRS gains across sales levels. The solid line shows the unconditional probability of being an FRS gainer is first increasing but quickly reaches a plateau after around £30,000 annual sales. The dashed line shows the probability of remaining a gainer conditional on being a gainer in the previous year. While this figure confirms the earlier fact that the conditional probability is much higher than the unconditional one (table 1.5), it reveals lower persistence of gains for very small traders and slightly higher than 70 percent conditional probability of gains for larger traders. Interestingly the conditional probability also reaches a plateau after £30,000 annual sales and there is little change in persistence of gains across sales levels after this point.

Figure 1.7 plots distribution of number of years gaining conditional on gaining once. Figure 1.7a shows the fraction of gainers that gained for less than 50 percent, exactly 50 percent, more than 50 percent and exactly 100 percent of the times they submitted returns. 34 percent of FRS gainers gain for all years while only 30 percent gain less than 50 percent of the times[20]. Figure 1.7b shows separate histograms for traders with different lifespans. For almost all lifespans the highest share is for traders gaining during their entire lifespan (far right dots for each curve). In summary these figures show a considerable share of traders gain during all years in the data, while many others have multiple years of gaining.

Fact 3: Gains are not small

Figure 1.8 plots the distribution of FRS tax gains for eligible VAT traders. The gains distribution has a mode at zero with 4.8 percent of the mass falling between £-100 and £100 FRS gains. This is due to HMRC's targeting of flat rates to make the average traders indifferent between FRS and VAT. A closer look at FRS gainers,

[20]In this figure, I have assumed traders who join FRS after x-year of gaining continue to gain while on FRS and put them in the 100 percent gains bin. Dropping the traders who join will change the percentages to 33, 14, 25, and 28 percent for less than 50, exactly 50, more than 50, and 100 percent bins respectively.

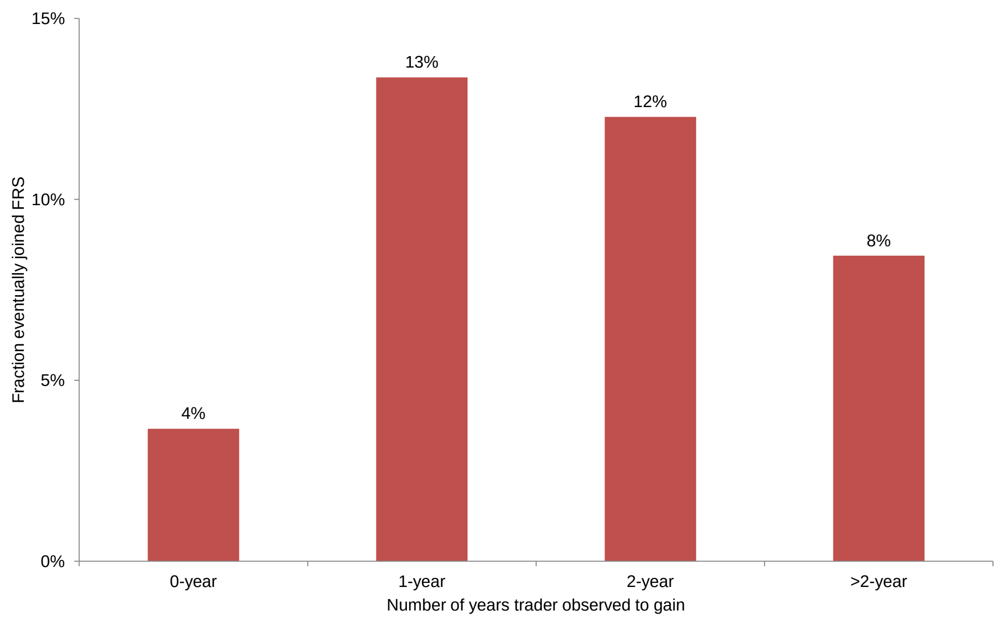

(a) Combined all lifespans

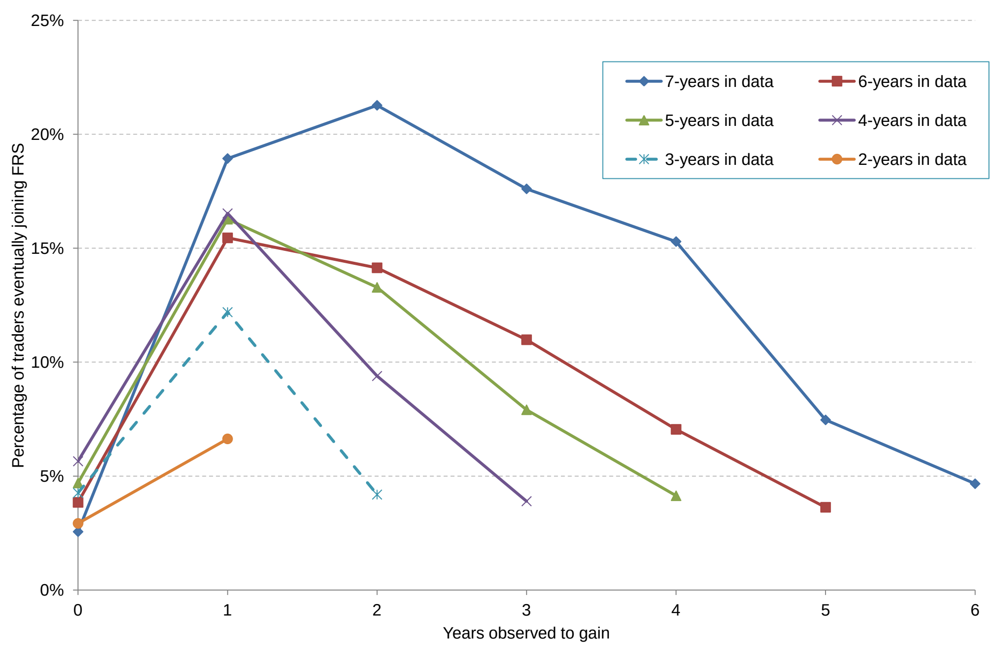

(b) Separately for different lifespans

Figure 1.5: Fraction of traders eventually joining FRS after x years of gaining

Notes: Figure shows the fraction of traders ever observed on FRS among different sub-samples of traders. The figures are based on trader-level dataset where there is one observation for each trader and I record the number of years gaining and the number of years present in the data. This graph uses the pool of unique traders who are present at least for two years in the data. Figure (a) reports percentage of joining traders for traders gaining never, one year, two years, and more than two years during their lifetime. Figure (b) reports percentage joining for traders gaining a given number of years separately for different lifespans. Maximum lifespan is seven years but following trader over time results in at most 6 years of gains (horizontal axis) for those who could join the scheme in the seventh year.

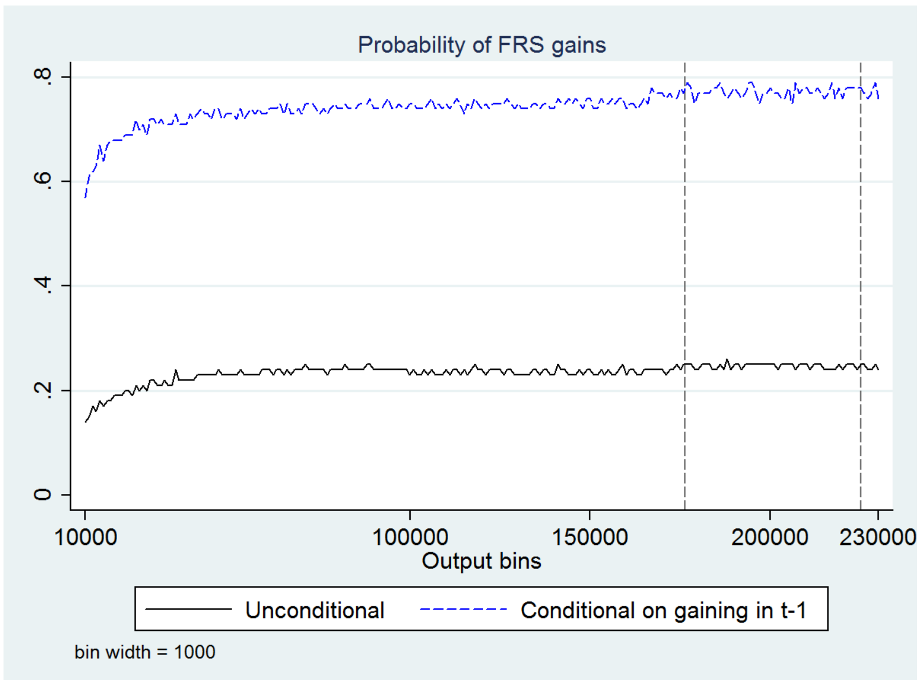

Figure 1.6: Unconditional and conditional probability of FRS gains

Notes: The solid line shows unconditional probability of being an FRS gainer within bins of gross output, i.e. the ratio of gainers to FRS eligible traders within bins. Dashed line shows the probability of gaining from FRS conditional on being a gainer last year, i.e. the ratio of traders gaining for a second year among last year gainers who remain on VAT (do not join FRS or exit). The sample here is bigger than the one reported in the tables because it includes traders above the FRS eligibility thresholds depicted by the vertical lines. I, however, exclude traders who are ineligible based on observable non-turnover criteria. The first line shows FRS joining eligibility threshold $(150,000 \times (1+0.175) = £176,250)$. The second line shows FRS continuation eligibility threshold $(£225,000)$. Figure uses returns-level dataset and combines all years.

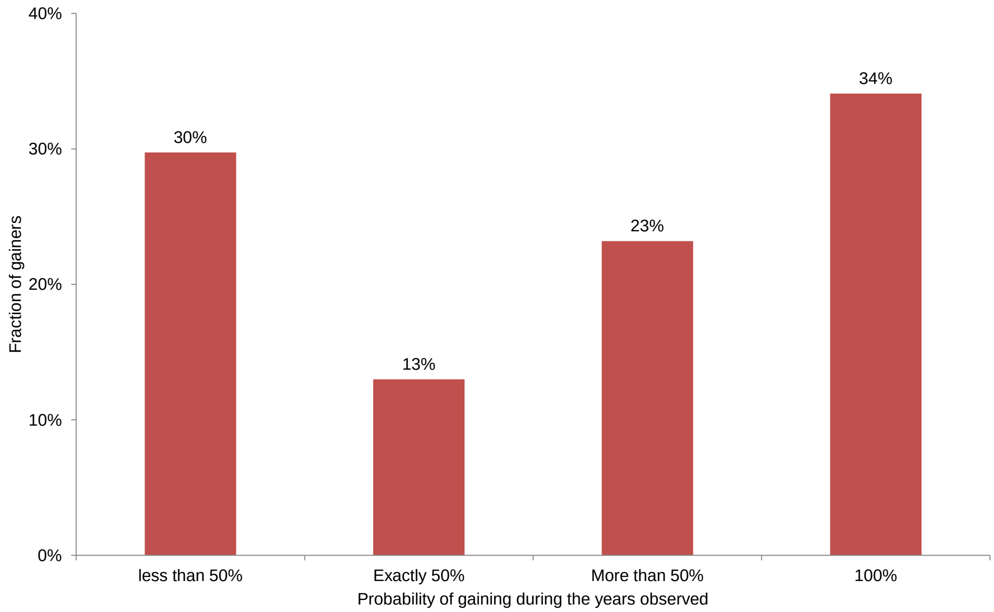

(a) Combined histogram

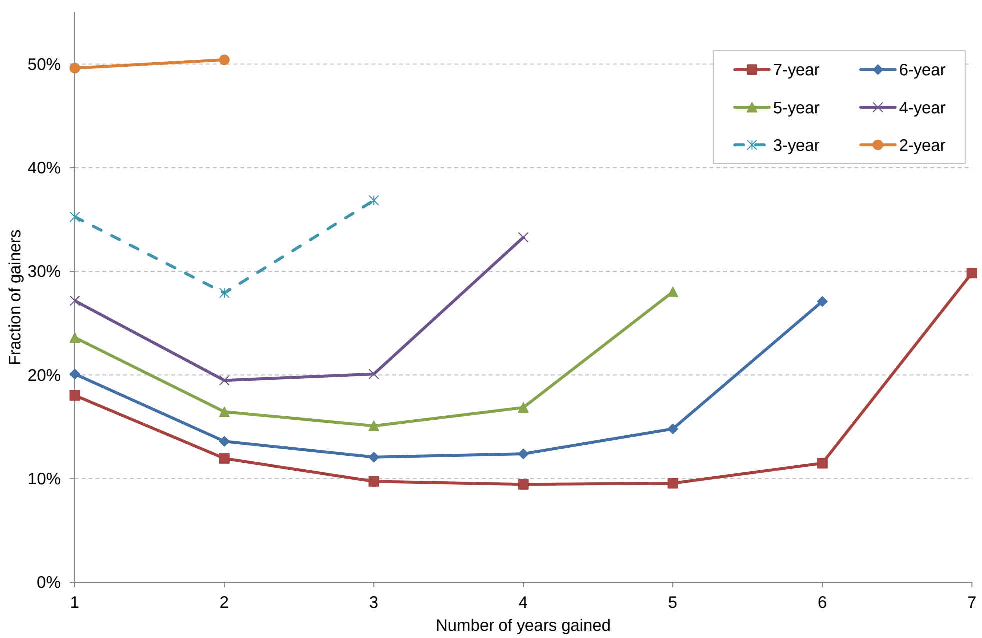

(b) Separate histograms for different lifespans

Figure 1.7: Distribution of number of years gaining conditional on gaining once

Notes: Figure shows distribution of the number of year gaining conditional on gaining once. Traders who joined FRS after gaining over certain years are assumed to continue to gain from FRS and hence are put in all year gaining bin. This graph uses the pool of 402,894 unique traders who are observed to gain at least once and are present at least for two years in the data. Figure (a) plots share of gainers that fall into four categories of gaining less than 50 percent, exactly 50 percent, more than 50 percent, and exactly 100 percent of the times they submit returns. Figure (b) shows separate histograms for traders with different lifespans and instead shows the distribution of number of years (rather than percentages).

i.e. the positive tail, reveals 92 percent of gainers have a gain of £100 or more and 46 percent gain £1000 or more.

Gains distribution reveals great asymmetry between gains and losses. Size of losses could potentially be much larger than gains: the first percentile of gains distribution shows a loss of £27,800 while the ninety ninth percentile shows a modest gain of £4,800. This is also in line with a high proportion of FRS losers (table 1.5 reports 74 percent of eligible traders lose out from the scheme). One might expect that given the way HMRC sets flat rates, this ratio should be closers to 50 percent[21]. But it should be noted that the gains distribution excludes the traders currently on FRS and includes eligible zero (and reduced) rated traders who would incur huge losses under FRS. I have no reliable information about how exactly flat rates were calculated but it seems HMRC excluded zero-rated traders from this calculation (see discussion of figure 1.10 too). Furthermore, FRS traders are likely to have had gains from FRS and exclusion of such traders in the gains distribution would shift the ratios in favor of losers.

In order to get a better sense of size of gains, figure 1.9 looks at FRS tax gains as a percentage of reported VAT liability across sales levels. The figure plots medians of relative tax gains distribution separately for FRS gainers (above zero) and losers (below zero) within gross sales bins of £1000. The top part shows fairly stable and non-negligible tax gains for FRS gainers. Gainers with annual sales between £9500 and £10500 (first bin) see a median reduction of 17 percent in their tax liability upon joining FRS. The median gain decreases to 12 percent for larger gainers but remains stable at this level. Perhaps not surprisingly, the bottom part confirms FRS losers incur large tax losses if they join the scheme. Median FRS losers with less than £50,000 annual sales would see an increase of 150 percent in their tax liability should they join FRS. This loss reduces to 100 percent for higher annual sales.

Fact 4: Gains are concentrated

To see the type of activities benefiting from FRS, table 1.6 lists ten sectors with highest number of FRS gainers. These sectors comprise 51% of all FRS traders and 41% of all FRS gainers. This table shows FRS is suitable for a concentrated number of sectors. The list includes management consultancies, computer consultancies,

[21]Obviously, this assumes mean and median of VAT liability distribution within flat rate categories are the same. If the VAT liability distribution is skewed, then targeting average VAT liability within sectors would not necessarily make 50 percent of eligible traders gainers.

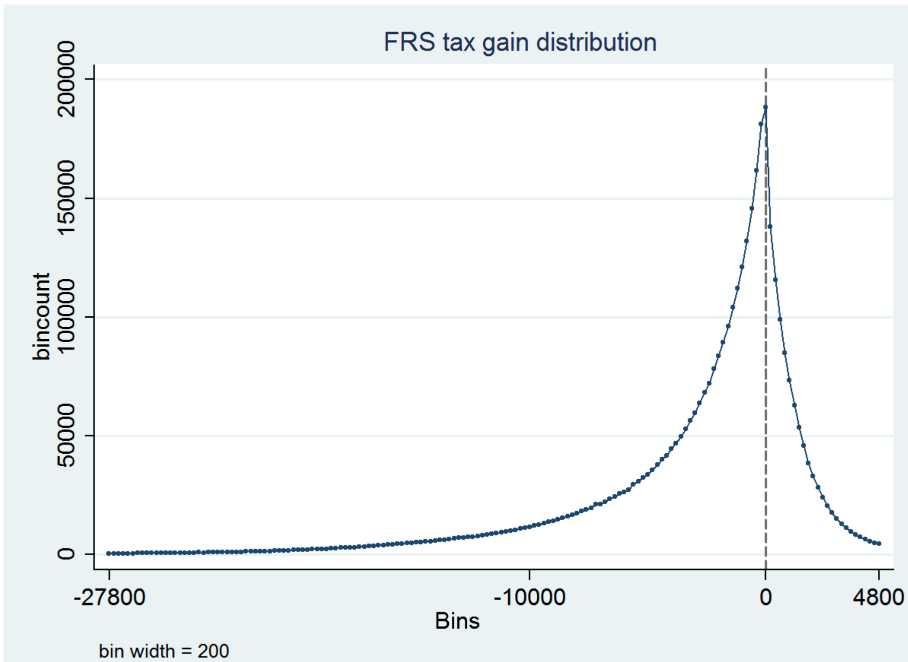

Figure 1.8: Distribution of FRS tax gains for gainers

Notes: Figure shows distribution of FRS tax gain for current VAT traders, positive numbers show gains from switching to FRS while negative numbers show losses. The figure uses returns-level dataset and combines all available years of data. Sample size is the sum of observations in column (1) of table 1.5, i.e. eligible VAT traders assigned a flat rate. Figure restricts to the first and ninety ninth percentiles of the gains distribution and removes traders with less than £1000 annual turnover (similar figures obtained without this or with £10,000 threshold.).

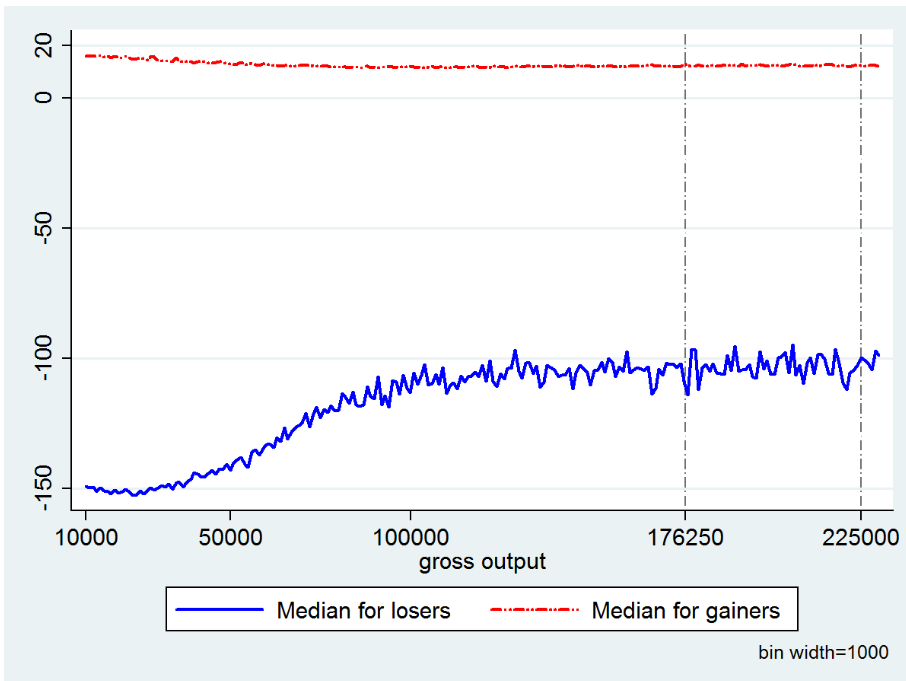

Figure 1.9: Medians of FRS gains as a percentage of VAT liability

Notes: Figure splits the FRS tax gain distribution at zero and plots medians over gross output bins for FRS gainers and losers separately. Solid line show medians of FRS gains for FRS losers and dashed line represent medians of FRS gains for FRS gainers. The sample here is bigger than the one reported in the tables because it includes traders above the FRS eligibility thresholds depicted by the vertical lines. I, however, exclude traders who are ineligible based on observable non-turnover criteria. The first line shows FRS joining eligibility threshold ($150,000 \times (1 + 0.175) = £176,250$). The second line shows FRS continuation eligibility threshold ($£225,000$).

business support activities, and take away food shops. Interestingly, most of these sectors have flat rates close to the high end of the range of applicable rates. Gains seem to be more persistent for these sectors: 77% of gainers who remain on VAT continue to gain in $t+1$ (compared to 70% for all gainers in table 1.5). Conditional median of gains (columns (6) and (7)) reveals non-negligible potential gains from joining FRS.

Figure 1.10 generalizes the patterns in table 1.6 by looking at distribution of FRS traders, gainers, and eligible VAT traders across flat rate categories. Dots in the figure show proportion of the specified group that falls in the given flat rate category. For example, the two far right solid blue circles show that the last two flat rate categories contain 31 and 26 percent of all FRS traders. This figure shows proportion of eligible traders, FRS traders, FRS gainers, and the flat rate percentages show positive correlations[22]. In other words, it seems there is a high concentration of FRS traders, gainers, and eligible traders in the higher flat rate categories. This pattern is partly due to the concentration of total observations in these categories. The three most populous flat rate categories are those with flat rate percentages equal to 6, 12.5, and 13 with a respective share of 17, 14 and 13 percent of total observations (eligible plus FRS traders). All other sectors have less than 9 percent of traders. The other factor that explains this positive correlation is the positive correlation between FRS traders and FRS gainers (both as a % of eligible traders) within 5-digit SIC2007 codes. Sectors with a higher percentage of FRS traders also have a higher percentage of FRS gainers[23].

This counter-intuitive pattern seems to be an artifact of HMRC's conservative approach in setting the flat rate percentages. Using returns submitted by FRS eligible VAT traders between 2004 and 2007 financial years, I calculated the average of net VAT to gross sales within 5-digit SIC2007 codes, restricting to traders with a positive net VAT. This average ratio should approximate the statutory flat rates based on HMRC guidance on calculation of flat rates. But when I compare calculated flat rates to statutory rates, I find that some sectors have statutory rates that are higher than the calculated ones[24]. These are mostly sectors with majority zero-rated traders

[22] The correlation coefficient between proportion of FRS traders and FRS gainers is 0.76; for FRS traders and eligible traders it is 0.36; for FRS gainers and eligible traders it is 0.70; for flat rate percentages and FRS traders it is 0.62; and for flat rate percentages and FRS gainers it is 0.54.

[23] Notice, this is the share of FRS traders and gainers from all traders in a given 5-digit SIC2007 code which is different from the share of population falling under each sector. Figure 1.10 is an aggregated version of the latter while table 1.6 is showing some evidence based on the former.

[24] The fact that some traders are on FRS during the time I am calculating the flat rates implies that calculated rates underestimate the statutory ones. The implicit assumption here is that this

Table 1.6: Ten sectors with highest number of FRS gainers

Sector	τ_F (2004-7)	% FRS	% gainer	% gainers join FRS in $t+1$	% gainers gaining in $t+1$	Conditional Median of gains (£)	Conditional Median of gains % VAT
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Management consultancy	12.5	35	36	5	74	522	7.5
Renting and operating of Housing Association	12	3	52	0	85	642	15
Computer consultancy	13	45	36	7	79	643	7.4
Other personal service activities	10	13	31	2	77	849	15
Other business support service activities	11	17	30	3	79	795	14
Other engineering activities	12.5	48	35	6	76	530	7.3
Take away food shops	12	31	39	5	84	808	7.2
Freight transport by road	9	17	29	1	67	461	8.5
Maintenance and repair of motor vehicles	7.5	10	29	2	76	841	13
Artistic creation	11	20	34	3	73	516	11

Notes: Table uses observations from 2004-2010 financial years. Column (1) reports the assigned flat rate during 2004-2007 financial years. Column (2) shows the percentage of FRS traders out of all eligible traders in each sector. Column (3) is the fraction of eligible VAT traders who gain from FRS in each sector. Column (4) is the fraction of FRS gainers who join FRS in the following period. Column (5) reports two year gainers as a percentage of last year gainers who remain on VAT and are still eligible for the scheme. Column (6) is the median of current FRS tax gains for the population of FRS gainers in the last year who remain on VAT. Column (7) is the same conditional median as in column (6) but for tax gain as a percentage of VAT liability.

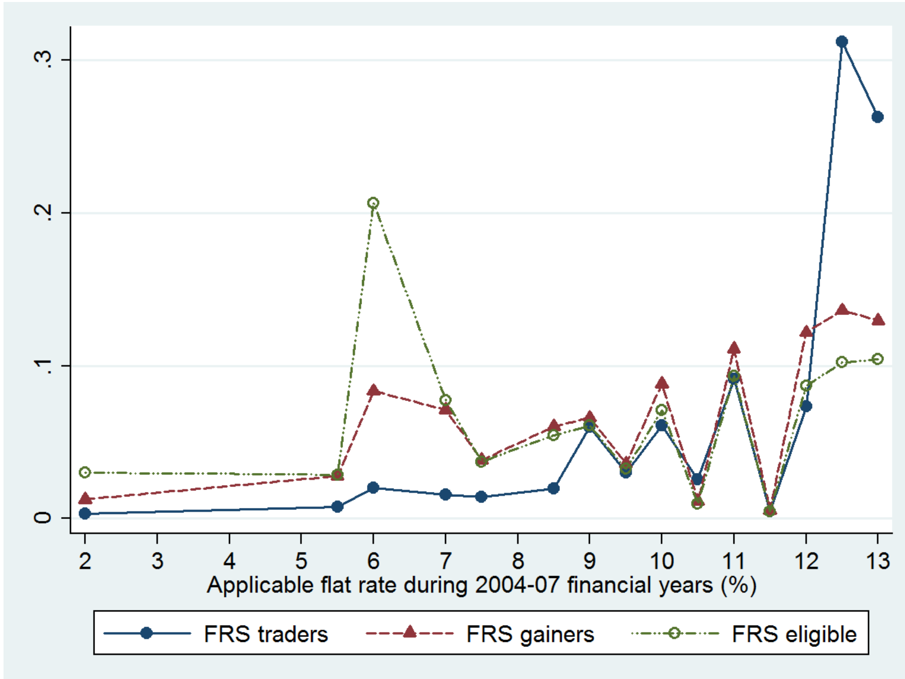

Figure 1.10: Distribution of FRS traders, FRS gainers, and eligible VAT traders across flat rate categories

Notes: Figure shows distributions across flat rate categories. Solid line shows fraction of FRS traders that fall in each flat rate category, dashed line shows fraction of FRS gainers in each FRS category, and dot-dash line shows the fraction of eligible VAT traders within each flat rate. Flat rate categories are based on the applicable rates during 2004-7 Flat rates range from 2 to 13.5 percent during 2004-2007, but 13.5 percent is excluded as I could not assign it. There are, therefore, 15 distinct flat rates. The sample is the returns-level dataset and covers 2004-2010 financial years.

or those with high input use (low share of value added) that feature a large number of traders with negative net VAT (repayment traders). Such sectors are unlikely to have a high number of FRS gainers if the calculation of flat rates ignores the repayment traders. On the other hand, sectors with mostly standard-rated traders (e.g. management consultancy) would receive a statutory flat rate closer to the sectoral average and hence are more likely to have a higher number of FRS gainers and FRS traders.

---

underestimation would not be able to explain the observed discrepancy between calculated and statutory rates. To justify this assumption, I note that in 2007 only 17 percent of eligible traders were on FRS. Furthermore, if I repeat the calculations restricting to only 2004 (when only 9 percent of traders were on FRS) the same pattern emerges between calculated and statutory flat rates.

1.5 Uncertainty

Traders decide to join FRS before gains are realized. Assuming risk neutrality, basic economic theory suggests they should join FRS when expected after tax profits are greater under the scheme. So far, I have shown some traders are observed to gain. But this is not necessarily equivalent to expected gains. Therefore, inaction of identified FRS gainers could simply be an artifact of expected FRS losses, not sub-optimal choices. In this section I first show that observed FRS gains influence the joining decision of a sub-sample of traders. Then I reinforce fact 2 from the previous section on persistence of FRS gains to show that gaining once is a strong signal of expected gains. Finally, I discuss implications of risk averse preferences and consider a few features of the scheme that might alleviate concerns.

Figure 1.11 shows that the probability of joining FRS rises sharply around zero last year gains. In other words, a visibly higher proportion of FRS gainers join the scheme compared to FRS losers. This pattern confirms that calculated gains are not irrelevant and influence the joining decision of a sub-sample of traders. Under the assumption that the responsive traders are not making a mistake themselves, I can conclude that observed gains are equivalent to expected gains for these traders. However, this figure might be less useful in ruling out uncertainty for the whole sample because the responsive traders might have different risk preferences or face lower levels of uncertainty.

To show that observed FRS gains signal expected gains I complement the evidence on persistence of FRS gains by looking more closely at the distribution of FRS gains conditional on past gains. Figure 1.12a plots twenty fifth, fiftieth (median), and seventy fifth percentiles of current FRS gains for traders falling in £500 bins of last year gains. The gains distribution shows high degree of serial correlation. The whole distribution of FRS gains shifts to the right for traders with higher past FRS gains. The comparison of the median line (solid black) with the 45 degree line (one-to-one dependence of gains over time) shows that the median gains and losses are slightly less than the absolute value of last year's tax gain. But size of the gains are quite comparable. For example the median gains for traders with last year tax gains between £5750 and £6250 is equal to £4800 and the 75 percentile is £6,000. The twenty fifth percentile of gains distribution is positive for traders with last year gains falling in $[750, 1250)$ bin or beyond.

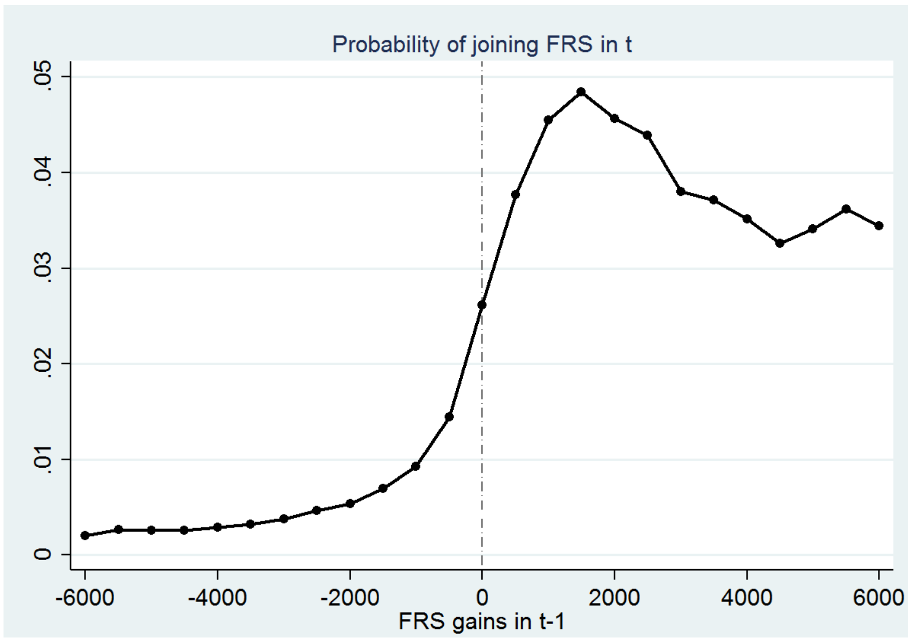

Figure 1.11: Probability of joining FRS conditional on last year gains

Notes: Figure depicts probability of joining FRS in year t conditional on falling in a given bin of FRS tax gains in year $t-1$. This is the ratio of the number of traders joining FRS to the number of traders remaining on VAT in year t within FRS tax gain bins of year $t-1$. Sample includes all traders who are eligible to join FRS during 2004-2009 financial years and do not exit the data in the following year. Figures restrict to last years gains being between £-6000 and £6000 and categorizes traders in to £500 bins of last year gains.

Table 1.7: Linear probability model of FRS gains

Dependent Var: dummy for gainer	(1)	(2)
L.gainer	0.647	0.617
	(.0078)*	(.0068)*
SIC2007 dummies	NO	YES
Year dummies	NO	YES

Notes: Table shows coefficient estimates from an OLS regression of a gainer dummy on covariates. Gainer dummy is equal to one if trader is observed to gain from FRS in a given year and zero otherwise. Columns (1) and (2) control for trader's VAT registration time (two dummies capturing whether VAT registered between 1 April 2002 and 1 January 2004 and after 1 January 2004), ownership status (two dummies capturing incorporations and partnerships), Average log of gross output, average and standard deviation of FRS gains as a percentage of VAT liability, fraction of years trader was eligible for FRS, and a dummy for monthly returns. Column (2) further includes SIC2007 and year dummies and 9 dummies capturing the 2004 FRS density decile for registered outcode of trader. Standard errors are adjusted for SIC2007 clusters and shown in parenthesis. * shows if coefficient is significant at 1 percent level. The sample for both regressions is 3,449,070 returns during 2005-2010. It includes traders that were at least eligible for FRS once during 2004-2010 and drops sectors with less than 1000 observations during the 7 years of the sample. Notice the sample only includes traders NOT on FRS and those I could calculate whether they gain from being on FRS.

Figure 1.12b shows FRS gainers as a percentage of traders within bands of last year gains (the x-axis is the same as in figure 1.12a). The figure shows less than 20 percent of last year FRS losers become gainers. Perhaps more importantly percentage of gainers rises sharply right after zero to more than 70 percent. The fraction of gainers increases to 80 percent for traders gaining between £750 and £1250 during last year and continues to increase as the size of past gains increases.

To see the robustness of the persistence conclusion, table 1.7 shows the results of regressing an FRS gainer dummy on lag of the dependent variable and other covariates. The coefficient estimate of last year gains is highly significant and shows the probability of gaining from FRS increases by 65 percentage points for last year gainers. Controlling for sector and year dummies reduces the coefficient to 62 percentage points. While these regressions suffer from all sorts of endogeneity issues, they confirm that being an FRS gainer in the past is an important correlate of current gains even after controlling for sector and year dummies and other observable characteristics.

Both figures 1.12a and 1.12b and table 1.7 indicate very high persistence of FRS gains and therefore suggest observed gains are a signal of expected gains. To assess the relative size of gains, figure 1.13 looks at twenty fifth, fiftieth (median), seventy fifth percentiles, and mean of gains as a percentage of VAT liability. This figure

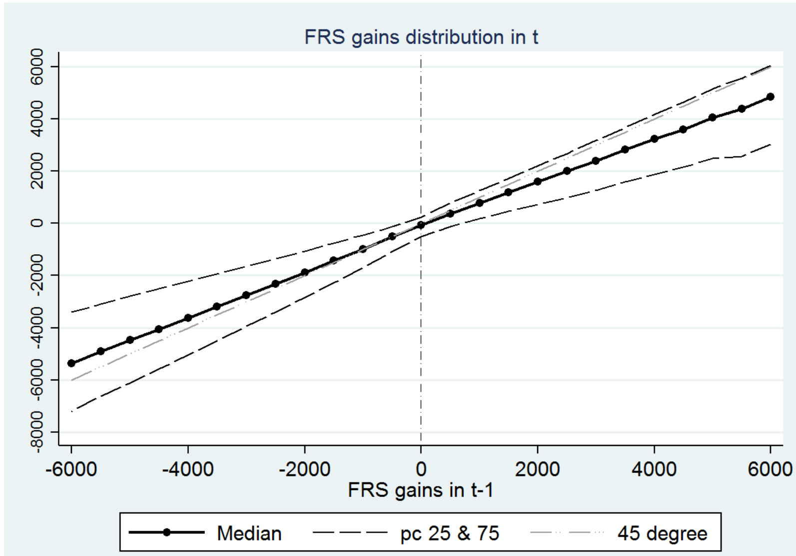

(a) Percentiles of tax gains in year t in bins year $t-1$ gains

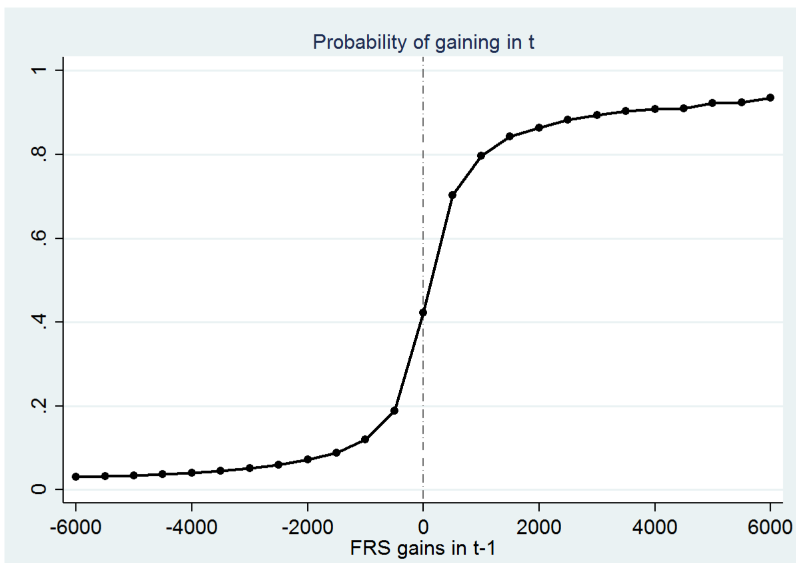

(b) Probability of FRS gains in bins of year $t-1$ gains

Figure 1.12: Impact of last year FRS gains on current gains

Notes: Figure (a) shows twenty fifth, fiftieth (median), and seventy fifth percentiles of FRS tax gain distribution in year t for VAT traders who were eligible for FRS in year $t-1$ within FRS tax gain bins in year $t-1$. Solid black line shows median and dashed gray lines show twenty fifth and seventy fifth percentiles. The solid gray line shows the 45 degree line. Panel (b) shows probability of having non-negative tax gains from FRS in year t conditional on being in a given bin of FRS tax gains in the previous year. This is the ratio of the number of traders gaining from FRS to the number of traders remaining on VAT in year t within FRS tax gain bins of year $t-1$. In both figures sample includes all traders who are eligible to join FRS during 2004-2009 financial years and do not exit the data in the following year. Figures restrict to last years gains being between £-6000 and £6000 and categorizes traders in to £500 bins of last year gains.

restricts to traders who have gained a year earlier and shows the dependence of the distribution on sales. Median gains are fairly stable at around 10 percent of VAT liability[25]. Seventy fifth percentile is also stable and shows 25 percent of last year FRS gainers save more than 20 percent on tax payment upon joining FRS. Twenty fifth percentile of the gains distribution is negative up until £40,000 annual sales but becomes positive for larger traders[26]. I have plotted mean of gains distribution to shed light on expected gains for FRS gainers. Assuming that gains distributions for last year gainers in the same sales bin are identical, the mean of FRS gains in each sales bin is equal to expected gains for traders in that bin. Therefore, I can use the realized gains for this group to back out expected gains for individual traders[27]. The mean coincides with twenty fifth percentile of FRS gains. For traders with gross sales less than £60,000, mean FRS gain is negative but traders larger than this level have positive mean. This suggests expected FRS gains for these traders.

So far I have assumed traders are risk neutral but would the same conclusions apply if traders are risk averse? Risk aversion could be important because as figure 1.13 shows the mean of FRS gains is almost 9 percentage points less than the median. In other words, there is a probability of incurring large losses even for last year FRS gainers. Therefore, while the mean of FRS gains is positive, the risk involved in opting in the scheme prevents risk averse traders from joining. This story suggests FRS liability is more volatile (involves higher uncertainty of after tax profits) compared to VAT liability. The summary statistics in table 1.3 shows coefficient of variation for net VAT is 0.64 for eligible FRS gainers (panel C) while it is 1.11 for FRS traders (panel B). This shows FRS traders face greater dispersion in distribution of tax liability compared to eligible gainers which is in line with the above reasoning. It is not, however, clear that this gap is entirely due to greater uncertainty of FRS liability. For example, coefficient of variation for gross sales shows a similar pattern. It is 0.61 for eligible gainers and 1.08 for FRS traders.

Two features of FRS alleviate some of the concerns arising from risk averse preferences. Infrequent large FRS losses (and higher volatility) could be a result of investments in capital goods. For example, management consultants might buy new computer systems every 5 years or take-away food shops might invest in new stoves

[25] The median gains as a percentage of turnover is also stable at around 1.5% (results not shown).

[26] 25th percentile fluctuates between a min of 0.2 percent and a maximum of 2.8 percent for traders larger than £40,000 with an average of 1.5 percent. This suggests on average 25 percent of last year FRS gainers have a gain of 1.5 percent or less (maybe negative) in the current year.

[27] Obviously this is a crude way of estimating expected gains as there are very few controls (sales). Table 1.7 below includes covariates but uses a gainer dummy as the dependent variable rather than a measure of size of tax gains.

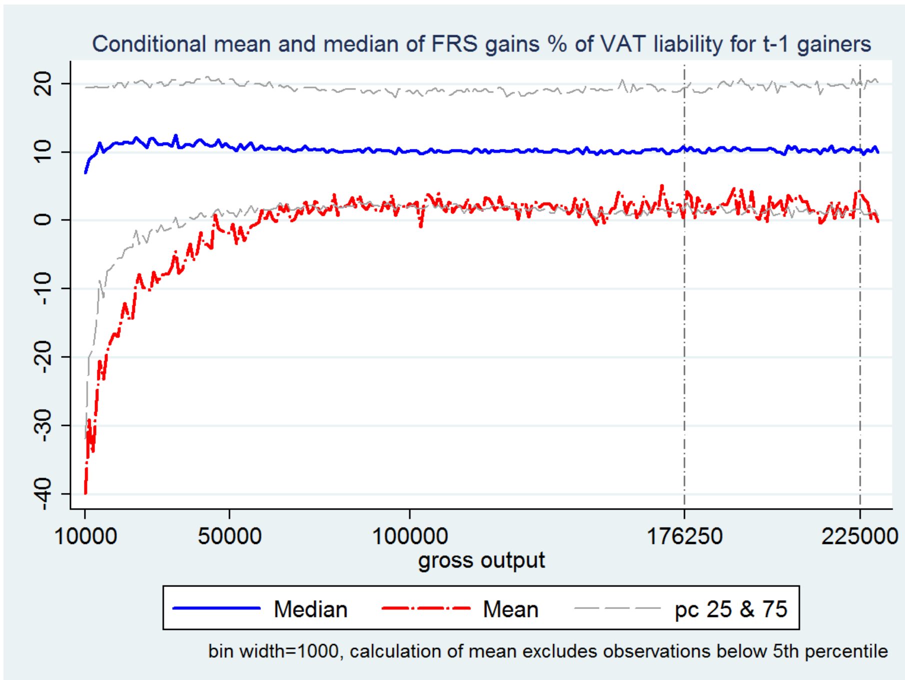

Figure 1.13: Percentiles of FRS gains as a percentage of VAT liability in t for traders observed to gain in $t-1$

Notes: Figure shows twenty fifth, fiftieth (median), seventy fifth percentiles and mean of FRS tax gain as a percentage of VAT liability distribution in year t for VAT traders who are observed to gain from FRS in year $t-1$. Traders are grouped in to bins of gross output in year t and the statistics of the gains distribution are calculated separately for each bin. The gray dashed lines show 25th and 75th percentiles, while the solid black line is the median. The mean is coinciding with the 25th percentile for most of sales level and is indicated by dashed blue line. The sample here is bigger than the one reported in the tables because it includes traders above the FRS eligibility thresholds depicted by the vertical lines. I, however, exclude traders who are ineligible based on observable non-turnover criteria. The first vertical line shows FRS joining eligibility threshold while the second vertical line shows FRS continuation eligibility threshold.

every 10 years. These investments will imply large losses if traders could not recover input VAT. I do not observe these investments separately in the data and therefore assume traders cannot recoup any input VAT when I calculate FRS gains. But the rules of the scheme *allow* reclaiming of input VAT on capital expenditures exceeding £2000. Incorporating this possibility might remove the outliers in figure 1.13 and move the mean closer to the median.

The other feature of the scheme is its easy and quick leaving procedure. Traders can leave the scheme at the end of VAT periods (a quarter for most). Therefore, if traders could predict large upcoming purchases that do not qualify for FRS input recovery, they can simply leave the scheme. Inaction of gainers is justified only when traders face large urgent (unpredictable) purchases that happen with small probability and do not qualify for FRS input recovery. For example, traders might need to purchase large stocks of consumable inputs that could not be postponed until they leave FRS[28].

1.6 Evidence on type of frictions

If uncertainty cannot fully explain inaction of FRS gainers what are the potential hurdles that prevent these traders from joining the scheme? In this section, I use temporal and spatial correlations in the data to justify a combination of information frictions and learning as the most prominent reasons for inaction. I define *information frictions* to include both lack of knowledge about rules required to calculate FRS liability and unawareness of the existence of the scheme[29]. *Learning* suggests traders know about the scheme but are not certain about its benefits. Therefore, they might wait for some time to learn about the optimality of the scheme[30]. I will argue that *inertia*, i.e. sluggish responsiveness to known expected gains, could not fully explain the observed patterns[31].

[28] It is worth noting that demand fluctuations would not necessarily generate higher volatility of FRS liability. If traders use a fixed proportion of inputs to deliver their supplies, an increase in demand increases input use but does not change the share of value added and therefore does not change the relative merit of FRS and VAT.

[29] The evidence is silent on deeper reasons responsible for lack of knowledge: e.g. high cost of acquiring information, biased beliefs about suitability of VAT, tendency to ignore non-default options, and lack of salience of VAT.

[30] Notice, learning could still be important even when there is no objective uncertainty. For example, traders might not know the objective distribution of FRS gains. They would update their prior beliefs after a few observations and learn that FRS is optimal for them.

[31] Evidence from Cash Accounting Scheme (CAS), another small business VAT scheme introduced in 1987, suggests lack of awareness might be key. Traders on CAS pay VAT when they receive money

Before considering the evidence on type of frictions, it is useful to outline potential ways traders could learn about FRS. VAT traders could learn about FRS through a) HMRC, b) tax agents and consultants, and c) business partners and peers. Traders are engaged with HMRC during VAT registration, submission of returns, and audit visits. But chances of learning about FRS is highest at the time of VAT registration because other occasions focus on existing circumstances rather than pointing to new possibilities[32]. Registration is a time of intense learning about VAT which could raise chances of knowing about FRS[33]. Channels (b) and (c) could be operative at all times but they could be stronger during registration. Channel (b) might be less important because less than half of VAT traders use agents[34] and tax and accountancy associations did not support the scheme initially[35].

The importance of registration period for acquiring VAT knowledge, suggests traders registering after FRS was introduced, are more likely to know about the scheme (*hypothesis 1*). Additionally, peer effects suggest traders with FRS-aware partners are more likely to know about the scheme (*hypothesis 2*). In the absence of awareness measures, I rely on estimates of probability of joining FRS for various groups to assess the validity of hypotheses 1 and 2.

I split eligible VAT traders into three groups based on date of VAT registration: a) *Pre-FRS* traders who registered before April 2002 (date of FRS introduction), b) *Early-FRS* traders who registered on or after April 2002 but before January 2004, and c) *Late-FRS* traders who registered on or after January 2004, when flat rates and eligibility thresholds were revised favorably. Hypothesis 1 suggests pre-FRS traders should have the lowest chance of joining because during their registration FRS was absent. In contrast, late-FRS traders might learn about the favorably revised FRS

from customers and reclaim input VAT when they fully pay for the purchase. Based on a telephone survey of around 1500 traders in 2006, HMRC reports 28 percent of eligible traders *have not heard* of CAS (HMRC (2006)).

[32] Among the numerous VAT guides, HMRC publishes one to help traders filling their returns (Notice 700/12 Filling in your VAT return). Interestingly, there is no mention of FRS here until October 2011 revision.

[33] HMRC's website contains a section on special VAT accounting schemes, where FRS is described. VAT experts indicated from October 2012, traders registering online would face the FRS option on the entry form.

[34] Returns data does not show use of agents but National Audit Office (2010) reports around 43 percent of VAT returns were submitted by agents during 2009-10. Furthermore, GfK Business (2008) reports 48 percent of businesses use tax agents for any VAT related issues, while 83 percent of incorporated businesses use agents for corporation tax affairs.

[35] In response to FRS consultation in 2001, many tax and accountancy associations argued FRS diminishes the accounting discipline VAT imposes on traders. 54 responses were received from a total of 225 copies sent out to trade associations, professional bodies, and individual businesses (HM Customs and Excise (2002)).

during registration, and hence should have highest joining probability. Hypothesis 2 implies traders registering later (e.g. late-FRS) are more likely to have FRS-aware partners as the take-up of the scheme was increasing.

To further support hypothesis 2, I use the registered outcodes of traders and define *FRS density* to be the ratio of FRS traders to all eligible ones in each outcode during 2004-05 financial year[36]. Registering in high FRS density outcodes implies greater chance of having an FRS-aware partner and hence higher joining likelihood if information frictions matter. I use the deciles of FRS density distribution and compare joining probabilities for traders registering in different deciles. I restrict attention to traders joining FRS after 2004-05 financial year because this is the year I use for defining FRS densities.

I take a survival time approach, and look at the probability of joining FRS over time. Consider a random variable $T \in [0, \infty)$ representing the time traders join FRS and t as a realization of this random variable. I use Kaplan-Meier (KM) non-parametric method to estimate the conditional CDF of T, $F(t \mid X)$, where X is a vector of categorical covariates. In survival analysis terminology this is known as the failure function. The KM estimation method relies on the fraction of traders experiencing the event of interest. Starting from a total number of traders, n_1, who have the option of joining FRS at time zero, the probability of joining on or before first period is estimated by $\frac{d_1}{n_1}$, where d_1 is the number of traders joining FRS in the first period. For the second period onward it is easier to consider the probability of not joining FRS which is simply the multiplication of probability of not joining in the first period by the probability of not joining in the second period, $\frac{n_1-d_1}{n_1} \times \frac{n_2-d_2}{n_2}$, where $n_2 = n_1 - d_1 - c_1$ is the total number of traders who are still on VAT in the second period and c_1 is the number of traders exiting (censored) the data during the first period. The cumulative probability of being on the FRS by the end of the second period is $1 - \frac{n_1-d_1}{n_1} \times \frac{n_2-d_2}{n_2}$. In general, the probability of joining FRS on or before jth period is estimated by $1 - \prod_{i=1}^{j} \frac{n_i-d_i}{n_i}$, where $n_i = n_{i-1} - d_{i-1} - c_{i-1}$ for $i > 1$.

To complement the non-parametric evidence, I estimate semi-parametric Cox proportional hazard models (CPH) and verify the non-parametric estimates hold after controlling for observables. The hazard rate is defined as the probability of joining FRS in an infinitesimal interval around t conditional on not having joined before t,

[36] Postcodes in the UK consist of two alphanumeric parts. Outcode (postcode district) refers to the first part. For example, WC2A is the outcode associated with WC2A 2AE. The geographical area covered by outcodes varies substantially. I use FRS density to make outcodes comparable.

divided by the length of the interval as it approaches zero. Equation (1.3) shows the definition of hazard rate and its relationship to CDF and PDF of T.

$$h(t \mid X) \equiv \lim_{\epsilon \to 0} \frac{\Pr\left[T \in [t, t+\epsilon) \mid T \geq t,\ X\right]}{\epsilon} = \frac{f(t \mid X)}{1 - F(t \mid X)} \tag{1.3}$$

CPH postulates that the effect of covariates enter as a time separable exponential term as follows

$$h(t \mid X) = h_0(t) \exp\left(\beta' X\right) \tag{1.4}$$

where $h_0(t)$ is the baseline hazard function and determines the evolution of hazard rate over analysis time when $X = 0$. The model is semi-parametric because the partial likelihood estimation leaves the baseline hazard unrestricted. In the next subsection, I provide KM non-parametric estimates of joining probability and in the second subsection, I show results of CPH estimation.

1.6.1 Non-parametric estimation

Figure 1.14 plots non-parametric estimates of the cumulative joining probability for pre, early, and late-FRS traders with the shadings around the lines showing 95 percent confidence intervals[37]. Figure 1.14a estimates joining probability for all eligible traders within the three registration groups. The horizontal axis captures the number of months since the FRS option was available to the traders. For early and late-FRS traders date of VAT registration is set as the zero while for pre-FRS traders the zero is the date of FRS introduction. Consistent with hypothesis 1, the figure shows late-FRS traders have higher probability of joining FRS with early and pre-FRS traders lagging behind. The difference between all groups is statistically significant at 5 percent level. For late-FRS traders the probability of joining FRS jumps to 17 percent right at the time of registration while the same measure remains close to zero for pre-FRS traders. For all groups, the subsequent increase in the joining probabilities is small relative to the initial jump.

The caveat here is that late-FRS traders face a more attractive FRS during the first year of VAT registration (due to the 1 percentage point discount on flat rates introduced in January 2004). Therefore, the higher joining probabilities for this group could be a result of greater benefits from FRS. To alleviate concerns I notice that the three groups face identical FRS incentives after t =24 months, yet the

[37] This is $F(t \mid X)$ where X contains only one categorical variable indicating the three registration periods (pre, early, and late-FRS).

probabilities do not converge. Furthermore, early-FRS traders face similar incentives as pre-FRS group from the outset, but the former shows a 2 percentage points increase in the joining probability at $t = 1$ while the latter does not.

Figure 1.14b focuses on FRS gainers which is perhaps a more relevant population for the analysis of joining probabilities[38]. The x-axis here shows months since the end of the first financial year traders are observed to gain. It is worth noting that all FRS gainers face a similar FRS structure because in order to be observed in this sample they have to be on VAT at least for one year and hence the temporary FRS discount has expired. Here again late-FRS gainers have a significantly higher joining probability with early and pre-FRS gainers lagging behind.

Figure 1.15 looks at the sub-samples within outcodes falling in first, fifth, and tenth deciles of initial FRS density. In figure 1.15a, analysis time is months since traders had the option of joining FRS. Consistent with hypothesis 2, traders registering in tenth decile have the highest probability with fifth and first deciles lagging behind. Joining probability jumps to 20 percent as soon as traders in tenth decile get the chance of joining but the jump is smaller for fifth and first decile traders. The joining probabilities increase almost in parallel for tenth and fifth deciles over analysis time but remain fairly stagnant for the first decile. The probability of joining FRS on or before the end of analysis time is 8, 14, and 30 percent for first, fifth, and tenth deciles. Similarly, for FRS gainers, figure 1.15b shows gainers registered in higher deciles of FRS density are significantly more likely to join FRS in all times after they gain.

Patterns in figure 1.14a could be consistent with inertia. New traders have invested fewer resources in accounting procedures and VAT familiarization, therefore they can invest in FRS accounting procedures. Existing traders are more reluctant to undertake new investments and hence have higher inertia. However, for inertia to justify observed patterns in figure 1.14b, one would need to assume traders with longer experience of VAT have higher inertia. This is a stronger assumption as all FRS gainers have set up normal VAT accounting procedures. But for inertia to justify spatial patterns in figure 1.15, one would need the more demanding assumption of less inertia for high FRS density areas. This assumption seems unreasonable unless high FRS density outcodes turn out to have a higher proportion of new traders. In the semi-parametric estimation I control for this possibility and show spatial patterns

[38]The sample of FRS gainers excludes traders joining right at the time of VAT registration and those with missing gains. Since the former group constitutes a big share of FRS traders I started the analysis by estimating joining probabilities for all eligible traders.

remain robust.

Both temporal and spatial correlations show a secular increase in joining probabilities over time which is more visible for FRS gainers. In other words, it seems that some FRS gainers realize that they could gain from FRS and join the scheme later on. This pattern could be consistent with inertia, learning, or gradual spread of information about the scheme. The above arguments suggest inertia may not be important but to show that learning is probably important I rely on semi-parametric estimates in the next section.

1.6.2 Semi-parametric estimation

So far I have looked at joining probabilities for various groups without controlling for potential confounding factors. For example, traders registering later might be registering in high FRS density outcodes. Therefore, patterns in figures 1.14 and 1.15 might not necessary reflect two distinct correlations. To rule out this possibility and other observable confounders, I estimate CPH models (equation (1.4)). Estimation results are reported as hazard ratios for ease of interpretation. For dichotomous variables hazard ratios are defined as the ratio of the hazard rate when the variable is equal to 1 to when it is 0, fixing other variables:

$$
\begin{aligned}
HR_i = \frac{h(t \mid x_i = 1,\, X_{-i})}{h(t \mid x_i = 0,\, X_{-i})} &= \frac{h_0(t)\exp\left(\beta_i \times 1 + \beta'_{-i} X_{-i}\right)}{h_0(t)\exp\left(\beta_i \times 0 + \beta'_{-i} X_{-i}\right)} \\
&= \exp(\beta_i)
\end{aligned}
$$

This suggests the rate of joining FRS is $HR_i = \exp(\beta_i)$ times higher for $x_i = 1$ traders relative to $x_i = 0$ ones. Alternatively the likelihood of joining FRS is on average HR_i times higher for $x_i = 1$ traders relative to $x_i = 0$ traders during the analysis period.

Table 1.8 reports estimation results when the start of analysis time is from the time traders have the option of joining FRS. The variables of interest are “gainer”, a dummy variable that is equal to 1 if trader is an FRS gainer, two dummies capturing early and late-FRS traders, and “initial FRS density”. In all specifications, I control for average and standard deviation of FRS gains over VAT liability for each trader, average logarithm of gross output, the ratio of number of years trader was eligible for FRS, dummies for sole proprietors and partnerships, and dummies for frequency of submitting returns. Standard errors are adjusted for clustering at 5-digit SIC2007

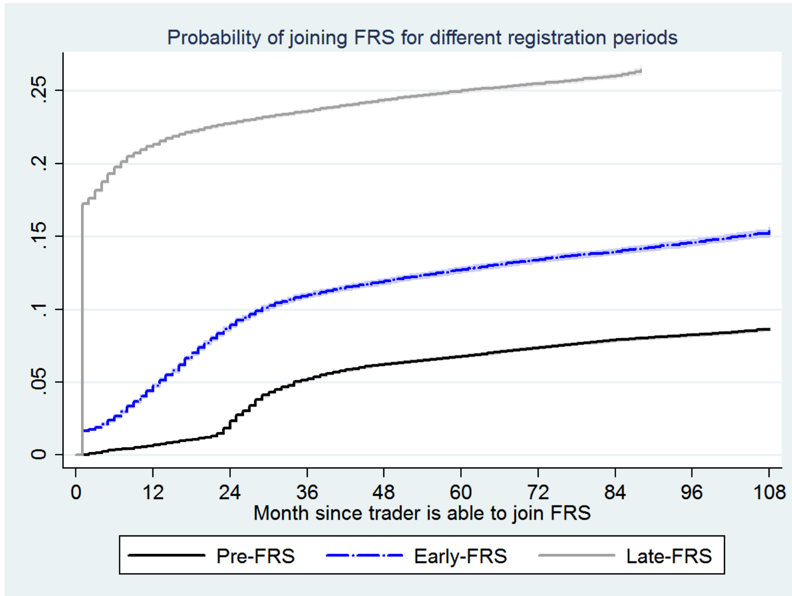

(a) Probability of joining since traders were able to join

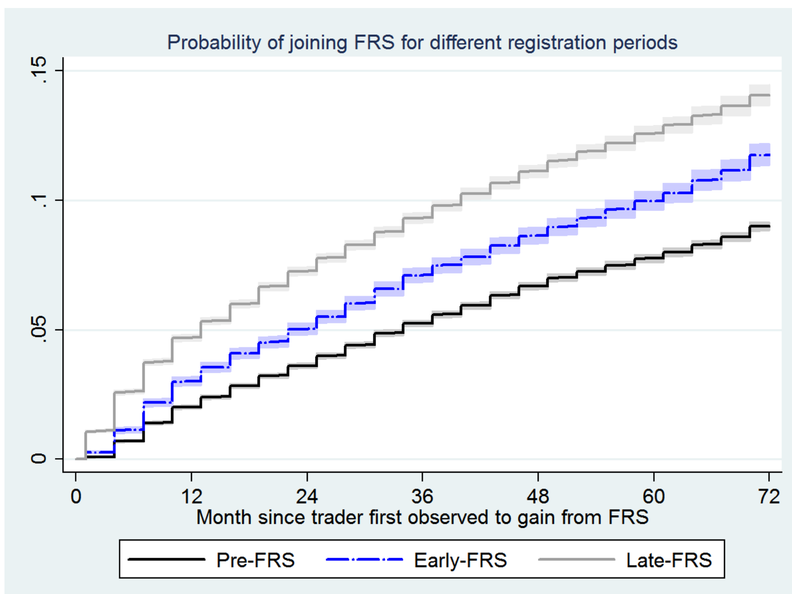

(b) Probability of joining since traders first gained

Figure 1.14: Probability of joining FRS for different VAT registration periods

Notes: Figures show Kaplan-Meier non-parametric estimates of probability of joining FRS on or before the analysis time for traders registering during different periods. Pre-FRS traders are those registering for VAT before April 2002. Early-FRS are traders registering between April 2002 and before January 2004. Late-FRS are traders registering on or after January 2004. 95 percent confidence intervals are shaded around the lines. Panel (a) shows joining probability since the time traders had the option of joining FRS. The zero of analysis time shows time of VAT registration for early and late-FRS groups but is fixed at April 2002 for pre-FRS traders. The initial ($t = 0$) number of traders that could potentially join FRS is 679,510 Pre-FRS, 180,416 early-FRS, and 943-241 late-FRS. Panel (b) shows joining probability as a function of months since traders first gained. This is the end of financial year where traders are observed to gain for the first time. The initial ($t = 0$) number of gainers that could potentially join FRS is 213,037 Pre-FRS, 52,145 Early-FRS, and 182,310 late-FRS traders.

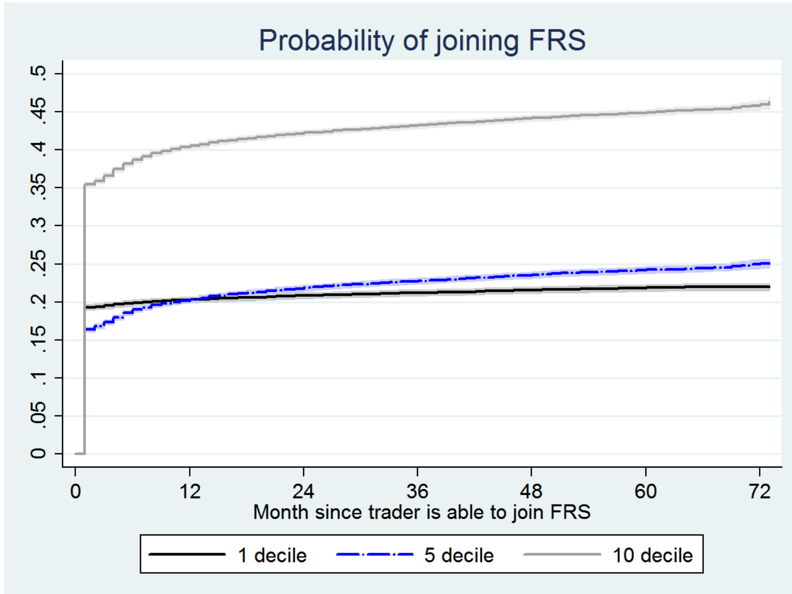

(a) Joining probability since traders are able to join

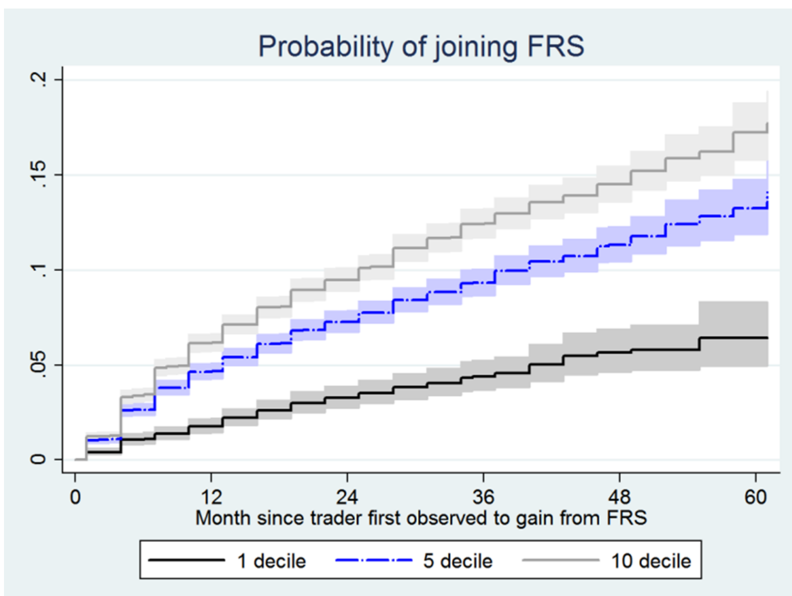

(b) Joining probability since traders gained

Figure 1.15: Probability of joining FRS for deciles of initial FRS density

Notes: Figures show Kaplan-Meier non-parametric estimates of probability of joining FRS on or before the analysis time for traders registering in outcodes featuring 1, 5, and 10 deciles of FRS density distribution in 2004-5 financial year. 95 percent confidence intervals are shaded around the lines. Panel (a) shows joining probability since the time traders had the option of joining FRS. The zero of analysis time shows either time of VAT registration or time of FRS introduction, April 2002, whichever is later. The initial ($t = 0$) number of traders that could join FRS are 59,094 in first, 76,803 in fifth, and 91,146 in tenth decile. Panel (b) shows joining probability as a function of months since traders first gained. This is the end of financial year where traders are observed to gain for the first time. The initial ($t = 0$) number of gainers that could join FRS are 6,484 first, 15,056 Fifth, and 15,856 tenth decile.

codes.

In column (1) a simple CPH model is estimated. The likelihood of joining FRS is 3.862 times (286 percent) higher for gainers relative to those never gaining. In line with figure 1.14a, early and late-FRS traders are respectively 55 and 178 percent more likely to join FRS compared to pre-FRS traders. Columns (2) to (5) estimate stratified CPH models with 5-digit SIC2007 codes and deciles of initial FRS density as grouping variables. Stratification allows baseline hazards to vary flexibly across SIC2007 by FRS density groups but restricts to identical covariate effects across strata[39]. Coefficient estimates are slightly reduced when I allow for stratification in column (2) but the main results remain robust. Gainers are still 202 percent more likely to join FRS. Early and late-FRS traders are 51 and 171 percent more likely to join relative to pre-FRS traders. In column (3) I remove traders with less than three years of returns data and the results are still robust.

Column (4) includes interactions of registration period dummies with gainer indicator. The interaction terms capture the change in the hazard rate for gainers registering in different periods. Early-FRS gainers are 45 percent more likely (significant) to join FRS relative to pre-FRS gainers ($1.45 = 1.851 \times 0.782$). Late-FRS gainers are 136 percent more likely to join FRS relative to pre-FRS gainers ($2.36 = 4.313 \times 0.546$).

Column (5) includes the ratio of the number of gain years to total observation years for each trader. The estimates here support coexistence of learning and information frictions. Estimates of hazard ratios for late and early-FRS traders remain by and large similar to previous specifications (information friction). Traders with one more year of gaining are on average 30 percent more likely to join the scheme (assuming 7 years of returns)[40]. This result suggests that learning plays a role in the joining decision, albeit somewhat smaller than the impact of registration periods (early and late-FRS traders have 52 and 182 percent higher likelihood of joining).

Column (6) only stratifies on 5-digit SIC2007 codes and instead includes a continuous variable for FRS density of the registration outcode of the trader. Here I restrict to traders registering from 2005-06 financial year onwards (hence remove early and late-FRS dummies)[41]. Increasing initial FRS density of the registration outcode of

[39] Stratification is similar to fixed effects in a linear regression. However, in CPH models stratification allows for more flexibility than inclusion of dummies. Dummies shift the hazard rate proportionately across categories but stratification allows independent time paths for each strata.

[40] To calculate this, I used the original coefficient estimate from column (4). Specifically, $\exp(\frac{1}{7} \times \ln(6.245)) = 1.30$.

[41] This restriction is put in place because by definition outcodes with higher FRS density in 2004 would show higher joining probabilities in that date.

Table 1.8: Estimates of hazard ratios (Cox proportional hazards model)

		Stratified: SIC2007 and deciles of FRS density				Stratified: SIC2007
	Baseline	Baseline	Traders with at least 3 returns	Interactions	% years gained	Include FRS density
	(1)	(2)	(3)	(4)	(5)	(6)
Gainer	3.862 (.405)*	3.025 (.227)*	2.924 (.269)*	4.319 (.385)*	0.865 (0.045)*	2.488 (0.276)*
Fraction of years gained					6.245 (.685)*	
Early-FRS	1.555 (.126)*	1.510 (.06)*	1.515 (.061)*	1.851 (.157)*	1.566 (.056)*	
Late-FRS	2.78 (.189)*	2.716 (.074)*	2.818 (.083)*	4.313 (.224)*	2.873 (.093)*	
Early-FRS× Gainer				.782 (.065)*		
Late-FRS × Gainer				.546 (.03)*		
FRS density						16.393 (0.132)*
FRS density × Gainer						0.407 (0.132)*
Observations	918,396	915,887	731,508	915,887	915,887	276,287
Number joining FRS	28,206	28,206	28,206	28,206	28,206	7,428

Notes: Table reports hazard ratios from estimation of Cox proportional hazard models. Controls included are average and standard deviation of FRS gains as a fraction of VAT liability for each trader, average of logarithm of gross sales, fraction of years trader was eligible, dummies for sole proprietors and partnerships, and dummies for frequency of submitting returns. Standard errors are adjusted for clusters in SIC2007 and reported in parenthesis. Stars show hazard ratio is significantly different from one at 1 percent level. Reported standard errors are calculated from original standard errors on coefficient estimates using delta method. This amounts to multiplying the original standard errors by $\exp(\beta_i)$. Test of significance, however, relies on the original z-score derived from the ratio of coefficients to the standard errors. Column (2) to (5) estimate stratified Cox models using SIC2007 and deciles of 2004 FRS density as strata. Column (6) only uses SIC2007 as a stratum and restricts the sample to traders registering from 2005-06 onwards.

traders by 0.05 increases the likelihood of joining by 15 percent for traders never gaining and 174 percent for FRS gainers. Overall, CPH estimations support the patterns presented in figures 1.14 and 1.15.

1.7 Conclusions

Results presented in this chapter show a significant number of small businesses with non-negligible tax savings fail to join FRS. I have provided evidence that observed FRS gains are a strong signal of expected gains. Therefore, it seems uncertainty cannot fully explain inaction. This, however, does not imply that traders' prior beliefs about suitability of the scheme are correct. Some traders might need to spend a few years observing gains before updating their beliefs about suitability of the scheme. Others might not even know about the scheme or have difficulty calculating potential gains. Temporal and spatial correlations in the joining probabilities indicate that both of these stories have some merit. Traders registering when FRS was in place are more likely to join the scheme. Furthermore, traders registering in high FRS density outcodes are more likely to utilize the scheme.

The conclusion that small traders are susceptible to errors in their business decision making resonates with the results of Devereux et al. (2014) who find small incorporated business are not completely shifting their incomes to the corporate base while in a frictionless world it is optimal to do so. Their preferred explanation for sub-optimal behavior is illiquidity of corporate profits and the need for having a stable flow of income (e.g. in the form of personal income). In this paper, however, I argued for presence of information frictions and learning. This means in the case of FRS, gainers would join the scheme if they get the right information or can resolve their doubts about optimality of the scheme more quickly. Accepting this view calls for a more effective role for the government to publicize business support schemes.

Chapter 2

Stimulus effect of the UK 2008 VAT rate cut

2.1 Introduction

The great recession 2008-2009 has lead to substantial fiscal stimuli and unprecedented expansionary monetary policy. Tax rebates, incentives for investment and consumption, and investment in infrastructure are common elements of recent fiscal stimulus packages around the world. Knowing whether fiscal policy could stimulate the economy during recessions and which elements are more successful are key issues in policy design. While there exist a large body of literature that studies the impact of fiscal policy, the debate about its effectiveness during recessions is far from settled. In fact as Auerbach et al. (2010) conclude "much of what has been learned recently concerns how such [fiscal] multipliers might vary with respect to economic conditions ...".

In the UK a temporary reduction of VAT rate was the main element of the fiscal stimulus package. On 24 November 2008 the Chancellor of the Exchequer announced that the VAT standard rate will be reduced from 17.5 to 15 percent from 1 December 2008 to 31 December 2009. The rate cut was heralded as timely, targeted, and reversible. The cut was estimated to cost £12.5 billion during the 13 months of its operation which amounts to approximately 15.5 percent reduction in VAT receipts or 2.2 percent fall in total tax revenue. VAT receipts data confirms the cut shaved off around £12 billion during 2009 calendar year. Figure 2.1 shows cumulative VAT receipts over each calendar month. VAT receipts in 2007, 2008, and 2010 all stand

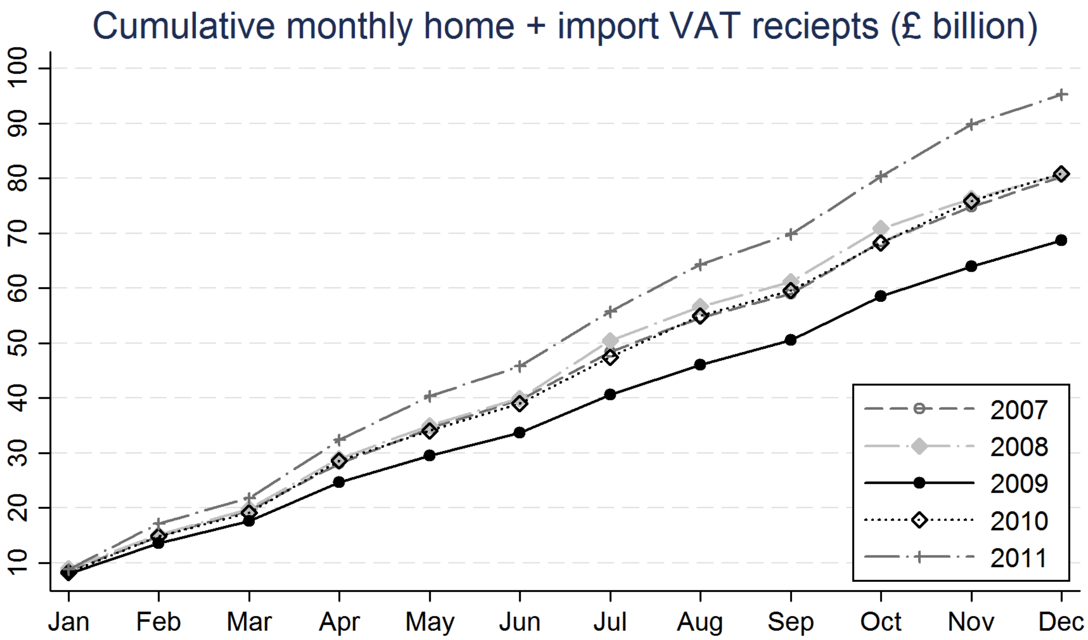

Figure 2.1: Total VAT receipts

Notes: This figure plots cumulative monthly VAT receipts from HMRC's VAT bulletin published in February 2014. The results from receipts data is not the same as those from returns. Furthermore, this figure includes import VAT whereas returns data used later on is for home VAT only. VAT rate was 17.5 percent prior to December 2008. It was reduced to 15 percent between 1 December 2008 and 31 December 2009. Reverted to 17.5 percent from 1 January 2010 and then increased to 20 percent from 4 January 2011.

around £80 billion. During 2009, i.e. 12 out of 13 months of the rate cut, VAT receipts are around £68 billion which shows an approximate fall of £12 billion (14.9 percent of VAT revenue in 2008).

The theoretical impact of the cut depends on whether traders pass-on the cut to customers or take home the reduction in tax liability. In the former case, the cut would result in income and substitution effects while in the latter the substitution effect would be absent. The income effect could result in higher consumption or extra savings (e.g. paying debt). But the consumption increase is not expected to be substantial due to the temporary nature of the cut unless individuals are credit constrained or myopic[1]. Two types of substitution effects could be present in the

[1]Under permanent income hypothesis, unanticipated temporary increases in income would be spread over the life cycle and therefore should have little impact on current period consumption. Credit constrained consumers would, however, consume any marginal income. Interestingly, many recent studies of US tax rebates find substantial consumption responses right after the receipt of rebates (e.g. Johnson et al. (2006), Parker et al. (2013), and Agarwal et al. (2007)). Most of the rebate money is consumed within a few quarters after receipt. For a review of the empirical literature on marginal propensity to consume out of income shocks refer to Jappelli and Pistaferri (2010).

case of price reductions. First, demand for standard-rated items would increase as their price relative to zero-rated items is lower (intra-temporal). Second, price of consumption is lower during the cut and consumers would shift purchases to benefit from lower prices (inter-temporal). Given about one third of standard-rated items are durable goods in the UK, the inter-temporal substitution effect could be strong because consumers can stock up and consume these items later[2].

The key difference between the UK VAT cut and the US tax rebates is that the former encourages consumption through price incentives (assuming some degree of pass through) while the latter works purely through an increase in after-tax income[3]. Absence of inter-temporal substitution effects for tax rebates could reduce their effectiveness as a stimulus policy. The same effect could also jeopardize nascent recovery if the economy has not returned to normal conditions after the expiry of the VAT cut[4].

A common issue to VAT cut and tax rebates is salience[5]. The VAT in the UK is quite complex, and it is not obvious that consumers know which products are subject to the standard-rate. Targeted advertisement by retailers at the time of the cut, however, increased the salience of the cut[6]. A related issue is the size of incentives. Small incentives might not be as effective in encouraging extra consumption. Under full pass through the standard-rate cut would reduce prices by 2.1 percent which might be insignificant in the face of large income drops during the recession (a £117.5 item would see a £2.5 price reduction)[7].

In this paper, I use administrative VAT returns data between 2002q1 and 2010q4

[2]The recession might moderate the inter-temporal substitution effect by tightening credit constraints and increasing income uncertainty. For a detailed discussion of the potential impact of the cut refer to Blundell (2009), Crossley et al. (2009), and Barrell and Weale (2009).

[3]While price incentives are widely used to promote business investment (e.g. R&D and investment tax credits), use of price incentives was more limited in the US (except for Cash for Clunkers program of 2009 and First-time Home buyer Credit).

[4]For the cash for clunker program in the US, Mian and Sufi (2012) find substantial demand shifting. They estimate that the two months program has led to 370,000 more car purchases but car purchases were lower for a period of 10 months after the program expiry. In other words, the cash for clunkers was ineffective in boosting medium run demand.

[5]For example if posted prices are tax exclusive (as in Chetty et al. (2009)) or the tax cut is applied at the till, consumers might fail to notice the price reduction. For tax rebates Sahm et al. (2012) find that a tax cut delivered through reduced withholding has half of the effect of a similar one-off tax rebate.

[6]Big retailers like Tesco heavily advertised the VAT cut and showed calculations of extra savings on their websites.

[7]In the context of US tax rebates Parker et al. (2013) find significant impacts on durable consumption for the larger 2008 rebates while Johnson et al. (2006) do not find a significant impact on durables from 2002 rebates. Both studies, however, find significant impacts on non-durable consumption.

from HM Revenue and Customs (HMRC) to estimate the stimulus impact of the standard-rate cut. Administrative VAT returns data is well suited for studying the impact of the cut for several reasons. First, I observe effective tax rates on sales and purchases and therefore, could identify standard-rated traders (i.e. treated). Second, amount of measurement error should be minimal relative to survey data because of potential penalties for mis-reporting. Third, I observe a large number of traders over the course of 32 quarters and therefore could control for a rich set of fixed effects (e.g. trader and date fixed effects plus sector by recession interactions) to alleviate concerns regarding confounding factors. The caveat of this data is that I do not observe quantities and prices separately. So I will not be able to separately identify price and quantity responses to the VAT cut.

The key challenge for estimation of the stimulus effect in presence of the great recession is to construct a valid counterfactual: how would sales, purchases, and value added have evolved for the group of treated traders had the standard-rate not been cut? The fact that a large part of consumer spending is zero-rated in the UK provides a natural solution to this challenge. I categorize traders into treatment and control based on pre-cut average effective output tax rates, i.e. sales VAT divided by sales. Traders with tax rates close to the standard-rate would potentially receive a benefit from the cut while traders involved in zero-rated activities would not. I use a difference-in-differences (DD) identification strategy and compare average growth rate of sales, purchases, and value added across standard and zero-rated traders before and during the VAT cut. The identification assumption here is one of parallel trends for the growth rates: had the standard-rate not been reduced the change in average growth rates during the cut period relative to the pre-cut period would have been identical for standard and zero-rated traders.

In returns data, growth rate of value added becomes negative from around 2008q1 and remains negative until 2010q2 (figure 2.3 panel a). The cut period covers 13 months affecting returns submitted in 6 quarters from 2008q4 to 2010q1. As far as the average impact of the recession is similar across the group of standard and zero-rated traders, the DD estimate would partial out any recessionary effects and delivers an unbiased estimate of the stimulus impact. This assumption is, however, unlikely to hold. In fact as figure 2.3 panel b) shows right around the recession time a clear gap emerges between average growth rates of standard and zero-rated traders with the standard-rated traders showing greater declines[8].

[8] The larger recession impact on standard-rated traders could be due to the fact that most durable goods fall in this category.

I use two strategies to isolate confounding recession effects. First, I allow average growth rates to differ for standard and zero-rated traders during the recession period (2008q1-2010q4) by including the interaction of a recession dummy with the treatment dummy. Estimated stimulus impacts under this strategy are effectively equivalent to dropping pre-recession data points. The estimated magnitudes reflect the differential change in average growth rate of standard-rated traders right at the cut period relative to to the six recessionary quarters before and after the cut. The identification assumption is now refined and this strategy would deliver unbiased estimates of the stimulus effect when the recession has a heterogeneous but time invariant impact on standard-rated traders[9]. Figure 2.3, however, suggests the deepest part of the recession coincides exactly with the cut period. Therefore, the differential impact of the recession could be changing over time with the greatest effect showing up right in the middle of the cut.

The second strategy I adopt for dealing with the recession is to allow heterogeneous recession effects for two-digit sectors by including sector by recession interactions in the regression. The identification of stimulus effect here relies on differential change of growth rates for standard-rated traders within the same two-digit sector right at the time of the cut. To the extent that recessionary effects are on average the same for standard and zero-rated traders within the same two-digit sector, the DD estimates from this specification would deliver unbiased estimates of the stimulus effect. This method allows for sector specific recession responses but assumes standard and zero-rated traders within the same broad sectors receive a similar recessionary effect[10].

DD estimates of the stimulus effect from the basic specification show implausibly large and negative numbers but once I employ either of the above strategies to deal with heterogeneous recession effects, the magnitudes become much smaller (sometimes positive) and insignificant. Impact of the cut on sales growth is estimated to be between -0.1 and 0.2 percentage points depending on the specification (all insignificant). Similarly, the cut has led to a reduction of purchases growth by 0.7 to 1 percentage point (insignificant).

A zero effect on gross sales and purchases is suggestive of a proportionate increase

[9]The common recession effect could vary over time. This assumption only requires the differential impact of the recession to remain constant.

[10]The first and second methods deal with different concerns and therefore it might be difficult to select one as the preferred method. The first method allows for a differential recession response for standard and zero-rated categories while the second method controls for two-digit sector specific recession effects. Since there is not a perfect correlation between two-digit sectors and standard-rated categories, the two measures control for different potential confounders.

in quantity demanded in response to tax induced price reductions. Under full pass through, the 2.5 percentage points reduction in the standard-rate would translate to a 2.1 percent price reduction. This price reduction would unambiguously lead to an increase in quantity demanded. But unless the price elasticity of demand is greater than 1, the resulting change in gross sales would be negative. Therefore, under full pass through a zero effect on gross sales suggests a proportional change in quantity demanded.

I check the robustness of my results for various sub-samples. First, VAT is eventually a consumption tax and one might expect that in a perfect VAT system intermediate production stages do not directly respond to rate changes[11]. Therefore, the retailers are a more relevant group to study the direct impact of the VAT cut. When I restrict attention to the retail sector, I find very similar results. I get small insignificant coefficient estimates for the impact of the cut on sales growth. Second, the change in the standard-rate could in principle affect both input and output effective tax rates. To remove potential confounding effects from the input dimension, I restrict the sample to traders that use solely standard or zero-rated purchases. Results from both of these sub-samples confirm the earlier findings. Finally, large traders might get more benefits from the cut by spending more on advertising the rate cut. When I restrict the sample to large traders, I still get estimates close to zero for sales and purchases.

Several papers study impact of VAT reductions on prices and sales in other countries using a similar DD strategy (I review the existing work on UK VAT cut in the next section). Turkey implemented VAT and special consumption tax cuts on certain durable and luxury goods during the financial crisis. Misch and Seymen (2013) compare changes in sales after the tax cuts for firms selling treated goods to unaffected firms. The Turkish tax cuts were implemented upon short notice and happened between March and September 2009 (a period of less than 7 months). Using three waves of Financial Crisis Surveys[12], they estimate that the group of affected firms had 39 percentage points higher sales growth relative to control firms. They justify the extreme magnitude of this coefficient based on the size of tax cuts. It is however, not entirely clear whether this could explain the result. The tax rate for passenger cars was reduced from 55 to 37 percent, but tax rate for white household goods was reduced from 6.7 to 0 percent. While they control for firm, industry by time and

[11]Traders in the intermediate stages might respond to VAT changes because the VAT system in the UK features extensive exemptions.

[12]This dataset is available at `www.enterprisesurveys.org`. The World Bank and International Finance Corporation commissioned these surveys in several countries.

region by time fixed effects it is not entirely obvious that their sample allows such a demanding specification. Their sample size is between 880 and 717 observations and the average number of observations per firm is 1.6.

Harju and Kosonen (2013) consider the reduction of VAT rate for restaurant meals in Finland from 22 to 13 percent in 2010. They carry out a difference-in-differences estimation using hotels and restaurants in neighboring countries as control groups for restaurants in Finland. The results show a low pass-through of 25 percent and they are unable to find a significant effect on sales.

Carbonnier (2007) uses two French VAT rate reductions for cars (33.3% to 18.6% in 1987) and housing repair services (20.6% to 5.5% in 1999) and estimates a pass through of 77 and 57 percent respectively in car sales and housing repair services. The estimates are in line with the market structure of these sectors as car sales is much closer to oligopoly. He also finds tax shifting is complete four months after the rate cuts with most of the change happening within the first two months.

In the next section, I briefly describe the VAT in the UK and discuss the timeline of the standard rate cut. Here I provide a review of existing studies that try to assess the success of the VAT cut as a stimulus policy. Section 3 discusses VAT returns data and present summary statistics. This section also justifies the definition of standard and zero-rated traders I use. In section 4 I explain the empirical strategy and discuss various specifications I use to isolate recession effects. Section 5 presents the graphical and regression evidence on the impact of the cut. A final section concludes.

2.2 Context

Businesses with annual taxable turnover above a threshold (£67,000 during 2008 financial year) must register for VAT in the UK[13]. Taxable turnover relates to total supplies of commodities and services under three different VAT rates. Table 2.1 shows the list of activities under each rate. Food, books, children clothes, exports, and other activities under the first column are zero-rated. This means the tax rate

[13]Business units of a single corporate body usually have one VAT registration. HMRC, however, allows separate VAT registrations for individual business units or divisions but there are strict conditions for separate registration. It seems most of big chain stores have a single VAT registration and submit one tax return in each accounting period. HMRC also allows group registration for a company with subsidiaries under some conditions. VAT returns data shows a very small fraction of returns relate to companies with group registration.

Table 2.1: Activities under different VAT categories

Zero-rated	Reduced-rated	Standard-rated	Exempt
Exports **Food** Books, newspapers and magazines Passenger transport Supplies to disabled and charities, Domestic water or sewerage services, Construction and sale of new domestic buildings, Children clothing, Cycle helmets, etc.	Mobility aids for elderly Energy saving and new energy items Domestic fuel and power Women's sanitary products Contraceptives Children's car seats Smoking cessation products Residential conversions and renovations	**household goods, and durables** **legal, accounting, consultancy services** catering, taxis, and everything not in other categories	Rent on domestic dwellings Supplies of commercial property Private education and Health services Postal services Burial and cremation Finance and insurance Betting, gaming and lottery Cultural admission charges

Notes: Exempt traders are not observed in my data because they cannot register for VAT.

on these supplies is zero but the businesses can still reclaim VAT on their purchases. Therefore, zero-rated traders would normally receive refunds from HMRC. The second column lists supplies under reduced-rated category. The tax rate on supply of these products is 5 percent during my sample. The observed number of traders under reduced-rated group is very small compared to the two other categories due to the narrow definitions. The last column of the table shows a non-comprehensive list of standard-rated activities. The tax rate on supply of household goods, most business services, and other standard-rated items is 17.5 percent prior to December 2008. Apart from the three tax rates, certain supplies are VAT exempt. Traders involved in exempt activities cannot register for VAT and are absent from VAT returns data[14]. These traders do not pay any VAT but cannot reclaim any input VAT. Last column of the table shows the list of exempt activities.

2.2.1 Standard rate cut

Figure 2.2 shows the evolution of the three statutory VAT rates during my sample. The zero and reduced rates are fixed during the whole sample period. The standard rate was at 17.5 percent since April 1991. In response to disappointing GDP figures in the three first quarters of 2008 and deepening financial crisis, on 24 November 2008 Alistair Darling, the Chancellor of the Exchequer, announced a 2.5 percentage

[14]However, traders involved in sale of taxable and exempt supplies can register for VAT. I drop sectors that qualify for exemptions from my analysis.

points temporary reduction in the standard rate[15]. The standard rate was reduced to 15 percent from 1 December 2008 to 31 December 2009[16]. The cut was heralded as best and fairest approach to boost the economy by "giving back" 12.5 GBP billion of tax to consumers:

> A reduction in the rate of VAT has been chosen as the main lever for the fiscal action as this change can be implemented rapidly (timely), will impact immediately on the purchasing decisions of firms and individuals to boost spending (targeted) and is reversible (temporary). A temporary reduction in the rate of VAT will lower prices for households and should provide help immediately, when they need it most. It will also incentivise them to bring forward the purchase of goods, which will help support firms and the people they employ as the economy slows. (part of Chancellor's announcement reported from Seely (2009))

The fact that there was only one week between announcement and implementation leaves little room for real behavioral responses before the cut[17]. Although the standard-rate was announced to return to pre-cut levels, there was some speculation about subsequent rate increases to compensate for the lost revenue. But the standard-rate was reverted to 17.5 percent as promised. However, on 22 June 2010 the new coalition government announced an austerity budget projecting a 2.5 percentage points permanent increase in the standard rate from 4 January 2011. The standard rate has remained at 20 percent since then. My sample covers returns submitted between 2001q3 and 2011q2 but I focus on the temporary rate reduction and its reversal. I refer to quarters from 2008q4 until 2010q1 as the cut period[18].

[15]The cut was quite unexpected. I could find speculative newspaper reports from 22 November 2008 onward. It seems the first report was by The Telegraph on 22 November 2008 but most of these reports were inaccurate. For example The Independent reports "The cut will be at least 2 percent, possibly to 15 percent, where it will remain for a "holiday" of one-and-a-half to two years."

[16]This is a reduction of 14.3 percent in the standard-rate but would result in a price cut of 2.1 percent under full pass through. Before the cut the tax inclusive price is $(1+0.175)p$, where p is the tax exclusive price. After the cut and under full pass through this will be $(1+0.15)p$. Therefore the change in the tax inclusive price is $^{0.025}/_{1.175} = 0.021$.

[17]There might be pure reporting responses due to the timing of returns submission. Returns are submitted with one month (and 7 days if online) delay. For example, returns relating to transactions between 1 August and 31 October 2008 are submitted on 30 November 2008. Therefore, these returns are submitted with the cut knowledge. I am, however, unable to identify any impact on average growth rates before the cut.

[18]A maximum of one month in 2008q4 returns and a maximum of two months in 2010q1 returns might cover the cut period.

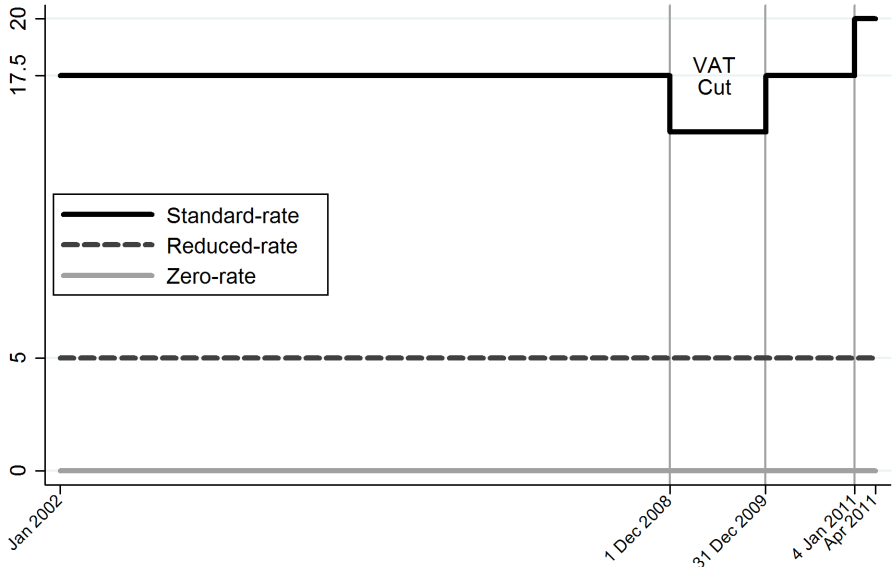

Figure 2.2: VAT rates over time

2.2.2 Assessments of the cut impact

Crossley et al. (2009) and Blundell (2009) try to predict changes in consumption as a result of the VAT cut given the existing evidence on elasticities. Due to the temporary nature of the cut, the authors expect a small income effect but a large inter-temporal substitution effect. Since luxuries fall in the standard-rate category in the UK, the authors suggest the inter-temporal elasticity of substitution will be around 1. Given that the cut lowered prices by 2.1 percent this would suggest an increase of 2.1 percent in demand for standard-rated items. This is consonant with my finding that gross sales did not change in response to the rate cut. Barrell and Weale (2009) use aggregate consumption data for six European countries with VAT rate changes (2 rate reductions and 7 rate increases) to estimate likely impact of a rate reduction in the UK. Their regressions show a 1 percentage point increase in the standard VAT rate has led to 0.3 percent increase in consumption before the rate rise and 0.5 percent reduction after the rate increase. Using a simulation model they conclude that the consumption would be increased by less than 1 percent while GDP will increase by less than half a percent[19].

The cut was speculated to be ineffective by many observers on the grounds that the resulting price change is insignificant in the face of drastic incomes falls[20]. Two

[19] Fernandez-de Cordoba and Torres (2011) get somewhat similar results using a calibrated general equilibrium model.

[20] For a discussion of the cut in the media see Iain Dale's blog, and his round up of other blog

survey based studies find little positive impacts. The Guardian reports the results of a PWC survey that shows 88 percent of consumers "said that the VAT cut had not prompted them to spend more on goods or services". ORC international, on behalf of HMRC, interviewed 2,005 VAT registered businesses during May and June 2009 to assess the compliance and commercial impact of the cut (ORC International (2010)). 78 percent of businesses responded that they passed on the VAT cut (almost all to the full extent) while 15 percent did not change their prices. However, 46 percent disagreed that the cut had a positive effect on their sales while 26 percent agreed there was a positive effect. Interestingly, there is not much of difference between businesses selling to other businesses or final consumers: respectively 17 and 21 percent stated there was a positive impact[21]. It is worth noting that in neither of the surveys people are asked questions about quantities. In the PWC survey the question is about consumer spending and in the HMRC study it is about sales. Therefore, a zero effect on spending and sales resonate with my findings and alludes to a demand elasticity of 1.

Chirakijja et al. (2009) use retail price index (RPI) to estimate the degree of pass through. They compare 28 two-digit categories containing standard-rated items to 36 groups with no tax change. They find the price index for standard-rated categories fell by 1.5 percentage points relative to other items right at the time of the VAT cut. Furthermore, there does not seem to be a lagged price change and all the impact materializes right at the time of the cut. This corresponds to a pass-through of around 71 percent[22]. This is close to the literature-based conclusion of 75 percent in Blundell (2009). The standard errors are quite large here and one can not reject full pass through at 5 percent significance level. Chirakijja et al. (2009) show evidence on the salience of the cut from a survey of consumer confidence. They show a larger fraction of consumers declare "it is a good time to buy large household appliances" right after the VAT cut while they still have a poor evaluation of overall economic situation.

Pike et al. (2009) discuss the difficulty of collecting price data after the VAT cut. Office of National Statistics (ONS) gathers price data partly from shelf labels[23] but

posts on the cut. Also BBC interviews show general skepticism about the cut effectiveness.

[21] The results show traders incurred additional costs to comply with the rate change. The median time spent on operationalizing and complying with the VAT cut was 2.7 hours. The mean was dominated by larger businesses and is much higher at 11.4 hours. I abstract from such costs in my analysis of the cut impact.

[22] The gross tax rate changed from 1.175 to 1.15 and therefore the change in that variable is 2.1 percent. Therefore pass-through is calculated as $\frac{1.5}{2.1} = 0.71$.

[23] ONS also gets some prices centrally for large chain stores or for services like utilities with no

many businesses did not change shelf prices and gave a discount at the till after the cut to save on relabeling costs. In the face of this evidence ONS has applied adjustments to collected price data which might have implications for Chirakijja et al. (2009) study. Pike et al. (2009) show that from local shops visited by ONS data collectors only 14 percent changed shelf prices while 43 percent applied the cut at the till, and 34 percent did not pass on the cut to consumers. This suggests a pass through rate of 66 percent for local shops which is slightly lower than earlier estimates.

Crossley et al. (2013) use UK Economic Accounts and Living Costs and Food Survey (LCFS) to analyze the evolution of different elements of consumption through three recessions of 1980, 1990, and 2008. They document that the fall in real GDP from peak (2008q1) to trough (2009q2) of the current recession was 7.1 percent which is by far greater than the earlier recessions. Interestingly, they confirm the intuition that real consumption of durables fall more than non-durables during recessions. But it seems the current recession is showing a different pattern. While initially the fall in durable purchases mimics that of earlier recessions, from 2009q1 until 2009q4 durable purchases starts to rise while non-durable purchases is flat. This period coincides with the VAT rate cut and also covers the car scrappage scheme. Therefore, they conclude that two schemes seem to be somewhat effective. The increase in durable purchases is reversed after 2009q4 and durable purchases start to decline further which could be consistent with inter-temporal shifting of durable demand (Figure 4 in the paper).

In order to show the difference between the results from returns data and national accounts (as in Crossley et al. (2013)), figure 2.3 shows a comparison of consumption growth from national accounts and average value added growth from returns data. Panel a) shows overall movements of the two series is very similar. For both series, 2008q2 is the first quarter that growth becomes negative. The two series continue to decline with equal rates until the beginning of the VAT cut where the fall in value added accelerates. The growth rates start to rise from 2009q2 until during 2010 when they become slightly positive.

In panel b) I consider average value added growth rates for standard and zero-rated traders separately (returns data). Growth rate of zero-rated traders is very volatile while that of standard-rated traders is more stable and follows the overall average in panel a)[24]. From this figure, it might seem that growth rate of standard-rated

regional variation

[24]The majority of traders are standard-rated, and therefore it is not surprising that the overall

traders start to pick up half way into the cut but it is impossible to claim this is the impact of the cut for two reasons. First, the growth rates for standard-rated traders remain negative and less than that of zero-rated traders throughout the cut. Second, the recession might have had a differential effect on standard-rated traders. The high degree of volatility in the zero-rated series prevents firm conclusions at this stage.

In panel c) I confirm the results of Crossley et al. (2013) by looking at durable and non-durable consumption growth from national accounts. Growth rate of durable consumption becomes positive half way into the cut while that of non-durable consumption remain negative for that period. It seems consumption data shows a positive impact of the cut on durable goods. However, looking at growth of durables before the cut, confirms higher volatility of this series and it could well be that durable consumption is just showing extreme volatility. Also growth rate of non-durables starts to pick up as well but at a slower rate. Therefore, this pattern might just be the natural evolution of the recession. Obviously, the correspondence between VAT rates and durable consumption is not clear cut and national accounts are not clearly comparable to returns data. For example, durable imports are part of durable consumption but are not directly included in returns data.

2.2.3 Other confounding policies

To tackle the financial crisis the government undertook many other policy reforms. For example, when the VAT cut was announced on 24 November 2008 the Chancellor also announced a rise in top income tax rate to 45 percent and an increase of 0.5 percentage point in national insurance both starting from 2011[25]. These tax changes were announced at the time of the VAT cut but become effective after the cut expiry[26]. Corporate tax rate was also reduced from 30 to 28 for profits greater than £1.5 million, and increased from 20 to 21 for profits less than £0.3 million for 2008/9 financial year. However, marginal corporate tax rates change even before the cut. Furthermore, standard and zero-rated sectors include all forms of ownership and not just incorporations.

growth rate of value added is similar to the standard-rated growth rate.

[25]The chancellor also announced £60 Christmas gift for pensioners (120£ for couples) on top of the usual £10 bonus.

[26]Under the assumption that a constant share of income is spent on standard and zero-rated categories, income changes would not create a heterogeneous impact on treatment and control. However, given the fact that necessities are zero-rated while luxuries are standard-rated, we might expect a differential spending response.

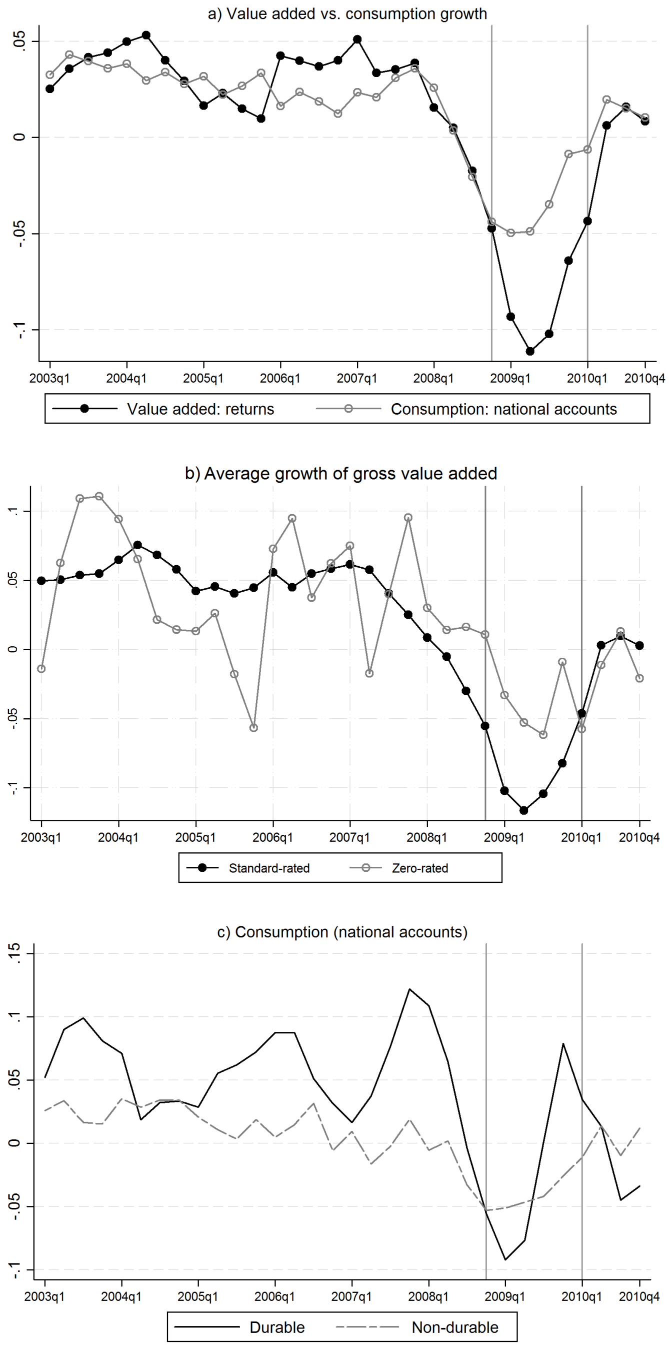

Figure 2.3: Change in value added and consumption (% on quarter a year earlier)

Notes: Figures show percentage change in variables in the given quarter relative to the same quarter one year earlier. In figure a) black line shows average change in log gross value added as observed in my data for the population of traders classified into standard or zero-rated on the sales dimension. Gray line shows percentage change in total consumption expenditure from national accounts (growth of ABJR series from ONS UK Economic Accounts data). In figure b) black line shows average change in value added for standard-rated traders while the gray line is for zero-rated traders. Figure c) uses classification of consumption into durables and non-durables from national accounts. Most durables are standard-rated items while non-durables (specially food) are mostly zero-rated. Two vertical lines show onset and end of VAT cut. Notice GDP is a real variable whereas gross value added is nominal.

Two policy reforms would have a more direct impact on the difference-in-differences estimation. First, in order to offset the effect of standard-rate cut on price of alcohol, tobacco, and fuel the government raised the excise duties on these items. Therefore, these products are essentially unaffected by the cut. Second, to help car manufacturers the government implemented a generous car scrappage scheme from May 2009 until March 2010. This scheme offered £2000 cash toward the purchase of a new car for customers with used cars with a minimum age of 10 years. I remove traders involved in both of these sectors in my analysis in order not to confound the impact of the VAT cut with these two changes.

2.3 Data

The data used in this paper is the universe of all returns submitted to HMRC between the first quarter of 2002 and final quarter of 2010[27]. Administrative VAT returns data is well suited for studying the impact of the cut for several reasons. First, I observe effective tax rates on sales and purchases and therefore, could identify standard-rated traders (i.e. treated). Second, amount of measurement error should be minimal relative to survey data because of potential penalties for mis-reporting. Third, I observe a large number of traders over the course of 32 quarters and therefore could control for a rich set of fixed effects (e.g. trader fixed effects and sector by date fixed effects) to alleviate concerns regarding confounding factors. The caveat of this data is that I do not observe quantities and prices separately. So I will not be able to separately identify price and quantity responses to the VAT cut[28].

Traders report net of tax sales, purchases, and the corresponding VAT on sales and purchases. I use three outcome variables to investigate the impact of the VAT cut: gross value added, gross sales, and gross purchases. I define gross value added as the difference between gross sales and gross purchases. Gross sales and purchases are respectively the result of adding up sales VAT to net of tax sales and purchases VAT to net of tax purchases (hereafter I drop the gross pre-fix). Sales, purchases, and

[27] I have access to returns from July 2001 to June 2011 but do not use observations prior to January 2002 or after December 2010. The beginning restriction is discretionary but does not affect any of the conclusions. The ending restriction is because from January 2011 the standard rate was increased to 20 percent. It is interesting to see the responses to the permanent increase in standard-rate from January 2011 but I do not have enough data points to identify any impacts.

[28] Office of National Statistics (ONS) compiles a set of price indices (e.g. producer price index) that could be merged with the returns data to convert nominal series into real (quantity) series. I have not yet pursued this avenue.

value added show high seasonal variation and are trending. In order to remove the trend and seasonality, I use the change in log of outcomes relative to the same quarter a year earlier as the main dependent variable and refer to it as percentage change in the original level variable[29]. I use the Δ_4 operator to denote this differencing: $\Delta_4 \ln y_{it} \triangleq \ln y_{it} - \ln y_{it-4}$, where y_{it} is the original outcome variable for trader i at date t.

I start from a total number of 66,375,762 returns between 2002q1 and 2010q4 and drop around 26 million observations (40 percent) in the following steps to arrive at a clean sample. I drop a) returns with zero reported sales ($\approx$10 million), b) majority exempt sectors (health, education, finance), alcohol, fuel, and tobacco related sectors, and sectors relating to wholesale and retail of cars ($\approx$6 million), c) forms of ownership other than sole proprietors, partnerships, and incorporations ($\approx$2 million), d) flat rate scheme traders ($\approx$4 million), and e) traders that could not be matched to a trader characteristics dataset ($\approx$ 4 million).

Table 2.2 reports summary statistics for the main variables. Mean value added is £150,133 but the distribution is very dispersed and the standard deviation is 300 times larger than the mean. Average sales and purchases are £627,955 and £482,127 respectively with a similar level of dispersion. The mean and standard deviation of log variables show much less dispersion (logs are defined only for positive values and therefore the number of observations is smaller). Standard and zero-rated traders constitute 55 and 14 percent of observations respectively. 31 percent of the sample are left unassigned. Standard-rated traders are on average smaller and have lower grow rates compared to zero-rated ones. Average growth rate of sales is respectively 1 and 2.1 percent for standard and zero-rated traders. Growth rates show high standard deviations. Standard-rated group is dominated by incorporated businesses whereas zero-rated traders are equally split between various forms of ownership[30].

Majority of VAT registered traders submit four returns during a year[31]. Due to the

[29] Notice the change in logs is equivalent to percent change in level variable only when the change is small. If the log change is $\beta = \Delta_4 \ln y_{it} = \ln y_{it} - \ln y_{it-4}$, then percent change in level variable is $\%\Delta_4 y_{it} = e^{\beta} - 1$ which is equal to β if β is small.

[30] The changes in number of observations from level variables to log versions and from log versions to differenced versions is due to negative values, and the requirement that traders need to have at least 4 quarters of data to get a non-missing growth rate. I have removed zero sales and therefore there is no change in number of observations for sales when I take logs. For value added, however, traders could have negative value added which would show up as missing when I take logs.

[31] A very small number of traders submit monthly or annual returns. The former group are often larger traders while the latter are smaller traders. I drop annual traders but keep monthly traders. To make them comparable to quarterly traders, I add up the value of three monthly returns in a calendar quarter.

Table 2.2: Summary statistics

	All traders			Standard-rated sales			Zero-rated sales		
Variable	Obs.	Mean	S.D.	% Obs.	Mean	S.D.	% Obs.	Mean	S.D.
Sales	38,952,778	627,955	6x10^7	55	445,605	2x10^7	14	867,202	6x10^7
Purchases	38,952,778	482,127	5x10^7	55	329,702	3x10^7	14	718,846	7x10^7
Value added	38,952,778	150,133	5x10^7	55	119,768	4x10^7	14	151,218	5x10^7
ln(sales)	38,952,778	10.71	1.73	55	10.82	1.63	14	10.51	1.94
ln(purchases)	38,276,569	10.13	1.90	55	10.10	1.85	14	10.31	1.89
ln(value added)	32,230,441	9.90	1.71	59	10.03	1.62	11	9.58	1.92
Δ_4 ln(sales)	32,450,522	0.0036	0.751	57	0.0100	0.664	15	0.0210	0.850
Δ_4 ln(purchases)	31,826,653	-0.0051	0.772	57	0.0006	0.759	15	0.0166	0.724
Δ_4 ln(value added)	24,501,010	0.0097	0.997	63	0.0142	0.931	11	0.0218	1.183
% incorporated	38,952,778	0.52		53	0.55		14	0.33	
% sole proprietor	38,952,778	0.29		53	0.29		14	0.32	
% partnership	38,952,778	0.19		53	0.16		14	0.35	

Notes: Summary statistics for level variables are calculated using observations between 2002q1 and 2010q4. Value added, sales, and purchases are all gross values, i.e. they include VAT.

sheer number of traders (around two million), VAT accounting periods are staggered within quarters. For around 38 percent of traders, accounting periods correspond to end of calendar quarters, i.e. end of March, June, September, and December. An equal share of the remaining traders submit returns at the end of the two other months in a calendar quarter. I collapse the data to trader-quarter observations, so I have one observation per quarter for each trader. For traders submitting returns that correspond to part of a calendar quarter I assume transactions are equally split between the three months covered in the return and take a weighted average of the two returns that have an overlap with the calendar quarter. For example, for traders submitting returns in February, I use a weight of two third on the February return and a weight of one third on the next period return (May return) to arrive at the adjusted returns for the first calendar quarter[32].

I define effective output tax rate, τ_o, to be the ratio of reported sales VAT to net Sales. I use average effective tax rates during the four quarters preceding the cut period to assign traders into treatment and control. Denoting the average of τ_o during 2007q4 and 2008q3 as $\bar{\tau}_o$. A trader is, respectively, defined to have standard and zero-rated sales if $\bar{\tau}_o \in [14, 18]$ and $\bar{\tau}_o \in [0, 4]$[33]. Figure 2.4 a) shows the distribution of effective output tax rates before and during the VAT cut. The before distribution (black line) shows more than 40 percent of traders have purely standard-rated sales (spike at 17.5) and 10 percent have pure zero-rated sales (spike at 0). Using the specified bands for treatment definition, 55 and 14 percent of observations are respectively assigned to traders with standard and zero-rated sales. About 31 percent of observations are left unassigned due to either traders not being in the data between 2007q4 and 2008q3 or traders with $\bar{\tau}_o$ outside the designated bands. The large spikes around standard and zero-rate suggest the banding I used for treatment definition are unimportant because it is unlikely that inclusion of few other traders, away from the spikes, has a significant impact on the results.

The binary classification of traders into standard and zero-rated groups leaves out traders with intermediate effective tax rates (i.e. those with $\bar{\tau}_o \in (4, 14)$). Bigger traders with a wide range of activities are likely to have intermediate tax rates

[32]The results with no adjustment or solely focusing on those with a perfect overlap with calendar quarters are very similar.

[33]In principle, I could use VAT law and relate 5-digit SIC codes to activities listed under different VAT rates. While this approach is feasible, certain zero-ratings cross the border of 5-digit SIC codes and make identification of pure zero and standard-rated sectors impossible. For example, most exports are zero-rated but any firm within any sector could be an exporter. Furthermore, there are exclusions within broad zero-rated sectors. For example, supply of food is generally zero-rated but items like ice creams, biscuits, cereal bars, etc are standard-rated.

and therefore might be excluded from the analysis. In order the see the impact of this exclusion, it is useful to check what percentage of total sales, purchases, and value added is removed from the analysis. It turns out that about 58-59 percent of total sales, purchases, and value added are included in the sample. This is 10 percentage points less than share of observations included in the sample (69 percent) and suggests that a higher number of large traders are being excluded from the analysis. I experimented with wider and narrower bands for definition of treatment traders and the quality of results remained unchanged[34].

One potential concern with the treatment definition is changes in composition of sales in response to the cut. For example, demand for standard-rated items might increase and therefore there might be an increase in effective output tax rates. The gray line in figure 2.4 a) shows the distribution of τ_o during the cut. There is essentially no change in fraction of zero-rated traders but the fraction of purely standard-rated traders is 10 percentage points lower than before the cut. While this might be suggestive of composition effects, closer investigation shows reduction in mass of purely standard-rated traders is due to transitions in and out of the temporary rate which distributes tax rates between 15 and 17.5. Furthermore, it seems the whole distribution is shifted to the left which is not consistent with a change in composition in favor of standard-rated items[35].

A related concern with the definition of treatment is stability of effective tax rates over time. Do the identified standard and zero-rated traders remain as such over time? To investigate this, table 2.3, shows transition probabilities for bands of effective output tax rate. The diagonal elements are the largest of row/column values and represent the probability of remaining in the same band as the last quarter. On average traders with τ_o within [0,4) remain in the same interval with 86 percent probability. Similarly traders with effective output tax rates within [17,18) continue to be in the interval with 88 percent probability.

A final concern with treatment assignment is the changes in the sample size. The mechanical effect of using average tax rates during 2007q4 and 2008q3 to assign traders to treatment is that there is a rise in the sample size up until these dates and a subsequent fall in the sample size after these dates. This is because of new traders' VAT registration prior to the treatment definition window. Newly registered traders

[34] I have carried out the analysis with [0,1] and [17,18] or [0,8] and [14,18] for zero and standard-rated assignments and the quality of the results are essentially unchanged.

[35] The temporary nature of the cut also reduces the possibility of a VAT induced product line switching for businesses.

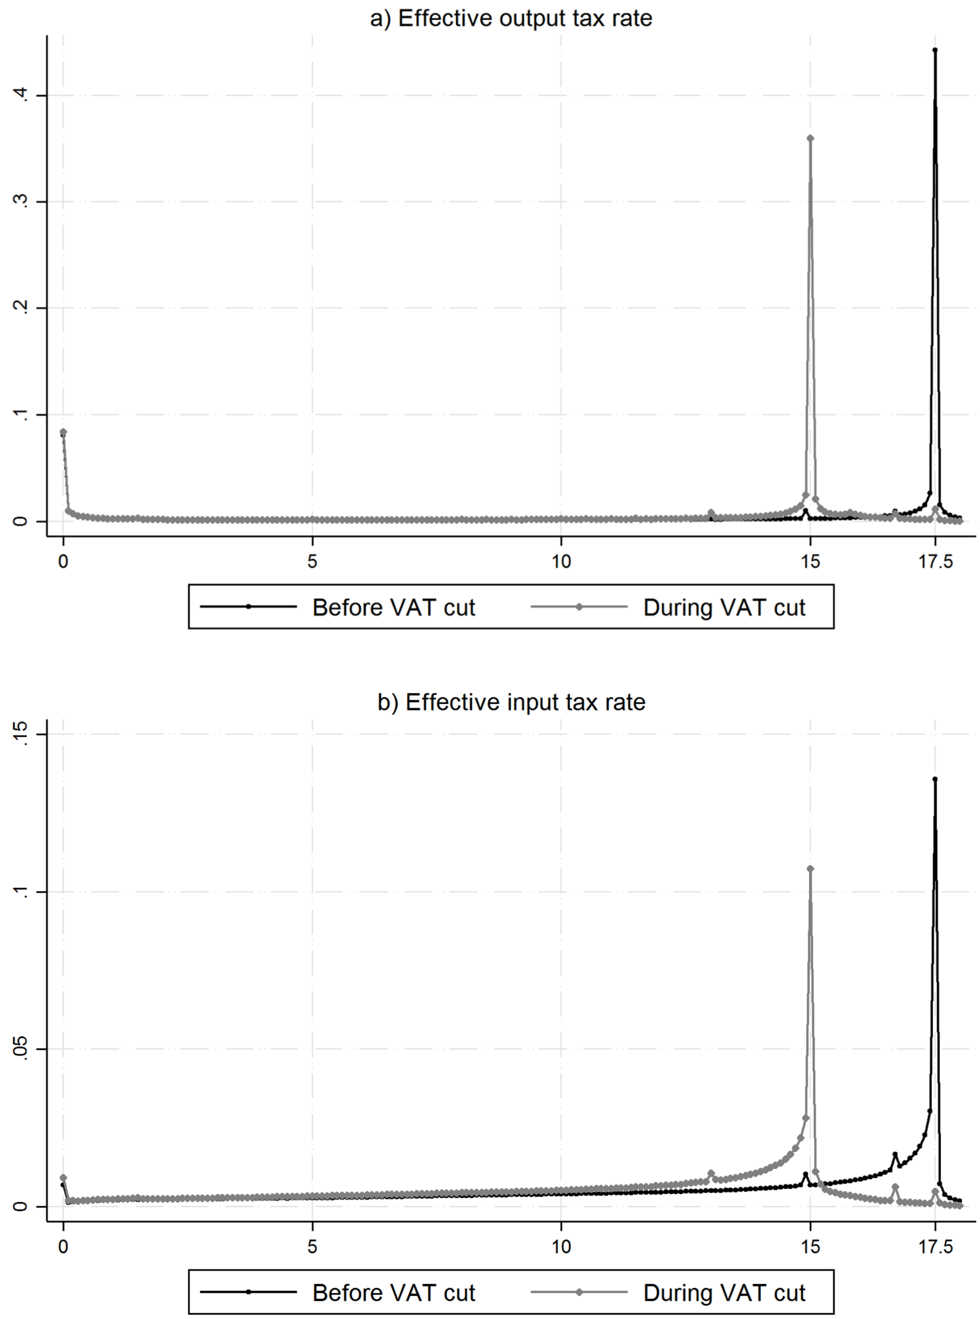

Figure 2.4: Distribution of effective output and input tax rates before and during the VAT cut

Notes: The bin width for distribution plots is 0.1 percentage point and the mass shows fraction of observations that fall in an interval centered around the indicated bin. Before period is from 2002q1 until 2008q3 and the during period is from 2008q4 until 2010q1. Effective output tax rate is calculated as the ratio of sales VAT to net sales (similarly for input tax rate). I am excluding observations between 2010q2 and 2010q4 from these graphs but their inclusion does not change any of the figures.

Table 2.3: Transition probabilities between bands of τ_o prior to VAT cut

Output tax rate		$\tau_o[t]$					
		[0,4)	**[4,6)**	**[6,15)**	**[15,17)**	**[17,18)**	**[18,∞)**
$\tau_o[t-1]$	**[0,4)**	**86.3**	4.05	5.5	0.9	2.84	0.4
	[4,6)	29.86	**35.21**	27.67	2.17	4.21	0.89
	[6,15)	6.54	4.61	**64.62**	9.91	12.76	1.57
	[15,17)	1.78	0.57	16.43	**45.82**	32.85	2.54
	[17,18)	0.91	0.19	3.35	4.97	**88.57**	2.01
	[18,∞)	2.8	0.81	8.75	8.76	49.79	**29.09**

Notes: the cells show probability that a trader with effective output tax rate within a given band switches to another band in the next quarter. Diagonal elements show probability of remaining in the same band.

are considered for treatment assignment but only show up in the analysis from the time they register for VAT. Furthermore, some of the traders who were present during 2007q4 and 2008q3 exit the sample after this date. To see the potential impact of a changing sample size, I have experimented with a balanced panel of traders that appear in all dates and the results remain the same. However, the balanced panel restriction removes about half of the traders. Therefore, I decided to present the more inclusive results with the unbalanced panel.

So far I have defined treatment based on whether sales are standard-rated. However, the standard-rate cut also affects input tax rates leading to potential behavioral responses. In principle, I could break down the sample to four groups based on standard and zero-rated distinction along the sales and purchases dimension. Two by two comparisons of these groups could deliver estimates of behavioral responses on the two dimensions of treatment (and their interactions). Two features of my setting, however, makes this approach unreliable. First, unlike sales, majority of businesses use a range of inputs and there seems to be little specialization in input use (as expected). Figure 2.4 b) plots the distribution of effective input tax rate, τ_i, defined as the ratio of purchases VAT to net purchases. τ_i distribution is much more dispersed and the standard and zero-rate spikes are much smaller. Before the cut 13.1 percent of traders used purely standard-rated inputs and less than 1 percent used only zero-rated inputs. This means categorization of traders to standard and zero-rated purchases leaves out a large part of the sample. Second, the joint density of τ_o and τ_i does not show enough concentration of traders for the four groups. Specifically the only group of traders with a large fraction is traders with standard-rated sales and purchases. The mass is particularly small for traders with zero-rated

Table 2.4: Joint density of effective output and input tax rates before cut period

		τ_i						Marginal
		[0,4)	**[4,6)**	**[6,15)**	**[15,17)**	**[17,18)**	**[18,∞)**	
τ_o	**[0,4)**	**4.6**	1.6	6.5	1.6	1.5	0.3	16.1
	[4,6)	0.4	**0.4**	1.1	0.3	0.3	0.1	2.4
	[6,15)	0.9	0.9	**8.6**	2.2	1.7	0.4	14.8
	[15,17)	0.4	0.4	3.8	**2.5**	1.5	0.3	8.9
	[17,18)	2.3	1.9	17.8	11.6	**19.9**	1.9	55.4
	[18,∞)	0.1	0.1	0.9	0.4	0.6	**0.3**	2.4
Marginal		8.8	5.2	38.6	18.6	25.4	3.3	-

Notes: Cells show fraction of observations that have an effective output and input tax rate within the bands specified by row and column headings.

sales and standard-rated purchases. Therefore, I do not carry out a two by two analysis and to control for potential responses arising from the purchases dimension, I restrict attention to the sub-sample of traders with standard-rated or zero-rated purchases[36] in some of the figures.

2.4 Empirical Strategy

The main empirical strategy used in this paper is difference-in-differences (DD) estimation using zero-rated traders as a control group for standard-rated ones. I compare the average change in outcomes for standard-rated traders during the cut period to the average change for zero-rated traders to estimate the stimulus effect. This strategy could be implemented in a regression as follows

$$\Delta_4 \ln y_{ist} = \alpha_i + \beta_t + \gamma_1 \text{SRated}_i \times \text{Cut}_t + \epsilon_{ist} \tag{2.1}$$

where $\Delta_4 \ln y_{ist}$ is the change in log of value added (or sales or purchases) for trader i operating in sector s in date t relative to four quarters earlier, α_i is trader fixed effects, β_t is date fixed effects, and $\text{SRated}_i \times \text{Cut}_t$ is the interaction of a dummy variable for traders with standard-rated sales with a dummy showing the duration of the cut period. The coefficient of interest is γ_1 and shows the differential change in growth rates for standard-rated traders during the cut period. Trader fixed effects

[36] This is defined in a similar way as standard-rated sales. A trader has standard-rated purchases if the average input tax rate during 2007q4 and 2008q3 is around the standard-rate, i.e. $\bar{\tau}_i \in [14, 18]$

control for anything that is constant over time and has an influence on growth rate of value added. For example, larger traders might have slower but more stable growth rates on average. Time fixed effects control for any event that affect standard and zero-rated traders to the same extent.

To claim that γ_1 is an unbiased estimate of the stimulus effect, I need to assume that in the absence of the cut the change in growth rates would have been the same for standard and zero-rated traders. This is obviously a contentious assumption given the fact that the cut was in response to the great recession. Some of the zero-rated activities relate to necessities like food while some of standard-rated traders are involved in sales of durable goods. The recession might have a stronger impact on standard-rated traders because of more elastic demand. Therefore, the confounding recessionary effect could induce a downward bias on the estimates of γ_1 from the basic specification in (2.1).

I experiment with several extensions of the basic specification to control for the confounding recession effect. From figure 2.3 it seems the recession started to impact aggregate variables from 2008q1. Therefore, in the first extension I introduce a second interaction term that allows for a heterogeneous impact of the recession on standard-rated traders from 2008q1 onwards

$$\Delta_4 \ln y_{ist} = \alpha_i + \beta_t + \gamma_1 \text{SRated}_i \times \text{Cut}_t + \gamma_2 \text{SRated}_i \times \text{Rec}_t + \epsilon_{ist} \tag{2.2}$$

where Rec_t is a dummy that is switched on from 2008q1 onward. Effectively, this specification relies on three quarters before the start of the VAT cut and three quarters after the end of the cut to identify the recession effect[37]. Under the assumption that the differential recession effect remains the same during and outside the cut period, estimates of γ_1 from (2.2) would give causal impact of the cut. If, however, the differential impact of the recession is changing over time the estimates here are still biased.

In the second extension, I allow two-digit sectors to receive heterogeneous impacts from the recession by including interactions of two-digit sector dummies with the recession dummy.

$$\Delta_4 \ln y_{ist} = \alpha_i + \beta_t + \gamma_1 \text{SRated}_i \times \text{Cut}_t + \delta_s \times \text{Rec}_t + \epsilon_{ist} \tag{2.3}$$

[37] The estimated magnitude of γ_1 from this regression is identical to a regression based on equation (2.1) where the sample is restricted to the post recession dates (other coefficient estimates would be different).

where δ_s is a set of 66 two-digit sector dummies. The potential factors controlled in this specification are slightly different from (2.2). As far as standard and zero-rated traders within the same two-digit sector are subject to the same recession effect, the remaining within sector differences between the two groups capture the causal effect of the cut. If however, sectors with majority standard-rated traders experience a greater recession impact right in the middle of the cut period, this specification would fail to give the causal estimates of the cut.

2.5 Results

In this section I first present graphical evidence on the response of firms to the VAT cut. In the second subsection I present regressions using firm level data. I investigate the responses of three dependent variables to the VAT cut: growth of sales, growth of purchases, and growth of value added. The main variable of interest is growth of sales because this variable is directly affected by the rate cut and would be the first to respond to changes in consumer spending. Purchases and value added would receive an indirect impact from the changes in sales via the production function.

I also present graphs and regressions for various sub-samples. I start from the largest sample and look at evolution of the dependent variables but this is not necessarily the most interesting sample. Under normal VAT accounting, any input VAT could be reclaimed. Therefore, while collection of VAT is through the production chain its final burden would be on consumers. If the VAT chain is unbroken and in the absence of evasion, only final demand should respond to VAT rate changes. Of course, final demand responses would feedback to intermediate demand but the response would be dissipated. Therefore, the second sub-sample I investigate is the sample of retailers (or business to customer sectors as I define later)[38].

The third sub-sample I study is the sample of traders with standard-rated purchases. For most of the results, classification into treatment and control is based on VAT rates on sales. The standard-rate change could, however, change the applicable tax rates on purchases. Focusing on sales and ignoring tax rates on purchases could lead to confounding effects. Therefore, to alleviate such concerns I restrict the sample to traders with a similar input tax rate. The final sub-samples I consider split the traders based size (turnover). Larger firms could change their prices more quickly

[38] I start off with the whole sample because the VAT system in the UK features certain exemptions that could result in responses for intermediate sectors.

or increase the salience of the cut through advertising. Therefore, there might be a larger impact for bigger traders. I split the sample into traders with average annual sales below and above £2.8 million[39] to investigate this possibility.

2.5.1 Graphical evidence

Figure 2.5 shows evolution of average growth rate of sales, purchases, and value added in three panels. Panel a) shows average sales growth for standard-rated traders (black line) is around 5 percent up until 2007q2. From 2007q3 the growth rate starts to decline and turns negative from 2008q3. After the start of the VAT cut (first vertical line) average growth of sales continues to fall and the decline only stops after 2009q2. By the end of the VAT cut (second vertical line) sales growth is -3 percent and right after the end of the cut, when you expect to see a backlash in standard-rate activities, it jumps to 2 percent and remains there for the following quarters.

While this pattern seems inconsistent with a positive impact of the rate cut, it is impossible to conclude anything in the absence of a valid counterfactual. The gray line in figure 2.5 shows the average growth rates for zero-rated traders which is used to build a counterfactual. The zero-rated series is much more volatile, partly due to a smaller sample[40]. Nevertheless, it seems average growth rate for zero-rated traders starts to fall approximately around 2008q4 but the decline in growth rates seem to be smaller (and much more volatile). During the recession a clear gap emerges between the two series which suggests standard-rated traders experienced lower growth rates during this time. It seems the start of the cut does not have an effect on the gap and the two series keep their distance. However, after the cut and towards the end of the recession the two series converge.

On the figure, I also report the estimated impact of the VAT cut using a simple DD. Specifically the reported number on the graph shows coefficient estimates (robust standard errors) on the interaction term from a regression of the plotted variable on a cut dummy, standard-rated dummy, and their interaction. For sales this number is -0.037 with a standard error of 0.023 which suggests growth rate of sales was 3.7 percentage points lower for standard-rated traders during the VAT cut period (insignificant). Including a recession interaction term would increase this coefficient to 0.03 but it remains insignificant. Therefore, it seems sales growth is not different

[39]This is one of several criteria that define a small business in the UK.

[40]Average per quarter observations for zero-rated traders is around 76,000 while for standard-rated traders it is 441,000.

for standard-rated traders during the cut. Obviously the deepest and possibly most heterogeneous part of the recession could happen during the cut and this estimate could only reflect that.

Panels b) and c) in figure 2.5 show evolution of average growth rates for purchases and value added[41]. A very similar pattern is observed here. Average growth rates for both of these variables fall more for standard-rated traders relative to zero-rated ones and the estimated stimulus impacts are around -3.5 and -5.5 percentage points for purchases and value added respectively. It is interesting to note that purchases series show a perfect co-movement right up to the recession where the gap emerges. Including a recession interaction in panels b) and c) would increase the coefficients to -1 and -5 percentage points for purchases and value added respectively. The impact on value added remains significant in the two specifications but is insignificant for purchases (at 5 percent). This suggests the recession induced gap between standard and zero-rated series remains constant before and during the cut period for purchases but expands for value added.

In summary, figure 2.5 shows during the recession a clear gap emerges between average growth rates of standard and zero-rated traders. The size of the gap seems to remain stable for sales and purchases while for value added it is expanding during the cut. Obviously one could argue that the recession induced gap would have expanded right at the time of the VAT cut had the VAT cut not been implemented. While I am unable to rule this possibility out, the convergence of the three series right after the recession and the absence of a set-back period supports a zero impact on sales and purchases. In what follows, I discuss similar graphs for the sub-samples of retailers, traders with standard-rated purchases, and large traders.

Restricting to retail sector

In order to see if final demand responses are different for standard and zero-rated traders, figure 2.8 looks at average growth rates of sales, purchases, and value added for the sub-sample of retailers[42]. All series have similar growth rates for standard and zero-rated retailers up until 2007q2. At this point a gap appears in the average growth rates with the growth rate of zero-rated sales and purchases increasing while the opposite happens for standard-rated retailers. During the cut period sales and

[41]Notice ln(value added) = ln(sales − purchases), therefore the magnitudes in panel c) are not directly calculated from panels a) and b).

[42]This is all 5-digit codes within the 74 two-digit SIC2007 code.

purchases growth rates are negative for both groups and they broadly follow a similar U pattern as for the overall averages in figure 2.5. DD estimates show that during the cut period the growth rates of sales and purchases is 1 and 1.6 percentage points lower for standard-rated retailers. At the time of the cut, value added growth jumps sharply for zero-rated retailers while that of standard-rated retailers only picks up slowly. While the movement for zero-rated traders seems a bit puzzling, it does not seem these figures are very different from figure 2.5. Furthermore, it is hard to see a positive impact of the VAT cut from these graphs.

Controlling for input tax rates

One potential concern with figure 2.5 is that the VAT cut affects tax rates on both purchases and sales. In figures 2.6 and 2.7 I respectively restrict to the sub-sample of traders with standard and zero-rated purchases and consider whether among these groups those with standard-rated sales saw a positive impact from the VAT cut. It is worth noting that the sample size here is much reduced because of the higher dispersion of effective input tax rates and therefore results may not be directly comparable to figure 2.5.

Figure 2.6 shows average growth rates for sales, purchases, and value added for the sub-sample of traders with standard-rated purchases. The zero-rated series are less volatile compared to figure 2.5 and overall it seems all three growth rates are similar up until the recession. Interestingly, for sales (panel a) and value added (panel c) standard and zero-rated series closely track each other even during the cut, while purchases (panel b) show a recession induced gap similar to what was observed in figure 2.5. Simple DD coefficients are also smaller here. Sales growth is estimated to be 1.5 percentage points lower for standard-rated traders during the cut, while purchases and value added growth are 3.7 and 0.8 percentage points lower. None of these coefficient are significant, though. A somewhat similar, but noisier picture emerges when I restrict attention to the group of traders with zero-rated purchases in figure 2.7.

Responses of large and small firms

Figure 2.9 shows the three dependent variables for large and small traders (as defined above). Large standard and zero-rated traders have very similar average growth rates (across panels) right up to the time of the VAT cut. During the cut, sales

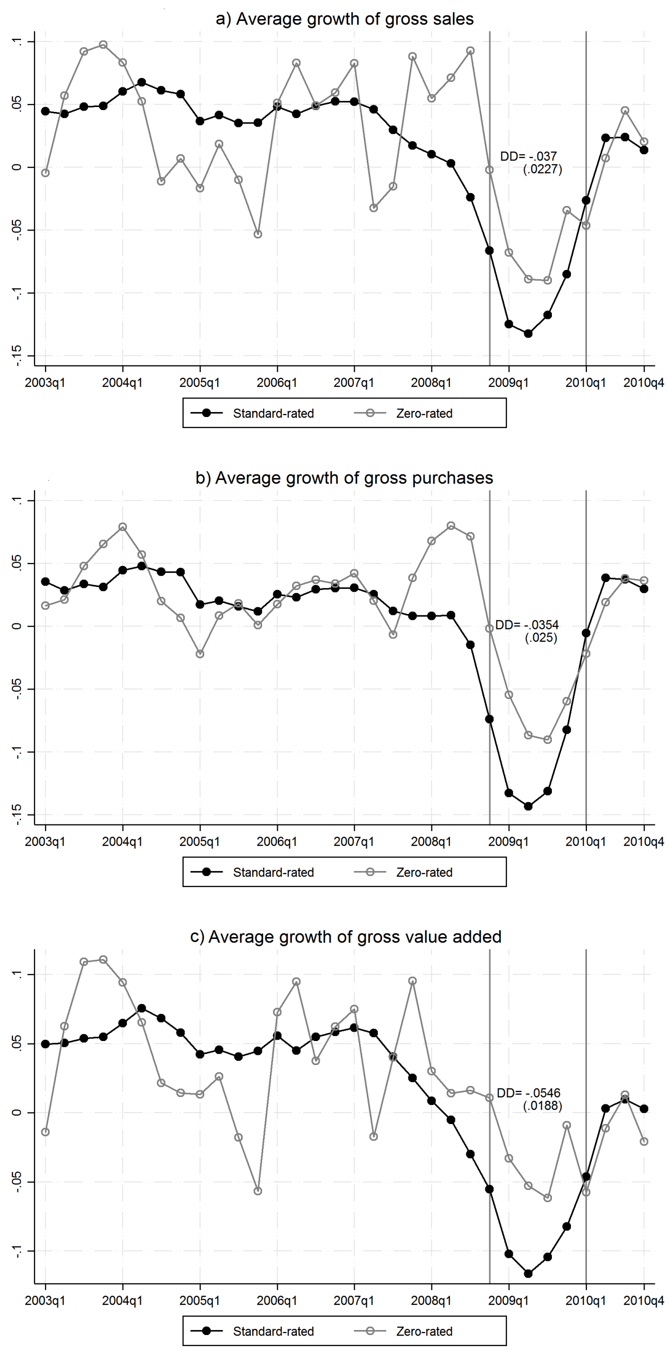

Figure 2.5: Change in log sales, purchases, and value added for standard and zero-rated traders

Notes: Graphs show average change in log of a) sales, b) purchases, and c) value added for standard and zero-rated traders over quarters. I use average effective output tax rate (ratio of sales VAT to net sales) during the four quarter preceding the cut (2007q4-2008q3) to define standard-rated traders. The trader is classified as standard-rated when this variable is between 14 and 18 percent and as zero-rated when it is between 0 and 4 percent. The reported DD estimates are coefficient estimates (robust standard errors) of the interaction term in a regression of plotted variable on standard-rated dummy, cut period dummy, standard-rated×cut period dummies, group specific quarter dummies and linear trends. The regression uses aggregated data with 64 observations. The first and second vertical lines mark 2008q4 and 2010q1 corresponding to the first and final quarter were the VAT cut has any effect. Total observations for log value added is 14,683,318 standard-rated trader returns and 2,454,875 zero-rated ones, approximately corresponding to 524k standard-rated and 82k zero-rated unique traders during the 4 quarters preceding the cut.

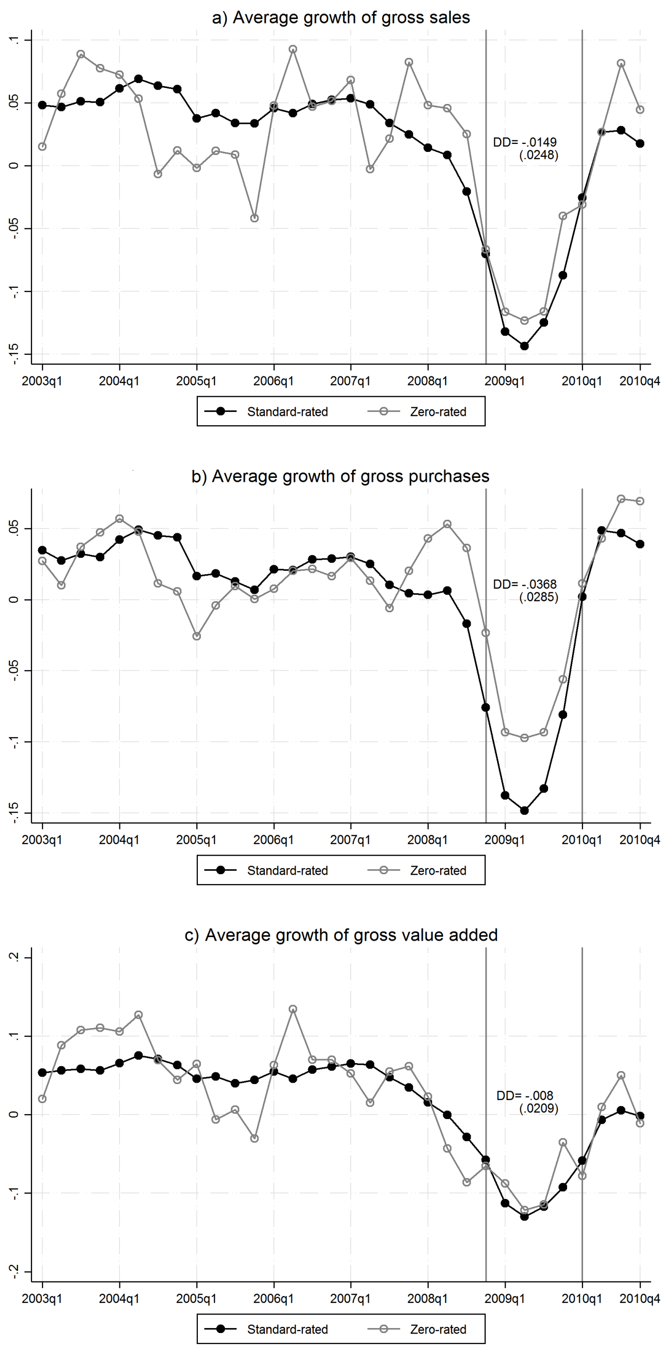

Figure 2.6: Change in log sales, purchases, value added (restrict to traders with standard-rated purchases)

Notes: This is a similar figure to figure 2.5 except for the fact that the sample is restricted to traders with standard-rated purchases (i.e. average effective input tax rate during 2007q4-2008q3 between [14,18] percent. See figure 2.5 for notes defining standard and zero-rated traders on the sales dimension.

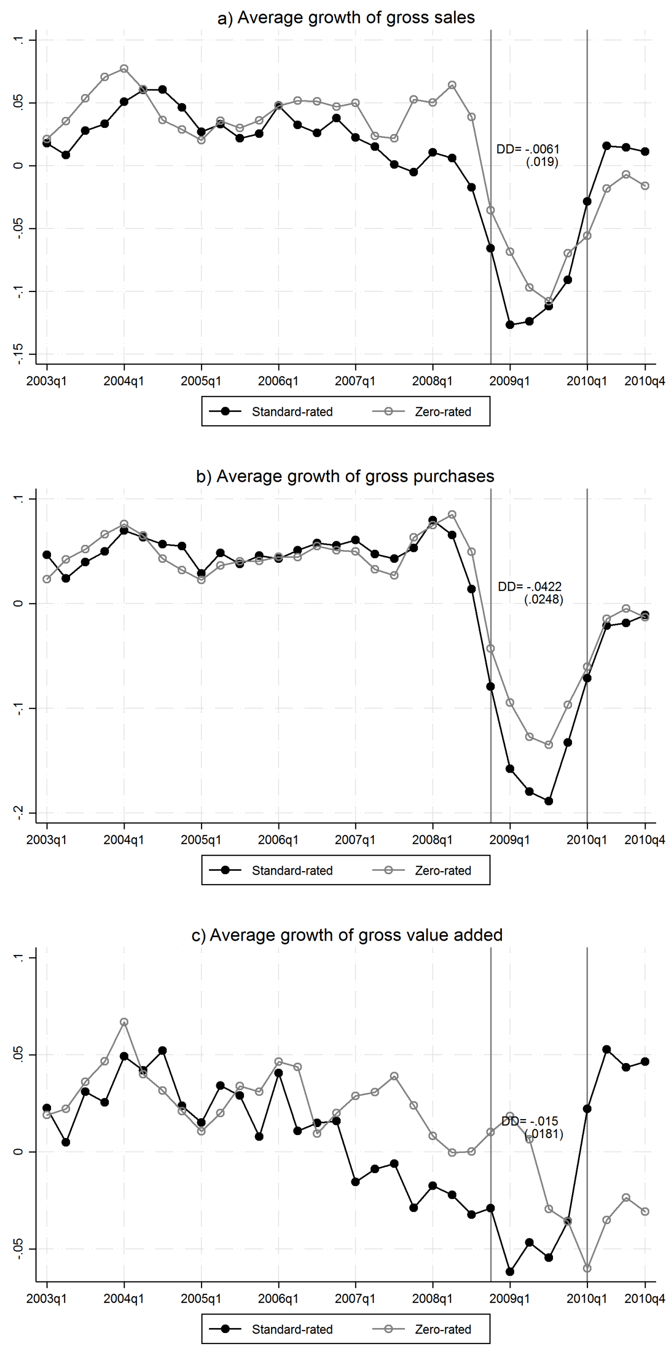

Figure 2.7: Change in log sales, purchases, and value added (restrict to traders with zero-rated purchases)

Notes: This is a similar figure to figure 2.5 except for the fact that the sample is restricted to traders with zero-rated purchases (i.e. average effective input tax rate during 2007q4-2008q3 between [0,4] percent. See figure 2.5 for notes defining standard and zero-rated traders on the sales dimension.

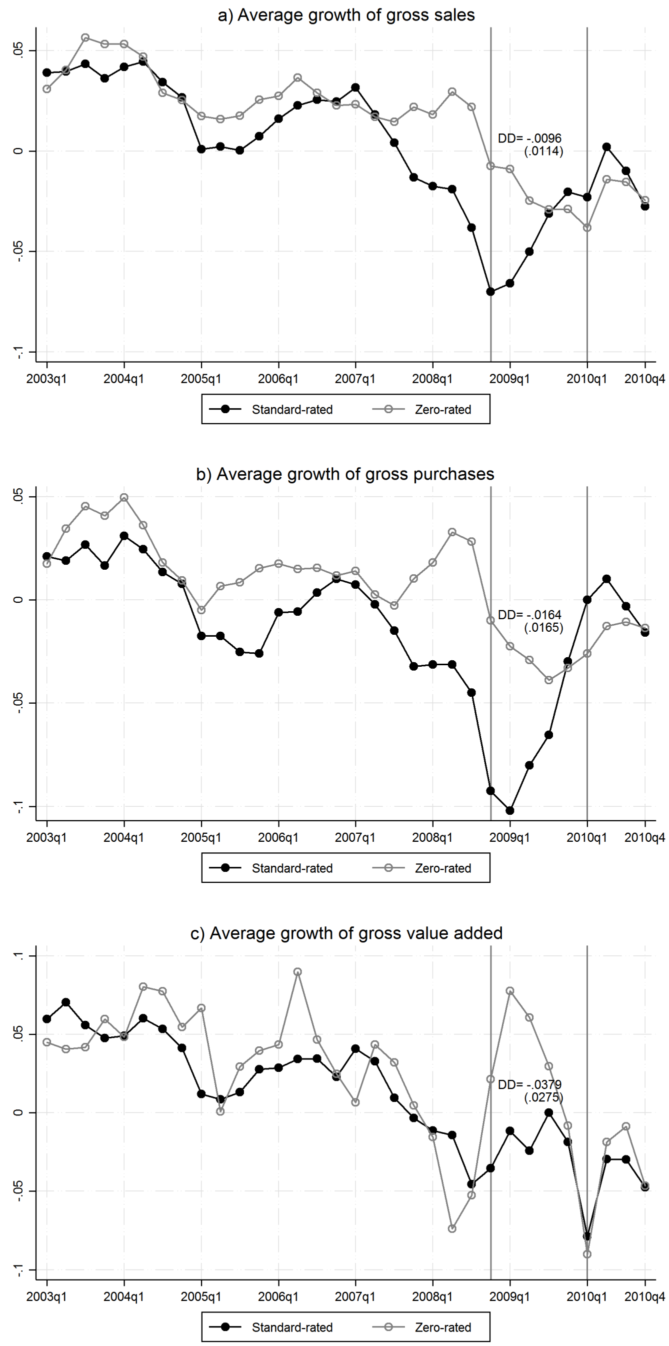

Figure 2.8: Change in log sales, purchases, and value added (restrict to retail sector)

Notes: This is a similar figure to figure 2.5 except for the fact that the sample is restricted to traders within the retail sector (SIC2007 codes of 47XXX). See figure 2.5 for notes defining standard and zero-rated traders on the sales dimension.

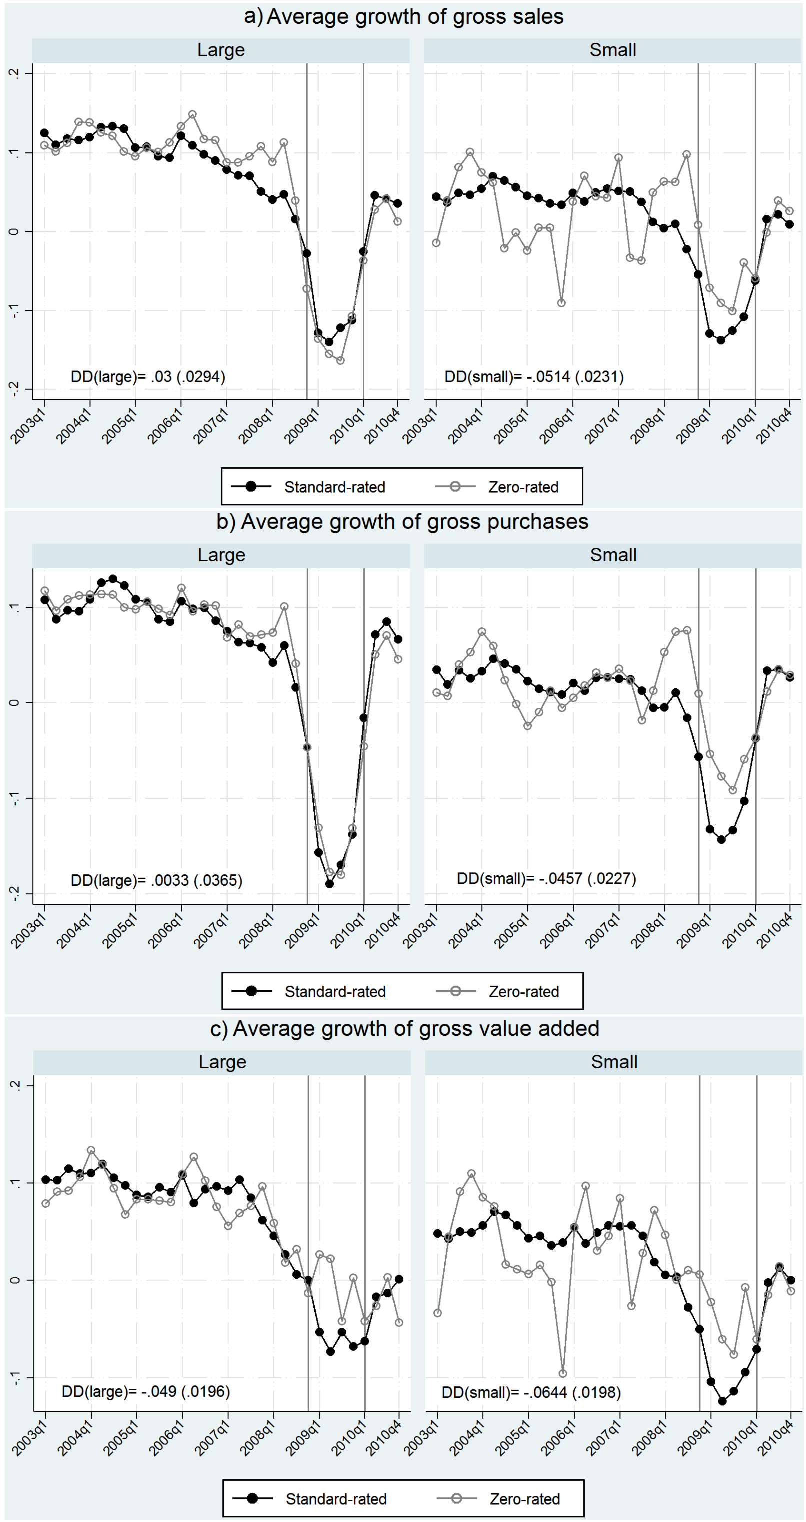

Figure 2.9: Change in log sales, purchases, and value added (Large vs. small traders)

Notes: This is a similar figure to figure 2.5 except for the fact that the sample is split into small and large traders. A trader is small if its average gross annual sales is less than £2.8 million during all years of appearance in the data. See figure 2.5 for notes defining standard and zero-rated traders on the sales dimension.

growth rate for large standard-rated traders is 3 percentage points higher (albeit insignificant). Purchases growth rate seems to be unaffected (DD estimate is 0.3 percentage points). However, when I consider growth rate of value added, the DD estimate and the figure reveal a significant negative impact of 4.9 percentage points[43]. The pattern of movements for small traders is very similar to the overall picture in figure 2.5 and does not support a positive impact of the cut.

2.5.2 Regression evidence

Table 2.5 shows the estimation results for the full sample. Columns (1) and (2) report DD estimates in the absence of recession controls. Column (1) uses a basic DD (OLS) specification and confirms coefficient estimates for the interaction term are similar to those reported in figure 2.5. Controlling for trader and date fixed effects in column (2) makes the estimated magnitudes larger (FE specification). For the three outcome variables these two specifications deliver implausibly large negative estimates of the cut impact.

Columns (3) - (5) control for recession heterogeneity. Columns (3) and (4) include an interaction of recession dummy with the standard-rated dummy in the OLS and FE specifications. The estimated stimulus effect for sales and purchases is now very small and insignificant. It seems the difference between standard and zero-rated traders emerges before the start of the cut and there is not a discernible impact on sales and purchases for standard-rated traders during the cut. To take magnitudes seriously, DD estimates suggest growth rates of sales and purchases for standard-rated traders are respectively 0.2 and -0.9 percentage points different from zero-rated traders (column (3)). The cut, however, continues to have a negative impact on value added growth (Panel C). Value added growth rate is 4.6 percentage points lower for standard-rated traders during the cut (significant at 5 percent). If the recession had a heterogeneous impact on standard-rated traders but the heterogeneity remained on average the same across the set of recessionary quarters, then group specific recession dummies would fix the identification problem and the coefficient estimate of S-rated

[43] As discussed earlier $\Delta_4 \ln(\text{value added}) = \Delta_4 \ln(\text{sales-purchases}) \neq \Delta_4 \ln(\text{sales}) - \Delta_4 \ln(\text{purchases})$. Therefore, I will not be able to relate panel b) and c) directly to panel a). Qualitatively, however, I expected a positive impact on value added. It might be that changes in number of observations and the way I calculated growth rates are responsible for these patterns. Specifically, not all observations present in panel b) and c) are used in panel a). For example, traders with negative value added would be excluded from panel a) because ln(value added) is missing for them. This reduction in sample size is non-trivial specially because the recession could change the number of traders with negative value added differentially across the two groups.

Sales ×Cut term in columns (3) and (4) would capture the causal effect of the VAT cut.

In column (5) instead of a recession interaction, I include the interactions of 65 two-digit sector dummies with the recession dummy to allow for a heterogeneous impact of the recession across sectors (SIC2d specification). This has a small impact on estimates. The fact that estimates remain remarkably similar across columns (3) - (5) is encouraging. However, none of these would deal with a time varying recession heterogeneity. In other words if the recession starts to disproportionately affect standard-rated traders right at the time of the VAT cut, estimation results under columns (3)-(5) would be biased. To increase reliability of estimates, I experimented with a more flexible specification where I controlled for sector by date fixed effects. This specification allows sectors to evolve freely in each quarter, effectively estimating stimulus effect from within sector-date differences between standard and zero-rated traders. While this specification is still vulnerable to within-sector heterogeneous impacts on standard-rated traders, it is quite rich and deal with many possibilities. Unfortunately estimating this specification with more than 2000 regressors is computationally demanding and I only ran this specification for a 10 percent sample of the data. The (unreported) results from these regressions were broadly similar to the results in column (5).

Table 2.6 reports estimates of the stimulus effect, i.e. coefficient of S-rated Sales×Cut, for various sub-samples. I have reported estimates from two specifications for the three dependent variables. FE and SIC2d specifications correspond to specifications in columns (4) and (5) of table 2.5. FE specification controls for recession heterogeneity using the interaction of S-rated Sales dummy and Rec dummy. SIC2d specification uses interactions of two-digit sector dummies with the recession dummy to isolate the recession effect. First row, reports the coefficient estimates for the whole sample as a benchmark while the other rows show estimates for other sub-samples.

The second row focuses on retailers to see whether traders at the end of the VAT chain receive a higher impact. I estimate a very small negative insignificant effect on sales and purchases. Value added growth, however, shows a larger reduction but still insignificant. Another way to identify traders dealing with final consumers is to use Input-Output tables and classify sectors based on the share of final demand. I use Input-Output tables from ONS for the year 2007 and classify sectors into being business to customer (B2C) if the share of final demand is greater than 50

Table 2.5: Regression results for the whole sample

	Basic		Recession control		
	OLS	FE	OLS	FE	SIC2d
	(1)	(2)	(3)	(4)	(5)
Panel A: Δ_4 ln (Sales), 21,598,298 observations (991,690 traders)					
S-rated Sales× Cut	-0.037*	-0.057**	0.002	-0.0003	-0.001
	(0.016)	(0.019)	(0.012)	(0.012)	(0.012)
S-rated Sales× Rec			-0.052**	-0.081**	
			(0.015)	(0.020)	
R-square	0.010	0.010	0.010	0.010	0.011
Panel B: Δ_4 ln (Purchases), 21,192,026 observations (980,959 traders)					
S-rated Sales× Cut	-0.036*	-0.050**	-0.009	-0.010	-0.007
	(0.015)	(0.017)	(0.012)	(0.012)	(0.011)
S-rated Sales× Rec			-0.035**	-0.058**	
			(0.012)	(0.015)	
R-square	0.007	0.006	0.007	0.006	0.007
Trader and Date FE	N	Yes	N	Yes	Yes
SIC2d×Rec	N	N	N	N	Yes
Panel C: Δ_4 ln (Value added), 16,555,223 observations (949,438 traders)					
S-rated Sales× Cut	-0.056**	-0.074**	-0.046*	-0.046*	-0.044**
	(0.015)	(0.016)	(0.018)	(0.018)	(0.015)
S-rated Sales× Rec			-0.014	-0.041*	
			(0.013)	(0.016)	
R-square	0.004	0.005	0.005	0.005	0.006

Notes: Table shows coefficient estimates and standard errors from estimation of five specifications for three outcome variables. The dependent variables in panels A, B, and C are respectively Δ_4 ln (Sales), Δ_4 ln (Purchases), and Δ_4 ln (Value added). Column (1) estimates a basic DD specification with treatment dummy, cut dummy, and their interaction. Column (2) includes trader and date fixed effects (specification 2.1). In columns (3) - (5) I control for recession heterogeneity. Column (3) adds a recession dummy and its interaction with standard-rated dummy to the OLS specification. Column (4) adds recession interaction to the fixed effects specification (specification 2.2). Column (5) includes interactions of 65 two-digit sector dummies with the recession dummy as in specification 2.3. Cut dummy is equal to 1 for all quarters between 2008q4 and 2010q1 and zero otherwise. Recession dummy is equal to 1 for all quarters after 2008q1 and zero otherwise. Standard-rated sales dummy is defined in the text. All standard errors are clustered at 5-digit SIC2007 codes (around 570 clusters).* and ** show coefficient estimates are significant at 5 and 1 percent confidence levels respectively.

Table 2.6: Coefficients and standard errors for DD estimate of the cut impact

Dep. Var.:	Δ_4 ln **(Sales)**		Δ_4 ln **(Purchases)**		Δ_4 ln **(Value added)**	
	FE	SIC2d	FE	SIC2d	FE	SIC2d
Sample	(1)	(2)	(3)	(4)	(5)	(6)
Whole	-0.0003	-0.001	-0.010	-0.007	-0.046*	-0.044**
	(0.012)	(0.012)	(0.012)	(0.011)	(0.018)	(0.015)
Retail	-0.0030	-	-0.009	-	-0.055	-
	(0.011)	-	(0.010)	-	(0.042)	-
B2C	0.041*	0.028	0.020	0.014	-0.012	-0.022
	(0.016)	(0.016)	(0.015)	(0.014)	(0.028)	(0.021)
S-rated purchases	0.015	0.014	-0.006	-0.016	-0.026	-0.007
	(0.020)	(0.020)	(0.018)	(0.016)	(0.057)	(0.041)
Large	0.019	0.023	0.0009	0.008	-0.062	-0.048*
	(0.017)	(0.017)	(0.016)	(0.017)	(0.033)	(0.023)
Trader and Date FE	Yes	Yes	Yes	Yes	Yes	Yes
S-rated Sales× Rec	Yes	N	Yes	N	Yes	N
SIC2d×Rec	N	Yes	N	Yes	N	Yes

Notes: Each cell shows coefficient estimates and standard errors on the interaction of cut and standard-rated dummies from a separate regression. Columns (1), (3), and (5) use FE specification allowing for heterogeneous impact of the recession on standard-rated traders (same specification as in column (4) of table 2.5). SIC2d specification (columns (2), (4), (6)) estimate equation (2.3) where I include interaction of two-digit sector dummies with the recession dummy (same specification as in column (5) of table 2.5). The first row replicates estimates of interaction terms for the whole sample from column (4) of table 2.5. The second row restricts attention to the retail sector (SIC2007 equal to 47XXX, 11 percent of the whole sample). The third row shows results from regressions on sectors identified as business to customer sector, i.e. with share of final demand greater than 50 percent from input-output tables 2007 (38 percent of whole sample). The fourth row shows results for the sub-sample of traders with standard-rated purchases (around 50 percent of whole sample). The fifth row restricts to large traders defined as those with average annual sales greater than £400,000 (32 percent of the whole sample). Cut dummy is equal to 1 for all quarters between 2008q4 and 2010q1 and zero otherwise. Recession dummy is equal to 1 for all quarters after 2008q1 and zero otherwise. Standard-rated sales dummy is defined in the text. All standard errors are clustered at 5-digit SIC2007 codes. The number of clusters are 570 for the whole sample, 558 for standard-rated purchases sample, 42 for retail sample, 255 for B2C sample, and 556 for large trader sample. * and ** respectively mark significance at 5 and 1 percent levels.

percent[44]. The third row shows the results for B2C sectors. Interestingly, estimates show a positive impact on sales and purchases growth during the cut period. FE specification shows sales growth was 4.1 percentage points higher (significant at 5 percent) for standard-rated traders. SIC2d specification delivers a smaller and insignificant effect here. Purchases and value added impacts are not significant at 5 percent.

The fourth row restricts to traders with standard-rated inputs and the fifth row considers large traders. Coefficient estimates for sales are between 1.4 to 2.3 percentage points which is larger than the whole sample but all are insignificant. For purchases the coefficient estimates are all very small and insignificant. Value added growth seems to be the only variable showing a significant reduction during the cut period.

To summarize, table 2.6 shows sales growth for standard-rated traders was either the same or slightly better than zero-rated ones across the sub-samples studied. Value added growth on the other hand show a negative impact across the sub-samples[45]. Most of the estimates are, however, not significantly different from zero at 5 percent. The findings that sales and purchases growth rates are not significantly different for standard-rated traders does not necessarily suggest the stimulus was ineffective in boosting the real economy. Under full pass through the VAT cut would mechanically lower gross sales for standard-rated traders by 2.1 percent, therefore a zero impact on gross sales suggests quantity demanded has increased by 2.1 percent. In other words, the price elasticity of demand for standard-rated products is -1 and price reductions lead to proportionate increases in demand. Similarly the 2.8-4.1 percentage points estimates for B2C sectors would translate to a 4.9-6.2 percentage points increase in quantity demanded under full pass through. These estimates correspond to an elasticity between -2.3 and -2.9.

[44] IO tables provide data on 110 sectors. This relates to two or three-digit SIC2007 codes. About 59 out of 110 sectors are identified to have a greater than 50 percent share of final demand. Some of these B2C sectors are a) products of agriculture, hunting and related services, b) preserved meat and meat products, c) dairy products, d) textiles, e) wearing apparel, f) furniture, g) gambling and betting, h) repair of computer and personal and household goods.

[45] The negative estimates for value added series are probably the least reliable because the definition of growth rates implies a much lower number of observations for this series. From table 1.3 we can see that total observations for growth rates of sales and purchases are around 32 million while for value added growth I have only 24 million observations. This is due to many traders having more purchases than sales (hence negative value added).

2.6 Conclusions

In this paper, I used a difference-in-difference estimation strategy to identify the stimulus impact of 2008 VAT rate cut in the UK. Graphical evidence suggests the recession had a stronger impact on standard-rated traders. During the recession sales, purchases, and value added growth is lower for standard-rated traders relative to zero-rated ones. This suggests a simple DD estimate would confound the recession effect with potentially positive effects of the VAT cut. Regression results confirm this intuition with implausibly large negative estimates.

However, after I allow the recessionary period to exert a heterogeneous influence on growth rates of standard-rated traders, the growth rate of sales and purchases seem to be similar for standard and zero-rated traders. In other words, the cut does not seem to have had a significant positive effect. This suggests the cut has boosted standard-rated demand just to compensate for the reduction in prices leading to an underlying price elasticity of -1. Restricting the sample to potentially more responsive groups (retailers and large traders) I estimated positive but insignificant effects on sales which alludes to a price elasticity less than -1. While these findings suggest the temporary standard-rate cut was effective in boosting real activity of the standard-rated traders (especially the retailers), my inability to provide direct evidence on price and quantity changes prevents firm conclusions.

It is worth emphasizing that the recession poses a real challenge to the identification of the impact of the VAT cut. The specifications estimated in this paper allow for some forms of heterogeneous recession effects but due to the high overlap of the recessionary period and the cut it is impossible to prove causality. In the most stringent specification, I allow two-digit sectors (66 of such codes) to have different average growth rates during the recessionary period. This specification relies on the variation between standard and zero-rated traders within the same two-digit sector during the cut to identify the stimulus effect. To the extent that the recession effects are sector specific I could claim causal estimation of the cut impact.

The other issue with the difference-in-differences methodology is its inability to detect across the board effects. For example, the cut reduces prices and frees up income to be spent on any item (income effect). If consumers decide to spend this on both zero and standard-rated products DD would not pick up any effect while the overall impact of the cut is positive.

Chapter 3

Educational Impact of Iran Iraq War

3.1 Introduction

Events such as wars, natural disasters, and pandemics could have long lasting effects on individual well-being. On the macro side these events could shift the equilibrium of the economy and leave local economies in a poverty trap. Empirical literature, however, was generally unable to provide support for this theoretical possibility[1]. From a micro perspective catastrophic events could impact exposed individuals in the long run even though they have no detectable aggregate effect. Young individuals who are still in the process of accumulating human capital are particularly vulnerable to negative shocks. Destruction of schools, interruption of classes, loss of teachers, loss of family members, and loss of household income are a few mechanisms that could reduce educational attainment of young individuals. Exposure to catastrophes could also affect individuals' health which itself could have an influence on educational outcomes. Given the large economic literature on importance of events during mother's pregnancy and before age of 5 for adulthood outcomes, the health mechanism is particularly relevant[2].

---

[1]For example, Davis and Weinstein (2002) and Davis and Weinstein (2008) find no evidence of multiple equilibria in the context of allied forces bombing of Japanese cities. Japanese cities converge to their pre-bombing population trends in the long run. Miguel and Roland (2011) are unable to uncover local poverty traps for heavily destroyed areas after the Vietnam War. Bosker et al. (2007), however, seem to find some evidence of multiplicity for German cities subject to WWII destruction.

[2]For example, Almond and Mazumder (2005) and Almond (2006) study the 1918 influenza pandemic in the US, Almond et al. (2009) investigate the effects of Chernobyl's radioactive radiations, Almond et al. (2010) consider Chinese famines, and Almond et al. (ming) and Almond and Mazumder (2011) study the impact of fasting during pregnancy on children. All these studies de-

This paper looks at educational attainment of Iranian children exposed to Iran Iraq War (IIW) 18 years after the end of the war. In September 1980 large scale Iraqi invasion of Iranian territory marked the beginning of an eight year war. By June 1982 the war displaced more than 1.6 million individuals across five war-hit provinces, approximately 22 percent of the population living in these provinces. Furthermore, most cities in these provinces came under aerial attacks or artillery fire. While there is a vast literature on analysis of motivations, operations, and strategic implications of the war, there is little work on economic impact of the war. I provide first reduced form estimates of impact of IIW on educational attainment[3].

I use a 2 percent sample of the 2006 Iran Population Census and compare high school graduation rates for children exposed to war to those not exposed. Date of birth and place of residence jointly determine whether a child was exposed to IIW. Therefore, I employ a difference-in-differences (DD) estimation strategy and compare war time cohorts across war and non-war provinces to pre-war cohorts. I distinguish between early childhood and school time exposure to war to provide separate estimates of the war impact on cohorts born during the war (1980-86 birth cohorts) and cohorts that went to school during the war (1963-1979 birth cohorts). The large literature on importance of early childhood events suggests that the chaotic war-time environment should have a negative impact on physical and psychological development of very young children. On the other hand the large scale displacement of individuals could have interrupted schooling and led to negative effects for school cohorts. I would be able to provide a comparison of early childhood and school time effects which might be useful in formulating mitigation policies for similar catastrophic events.

The DD estimates show that the probability of high school graduation is reduced by 4.8 percentage points (significant at 10 percent) for the cohorts born during the war while there is only a 1.9 percentage points (insignificant) reduction for the cohorts that spent some of their school years during the war. In my sample 38.8 percent of individuals have completed high school, therefore, these numbers correspond to 12.4 and 4.9 percent reduction in high school graduation. The estimates suggest early childhood effect is 2.5 times higher than the school time effect. The early childhood

tect significant large impacts of the early childhood event on adult human capital and labor market outcomes. For review of the literature see Almond and Currie (2011a) and Almond and Currie (2011b).

[3]My search of the Farsi and English literature has returned no study of educational impact of the war. Mofid (1990) and a few other Farsi publications provide aggregate estimates of the economic cost of the war. A handful of articles studied the impact of exposure to chemical warfare during IIW on health outcomes (e.g. Ahmadi et al. (2010), Ahmadi et al. (2009), Kadivar and Adams (1991), and Khateri et al. (2003)). Mahvash (2011) studies impact of the war on divorce patterns.

effect is robust with respect to several alternative specifications. For example, when I control for differential trends across war and non-war provinces, the war impact continues to be significant for early childhood cohorts but becomes insignificant for school cohorts.

To interpret these estimates as causal, I need to rule out several potential confounding factors. First, the 2006 Census does not provide data on wartime residence of individuals. Therefore, I rely on birth locations to identify whether individuals were in war provinces during the war. Furthermore, I only know the birth place of individuals who are living in their birth place in 2006. I denote these individuals as non-migrants and restrict the sample so I can define treatment status properly (61 percent of all individuals are non-migrants). Non-migrants in war provinces may not be comparable to non-migrants from non-war provinces because war induced many individuals, who would not have migrated otherwise, to migrate out of war provinces. If the well-endowed individuals are more likely to migrate and permanently settle out of war areas, the sample restriction would imply a downward bias for the estimates of war impact.

Two pieces of evidence relieve some of the concerns arising from the sample restriction. First, province-level migration figures from 1986 and 1996 Censuses suggest war provinces are being de-populated during the war and then partially re-populated after the war. Therefore, at least part of the migrants have returned to their home[4]. Furthermore, intra-province migration figures are higher for war provinces during the war compared to non-war provinces. This abnormal pattern continues to the 1996 Census round but in the 2006 Census, intra-province migration rates for war and non-war provinces become very similar. The intra-province migration patterns suggest many war migrants were settled in the same province during the war and probably have returned to their homes later. Second, there is not a discernible difference between the fraction of non-migrant individuals within each birth cohort across war and non-war provinces. In other words, the same share of individuals from each cohort is included in the sample across war and non-war provinces. Unfortunately, these patterns cannot fully rule out the possibility that war induced migration could be responsible for the estimated effects.

The rest of the confounding factors that might affect causal interpretation of my estimates are simultaneous events that might have a heterogeneous impact on cohorts

[4]The reverse pattern is observed for some of the provinces neighboring war provinces which shows temporary settlement of war migrants in neighboring provinces and their subsequent return once the war is finished.

in war provinces. The DD identification requires that in the absence of IIW the educational gap between war and non-war provinces stays the same for treated and control cohorts (parallel trends). Therefore, anything that happens at the same time as the war is a potential confounder.

The first simultaneous event is a dramatic increase in population growth during 1976 and 1986. Average yearly population growth during this decade was 3.9 percent while in the preceding and proceeding decades it was respectively 2.7 and 2.5 percent. Interestingly, the baby boom resulted in similar birth increases in war and non-war provinces. But, the larger negative impact of the war on early childhood cohorts may simply reflect inability of war provinces to accommodate the baby boom. To alleviate this concern, I collect province-level yearly figures on number of schools and students at the primary level. Once I include these measures of educational resources in the regressions the DD estimates remain the same.

The second simultaneous event is a series of ethnic rebellions that started right after the revolution. Short lived rebellions happened in Khuzestan, Azerbaijan, and Sistan but Kordestan uprising was the most prominent and continued until 1982. When I exclude Kordestan from the sample, the war effect changes slightly. Furthermore, I notice that the war impact seems to be increasing over time which is in constrast to an expected impact of the rebellions that finished by the end of 1982.

This paper speaks to the vast literature on the impact of early childhood circumstances on adult outcomes and confirms conclusions in the literature that very young children are more vulnerable and could suffer long lasting effects from catastrophic events. On the specific subject of conflict, there are several papers that estimate impact of conflict on educational attainment of children using DD methodology. In a cross country setting, Ichino and Winter-Ebmer (2004) compare Austria and Germany to countries not involved in WWII. They find that school age children exposed to WWII attained lower education relative to non-war cohorts. They also find significant earning losses 40 years after the war that could be attributed to lower educational attainment of these cohorts. Using within country variation Shemyakina (2011) estimates that Tajikistan civil war had a significant impact on enrollment of girls and their rate of finishing mandatory schooling but she does not find any impact on boys[5].

[5]Several other papers employ a DD methodology and find significant negative impact of conflict on educational attainment of children in various contexts: see Akresh and Walque (2008) for the impact of Rwandan Genocide, Blattman and Annan (2010) for child soldiering in Uganda, Merrouche (2006) for effect of landmines in Cambodia, and Chamarbagwala and Moran (2010) and

The rest of the paper is organized as follows. The next section gives a brief overview of Iran's education system and IIW. Section 3 and 4 describe the data and the identification strategy. In section 5 I present the graphical evidence and regression results. Section 6 discusses the issue of sample selection and other simultaneous events that could result in biased DD estimates. The last section concludes.

3.2 Context

3.2.1 Education system in Iran

Establishment of modern primary and secondary schools in Iran dates back to the beginning of the twentieth century. In 1910 the Ministry of Education was founded and one year later, with the passage of Fundamental Law of the Ministry of Education, primary education became compulsory and free of charge. This was equivalent to 6 years of education but due to lack of access and insufficient resources even literacy rates remained low (Arasteh (1962)). In 1943 the Law of Compulsory Free Public Education once again commissioned the government to expand free compulsory primary education to all areas within 10 years (Menashri (1992)). It also stipulated a fine for preventing children from attending schools. However, as figure 3.1 shows literacy rate only started to increase gradually. Starting from around 40 percent in 1940, the literacy rate rose to 88 percent in 1970[6]. Primary completion rates remained 15 percentage points below literacy rates during the same period.

In 1971 the education system was restructured to three levels: 5 years of primary, 3 years of intermediate, and 4 years of high school (table 3.1). This reform also extended free compulsory education to the end of the intermediate level. This change seems somewhat effective with the high school completion rate starting to accelerate after 1970 but intermediate school completion rate never exceeds 80 percent (figure 3.1). In 1992, and in response to high failure rates, the high school level was transformed to a unit-based system.

Chamarbagwala and Moran (2011) for impact of Guatemala civil war. There are also a few studies that look at other dimensions of human capital like health. As an example Akbulut-Yuksel (2010) finds significant impact of allied bombing on children educational attainment, health and adult labor market outcomes in Germany during WWII. She attributes the educational impact to the physical destruction of schools an teacher absence and the health impact to malnutrition during WWII.

[6] The rapid rise in literacy over cohorts is due to both expansion of education system over time and a successful adult education campaign after 1979 revolution.

Table 3.1: Evolution of education system in Iran

	Initial system	Re-organized system	New high school system
Period	1911-1970	1971-1991	1992-2006
Primary school	6 years	5 years	5 years
Intermediate school	-	3 years	3 years
High school	6 (3+3) years	4 years	4 (3+1) years

Notes: During the initial system, the upper high school level was reformed several times through additions of various majors. The 1971 re-organization of the system started in 1967 by changing primary schools to a 5-year system. The first intermediate schools opened in 1971, while the first 4-year high schools opened in 1974. The re-organized system was subject to various changes at the high school level. The new high school system was approved in 1990 to reduce grade repetition partly in response to the baby boom that inflated the number of students going to high school. Under the new high school system, students need to retake only the courses they could not pass during a school term. Whereas earlier students had to repeat the whole grade if they failed a number of courses. The role out of the new high school system started in 1992 with 10 percent of first year high school students. By 1998 all high school grades were functioning under the new system. Source: Menashri (1992) and various educational laws from the website of the Islamic Parliament Research Center.

Children start grade 1 of primary level at age 6 and with no grade repetition will graduate from grade 12 at age 18. At the end of each level students sit through centrally administered exams to obtain the relevant degree. For most of my sample high school diploma is awarded in grade 12[7]. Grade repetition was not uncommon during the period of analysis. Figure 3.1 shows a clear dip for the last few data points in each series which is indicative of grade repetition or late start. However, for high school, full grade repetition is less likely partly due to the introduction of the unit-based system[8].

3.2.2 Iran Iraq War (IIW)

Iran Iraq relationship was very contentious right from Iraq's independence in 1932. The major source of dispute was over the control of the bordering river, Arvand-Rud. However, except for a few skirmishes the relationship was by and large peaceful. The main agreement during this period was the Algiers Agreement in 1975 that

[7]In most pre-1992 years this corresponds to the 12th grade. Post-1992, diploma was awarded after successfully finishing 11th grade. In the new system the 12th grade was designed to prepare students for entering university.

[8]A study by the Islamic Parliament Research Center suggests around 2.5 and 7.8 percent of primary and intermediate students had to repeat a grade due to failure during 2002-2006 school years. After the introduction of the new high school system in 1992, grade repetition is very uncommon as students need to retake only the failed courses.

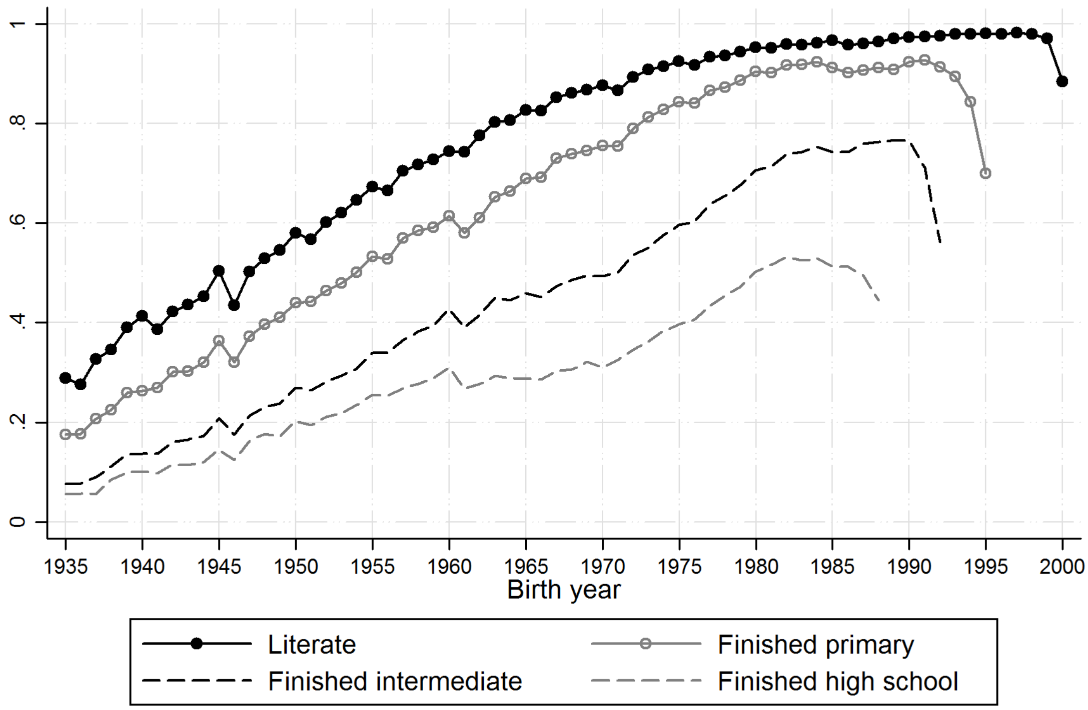

Figure 3.1: Expansion of modern education in Iran

Notes: Figure shows fraction of individuals with the specified degree for each birth cohort using the 2 percent sample of 2006 Census used in this paper. I restrict to individuals who are 6 years or more at the time of the census for the fraction of literate individuals. This corresponds to those born up until 2000. Similarly for fraction of individuals finishing primary, intermediate, and high school I respectively restrict to those aged 11, 14, and 18 years old in 2006 Census. This corresponds to 1995, 1992, and 1988 birth cohorts. The figure starts from 1935, corresponding to cohort of 71 years old individuals. For cohorts that studied under the old system the equivalent level is calculated. For example grade 8 of the old system corresponds to the final year of intermediate level.

set the frontier along the thalweg in Arvand-Rud allowing Iran to freely use the river's navigational routes. The 1979 Islamic revolution in Iran and the subsequent instability, however, encouraged Saddam to denounce the Algiers Agreement and to engage in an unprecedented large scale war lasting for about 8 years and claiming 213,255 lives on the Iranian side[9].

On 22 September 1980 Iraq started an ambitious ground invasion of Iranian territory along the 650 miles border. Until November 1980 Iraq captured vast swathes of Iranian territory including ten important cities and came close to a few major cities[10]. The advancement of Iraqi forces soon came to a halt and after some unsuccessful offensive during 1981, Iran was able to recover most of the occupied territory (including some major cities) until June 1982. From this time until the signing of UN's 598 resolution and the subsequent cease fire on 20 August 1988, there was virtually little territorial exchange and the war continued with attacks and counter-attacks along the border.

From the beginning till the end of the war all bordering villages and cities were battle fronts subject to constant shelling, aerial, and ground attacks. I use the five officially war hit provinces of Khuzestan, Ilam, Kermanshah, Kordestan, and West Azerbaijan as treated areas in my analysis (figure 3.2). However, many industrial and civilian centers well inside the country were targeted by aerial and missile attacks during the war, especially in 1985, 1987 and 1988 during the so called episodes of war of cities[11].

3.3 Data

The variables for my analysis are coming from a 2 percent sample of individual records of 2006 Iran Population Census from Statistical Centre of Iran (SCI)[12]. 2006 Census administered an extended questionnaire to about 20 percent of randomly selected households. Current data is a 10 percent extraction of this sub-sample. The

[9] Many books and articles are written on the background of the conflict and the development of the war during its 8 years. See Bakhash (2004), Cordesman (1987), Souresrafil (1989), Karsh (2002) and Hiro (1989) for detailed chronologies of war events and Cordesman and Wagner (1990) and Potter and Sick (2004) for in-depth analysis of war events.

[10] The captured cities are Khorramshahr, Susangerd, Bostan, Mehran, Dehloran, Ghasreshirin, Howeize, Naftshahr, Sumar, and Musian. The cities subject to continuous shelling are Abadan, Ahwaz, Andimeshk, Dezful, Shush, Islamabad, and Gilangharb.

[11] The exact timing and location of missile attacks could be used as an alternative identification strategy for studying the war impact. I am building a database of all missile and aerial attacks on Iranian cities to conduct this analysis.

[12] This is freely available from Statistical Centre of Iran in Farsi and from IPUMS in English.

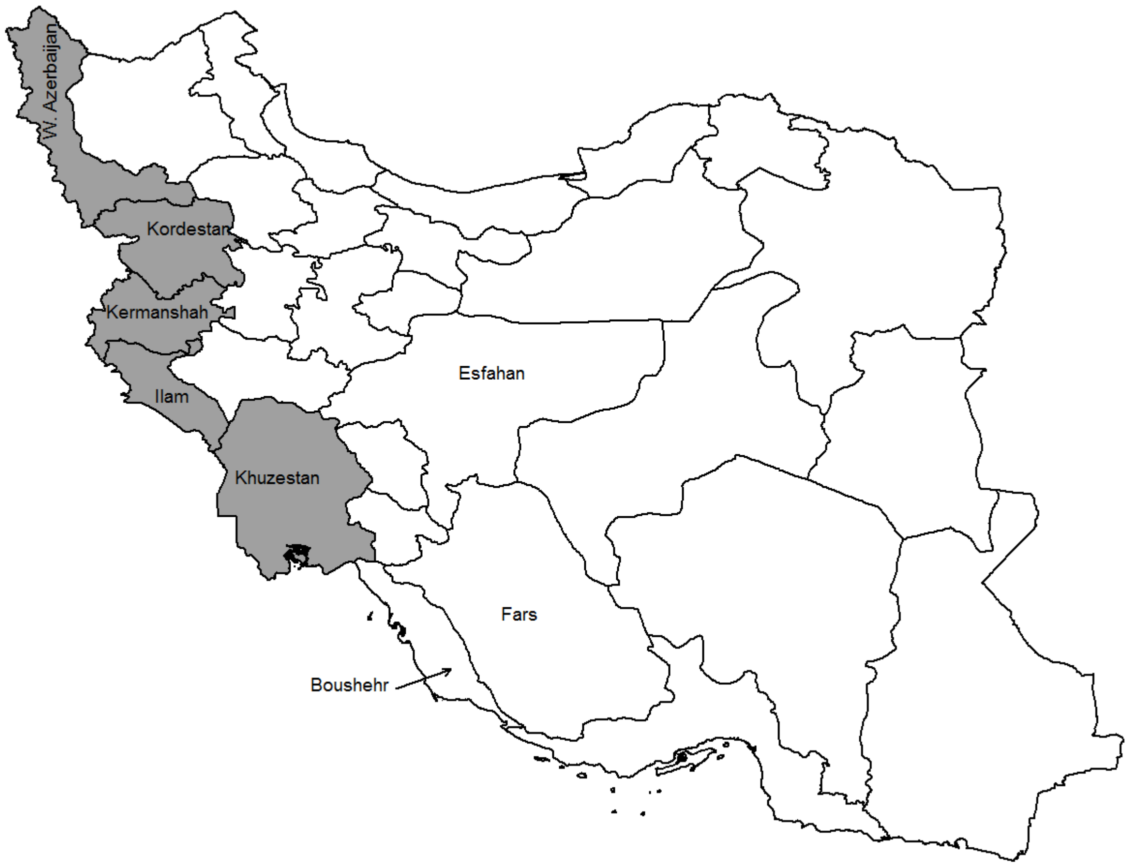

Figure 3.2: War hit provinces

Notes: Figure shows a map of Iran provinces. The grayed areas are the five provinces officially declared as war hit.

sampling unit is a household and it is stratified at district by urban location. It provides data on current residence, date of birth, migration during the past 10 years, educational attainment, employment status and other characteristics. The sample within each stratum is random but SCI provides individual probability weights (i.e. inverse of sampling probability) that could be used to calculate nationwide aggregate statistics. All calculations presented in this paper use these weights unless stated otherwise but the results without weights are very similar[13].

The main variable used for educational attainment is a dummy that shows whether the individual has finished high school. I focus on high school graduation because primary and intermediate graduation rates are quite high among the young cohorts (figure 3.1) and show little difference between war and non-war provinces[14].

I restrict the working sample to individuals aged between 20 to 66 years in 2006. Children aged 6 need to enroll for the first grade of primary school and with no grade repetition, they would finish grade 12 by the age of 18. To minimize the impact of grade repetition I restrict the sample to individuals aged 20 or more in 2006. On the other hand, very old cohorts have a small sample size and also have very low high school completion rates, therefore, I restrict to cohorts aged 66 or less. These individuals were expected to finish high school in 1958 when average literacy and high school graduation rate were about 72 and 28 percent (figure 3.1)[15].

The school year begins on 23 September each year and ends in June next year, therefore, the age conditions outlined above are based on age as of 22 September. This is also the way I define birth cohorts throughout the paper. For example, all individuals born between 23 September 1939 and 22 September 1940 are assigned to the 1940 birth cohort and will start primary school in 23 September 1946, i.e. 1946 school year. Therefore, the age restriction above is equivalent to constraining the analysis to 1940 and 1986 birth cohorts.

Panel A and B in table 3.2 show summary statistics for these variables in the full sample and with the above age restriction. The restricted sample has higher educational attainment. While on average 76 percent of individuals are literate in the full sample, the restricted sample has 84 percent literacy rate. Similarly 61 percent of individuals finished primary in the full sample while this number is 74 percent in the

[13] Average high school completion rates for very old cohorts is lower without using the weights (across war and non-war provinces).

[14] Educational variables are derived from a single coded variable in the original dataset. I could also use years of education but the mapping from the coded variable to years of education is less clear and is subject to greater error.

[15] Lowering or increasing this upper age by 10 years does not change the results significantly.

restricted sample. Rate of high school graduation is 0.13 percentage points higher in the restricted sample but unemployment rates are very similar (12 vs. 11 percent). Restricted sample is also on average older (due to removal of the large number of individuals aged 0-19).

To match the 2006 Census records to war measures I would need to have information on the residence of individuals during the war. Unfortunately, Census records data only on current residence and whether the individuals are living in their birth place. Therefore, I could only identify birth place of those living in their birth place in 2006[16]. Therefore, I restrict the sample to these individuals and use the term non-migrants to refer to them[17]. I acknowledge this could pose a serious challenge to my identification and discuss some supporting evidence in section 3.6.1.

Panel C in table 3.2 shows summary statistics for non-migrants born between 1940 and 1986. About 61 percent of individuals in sample B are non-migrants, i.e. appear in sample C. Interestingly, samples B and C do not show considerable differences for most of the variables in the table. The non-migrant sample seems to be slightly more educated, has higher unemployment rate, is younger, and has a lower share of urban individuals[18].

3.4 Empirical Strategy

I employ a difference-in-differences (DD) identification strategy to estimate the war impact on educational attainment. I compare the difference between average high school completion rates for cohorts exposed to war to those not exposed across war and non-war provinces[19]. The war started in 1980 and theoretically could impact all individuals under the age of finishing high school. The oldest cohort that could potentially receive an impact is the cohort of individuals in their 12th grade in 1980. These individuals are 17 years old in 1980 and therefore correspond to the 1963 birth cohort. The youngest cohort that is affected by the war is the 1988 birth cohort,

[16]Since the war has ended 18 years before the Census I am unable to use migration questions (which relate to past 10 years) to identify war time residence of all individuals.

[17]Technically some of these individuals could be return migrants, i.e. those who have returned to their birth places after a temporary leave.

[18]This is mainly due to high rural to urban migration rates in Iran.

[19]Implementation of a DD strategy for estimating war impact on high school completion is feasible because in the absence of recall bias it does not matter whether I measure completion rates in 2006 or exactly at the time the individuals have finished high school. Using a similar strategy for unemployment rate is not feasible because age has an effect on unemployment.

Table 3.2: Summary Statistics

Variable	All cohorts			Born between 1940-1986					
	A: Whole sample			B: Everyone			C: Non-migrants		
	Obs.	Mean	S.D.	Obs.	Mean	S.D.	Obs.	Mean	S.D.
Literate	1,357,394	0.76	0.43	711,779	0.84	0.37	463,552	0.85	0.36
Primary	1,357,394	0.61	0.49	711,779	0.74	0.44	463,552	0.76	0.43
Intermediate	1,357,394	0.41	0.49	711,779	0.54	0.5	463,552	0.56	0.5
High School	1,357,394	0.24	0.43	711,779	0.37	0.48	463,552	0.39	0.49
Years of schooling	1,357,394	5.96	5.28	711,779	7.63	5.52	463,552	7.85	5.46
Unemployment	441,083	0.12	0.33	357,405	0.11	0.31	242,719	0.13	0.33
Age	1,357,394	28.13	18.64	711,779	35.85	12.03	463,552	34.64	12
Male	1,357,394	0.5	0.5	711,779	0.5	0.5	463,552	0.49	0.5
Family size	1,357,394	4.72	2	711,779	4.52	1.93	463,552	4.59	1.98
Urban	1,357,394	0.69	0.46	711,779	0.71	0.45	463,552	0.65	0.48
Head is in birth place	1,357,394	0.56	0.5	711,779	0.56	0.5	463,552	0.84	0.37
Ind. is in birth place	1,357,394	0.7	0.46	709,219	0.61	0.49	463,552	1	0

Notes: Table shows actual number of observations, weighted mean and standard deviation of main variables for three sample. Sample A consists of all individuals in the data. Sample B restricts to individuals born between 1940 and 1986. Finally sample C restricts to individuals living in their birth places during the Census (in 2006). Sample C is the main sample for subsequent analysis.

i.e. individuals born in the last year of the war. Unfortunately, I will not be able to observe high school outcomes for this cohort in 2006. Therefore, the youngest cohort included in the treatment is the 1986 birth cohort.

I distinguish between two types of war exposure: early childhood exposure and school time exposure. The large literature on importance of early childhood events suggests there might be a larger impact on very young children. Therefore, I split the treated cohorts in to early childhood exposure (1980-1986 birth cohorts) and school time exposure (1963-1979 birth cohorts) leaving 1940-1962 birth cohorts as the control group[20]. Equation (3.1) shows a regression specification that implements the DD methodology with two treatment groups.

$$\begin{aligned} y_{ics} &= \alpha + \beta \text{War\_prov}_s + \delta_E \text{Early}_c + \delta_S \text{School}_c \\ &\quad + \gamma_E \text{War\_prov}_s \times \text{Early}_c + \gamma_S \text{War\_prov}_s \times \text{School}_c + \epsilon_{ics} \end{aligned} \tag{3.1}$$

where y_{ics} is either a dummy that shows whether individual i in birth cohort c living in province s has finished high school, War_prov_s is equal to 1 for the five war hit provinces, Early_c is equal to 1 for 1980-1986 birth cohorts, School_c is equal to 1 for 1963-1979 birth cohorts, and α is a constant. Coefficients of interest are γ_E and γ_S which respectively show the war impact on cohorts exposed to war during early childhood and during school. I cluster standard errors at province level to allow for correlated shocks for all cohorts within a given province[21]. I also estimate an extended specification where I control for province and cohort fixed effects as follows

$$y_{ics} = \alpha + \beta_s + \delta_c + \gamma_E \text{War\_prov}_s \times \text{Early}_c + \gamma_S \text{War\_prov}_s \times \text{School}_c + \Psi X_{ics} + \epsilon_{ics} \tag{3.2}$$

where β_s is a set of province fixed effects, δ_c is a set of cohort fixed effects, and X_{ics} is a set of individual or province level controls[22]. The identification assumption for

[20] I take a non-overlapping definition here despite the fact that 1980 and 1981 cohorts started their primary during the war and hence are in some sense twice treated. Similarly, cohorts born between 1975 and 1979 are aged 5 or less when the war breaks out and therefore, might be considered in the early childhood group too. Later, I use continuous overlapping measures of treatment by looking at the number of years under age of 6 spent during the war and number of school years spent during the war.

[21] Since I have 30 provinces, the small sample bias of standard errors might be significant. However, I experimented with district level clustering (336 clusters) and the standard errors were smaller for coefficients of interest. Therefore, I take the conservative approach of clustering at the highest level.

[22] The set of controls included are indicators for urban households, the gender of the individual, family size, and their interactions with War_prov dummy. It is not possible to include parents educational attainment as I do not observe that for individuals living apart from their parents. In

causal interpretation of γ_E and γ_S is that in the absence of the war the difference between high school graduation rates across war and non-war provinces would have been the same for control and treated cohorts. In other words, the DD estimation would identify the war impact from changes in the size of the war non-war educational gap for younger cohorts. Therefore any other factor that affects younger cohorts in war provinces differentially could pose a challenge to causal interpretation. DD is, however, robust to fix differences between provinces and country-wide cohort specific variation.

There are two types of concerns with identification here. First, the restriction of the sample to non-migrant individuals is likely to result in a downward bias in the estimation of the war impact simply because war might have induced more well-endowed households to migrate out of war provinces. The second category of concerns arises due to the conditions of Iran right after 1979 revolution. A baby boom generation (1979-1986 birth cohorts), various ethnic rebellion (e.g. Kordestan uprising), and terrorist activities in major cities are a few simultaneous events that could produce a bias in my estimates. I first discuss estimation results in the following section and in section 3.6 I try to address some of these concerns.

3.5 Results

I start by presenting average outcomes for treatment and control cohorts. Table 3.3 shows average high school graduation rates for treatment and control. In panel A I compare early childhood cohorts (1980-1986 birth cohorts) to the control cohorts (1940-1962 birth cohorts). Columns (1) and (2) show high school graduation rates respectively for war and non-war provinces. Treated cohorts have an average high school graduation rate of 43 and 54.6 percent in war and non-war provinces. Column (3) reports the difference between these numbers. Obviously not all of this difference is due to the war impact. Using the educational gap between war and non-war provinces for the control cohorts (5.5 percentage points), in the third row of column (3) I have calculated the DD estimate of the impact of the war. This suggests, high school graduation rates are 6.1 percentage points lower for treated cohorts in war provinces as a result of the war. Panel B shows cohorts exposed to war during their school time received a smaller impact compared to early childhood cohorts.

some of the robustness checks I include yearly number of schools and students in the province as additional controls.

In panel C of table 3.3 I compare two control cohorts as a placebo test. I compare 1940-62 birth cohorts to 1930-39 birth cohorts. Neither of these groups has received any impact from the war because they should have finished high school education by 1980. The DD estimate for high school graduation shows the war non-war gap has widened for younger cohorts by 2.9 percentage points (insignificant). This is a comparable number to the estimated effect for the cohorts exposed to war during their schooling and suggests probably the panel B treated cohorts are not impacted by the war. It is however, much smaller than panel A estimates of war impact.

3.5.1 Graphical evidence

Figure 3.3 plots evolution of high school graduation rates for birth cohorts in war and non-war provinces. The first, second, and third vertical lines mark 1963, 1980, and 1986 birth cohorts. In the figure I also report estimates of γ_E and γ_S from the basic specification (3.1) using aggregated data. Consistent with table 3.3, figure 3.3 shows a lower fraction of individuals finish high school in war provinces even before the war but the movements of the two series seems to be fairly parallel. For the cohorts exposed to war during their late education there does not seem to be a change in the gap between war and non-war provinces. However, the gap seems to be widening from 1972 birth cohorts (i.e. second grade in 1980). This pattern becomes clearer for 1980-86 cohorts. Reported coefficient estimates also show cohorts exposed to war during their early childhood in war provinces are on average 7 percentage points less likely to finish high school. This is more than three times the magnitude of the effect on school exposed cohorts. Interestingly the early childhood effect is significant while the school exposure effect is not.

3.5.2 Regression results

Table 3.4 shows regression results for various specifications for high school completion. Column (1) reports coefficient estimates from the basic DD estimation with no controls (equation (3.1)). Coefficient estimates are very similar to those reported in figure 3.3. Once I add gender, urban, and family size and their interactions with War_prov as controls in column (2), the coefficient estimates are slightly reduced. Column (3) shows the estimation results from the full specification (equation (3.2)) with province and cohort fixed effects. This is my preferred estimate of the impact

Table 3.3: Average rate of finishing high school

	High school graduation rate		
	Province		
	war	Non-war	difference
	(1)	(2)	(3)
Panel A: Early childhood exposure to war			
Treatment	0.430	0.546	-0.116
Born 1980-1986	(0.020)	(0.043)	(0.048)
Control	0.172	0.227	-0.055
Born 1940-1962	(0.010)	(0.047)	(0.048)
Difference	0.257	0.318	-0.061
	(0.022)	(0.014)	(0.025)
Panel B: School time exposure to war			
Treatment	0.304	0.382	-0.078
Born 1963-1979	(0.010)	(0.051)	(0.052)
Control	0.172	0.227	-0.055
Born 1940-1962	(0.010)	(0.047)	(0.048)
Difference	0.132	0.155	-0.023
	(0.013)	(0.006)	(0.015)
Panel C: Placebo experiment			
Treatment	0.172	0.227	-0.055
Born 1940-1962	(0.010)	(0.047)	(0.048)
Control	0.029	0.054	-0.026
Born 1930-1939	(0.003)	(0.020)	(0.020)
Difference	0.144	0.173	-0.029
	(0.009)	(0.027)	(0.029)

Notes: Columns (1) and (2) show average rates of finishing high school for cohorts born in war and non-war provinces. Column (3) reports the difference. The last row in each panel also reports the difference. Therefore, the third row of column (3) is the DD estimate of the war impact. Standard errors are clustered at province level (30 clusters) and reported in parenthesis below coefficients. Sample restricts to individuals living in their birth place. Note difference between numbers in this table and figure 3.3 is due to use of aggregated data in the figure and individual data in the table.

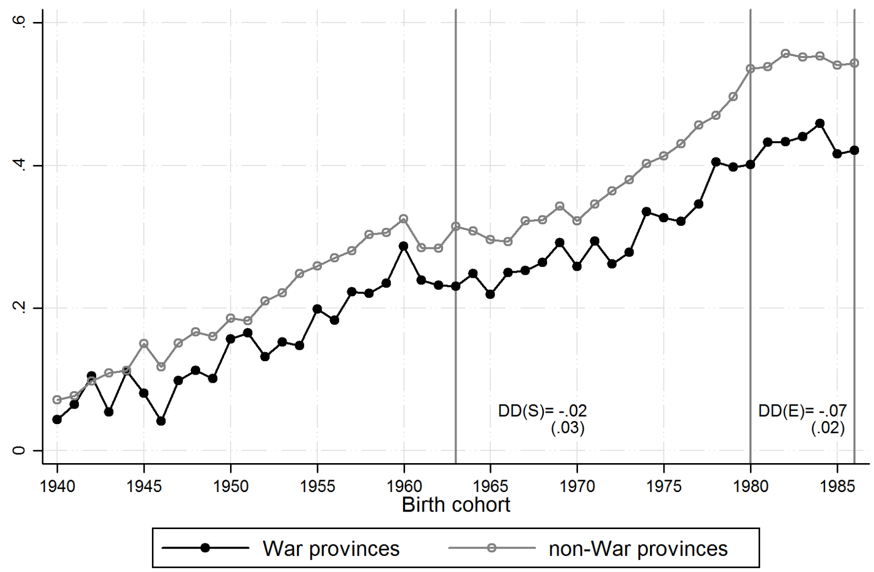

Figure 3.3: Average high school graduation rate for birth cohorts

Notes: figure shows average high school graduation rates for each cohort born in war and non-war provinces. The first, second, and third vertical lines mark 1963, 1980, and 1986 birth cohorts. reported DD(E) and DD(S) correspond to estimates of γ_E and γ_S from the basic specification (3.1) using aggregated data. Robust standard errors are reported under coefficient estimates. The sample is restricted to non-migrant individuals and averages are calculated using sampling weights.

of the war. Based on this probability of finishing high school is reduced by 4.8 percentage points for early childhood cohorts (significant at 10 percent). Given the sample mean of 39 percent high school graduation rate, this amounts to a 12 percent reduction[23]. The war impact on school age children is about one third of the early childhood effect and is insignificant.

In column (4) I estimate a much more stringent specification with household fixed effects. This specification would control for unobserved household characteristics that could impact educational outcomes. While this is an interesting specification it is subject to a potential caveat. Individuals in the control cohorts are most likely parents in the household while younger cohorts are still with their parents. Educational outcomes are shaped by household characteristics while the individual is still a child. But inclusion of household fixed effects in my sample would not control for the relevant fixed effects for parents. However, if there is high inter-generational correlation in educational attainment, inclusion of household fixed would correct for some of the factors that mattered for both parents and children education. With this potential caveat in mind, column (4) delivers the same estimate for the impact of

[23]The sample mean for 1980s cohorts is 0.528 and with this the effect is about 9 percent.

the war on early childhood cohorts but the impact on cohorts exposed to war during schooling is now vanished. The standard errors are slightly larger here and none of the coefficients are significant at 10 percent.

Column (5) uses continuous measures of exposure to war. I have calculated the number of years individuals have spent during the war while under age of 6 (early variable) and the number of years they have spent during the war while aged between 6 and 17 (school variable). Here individuals born between 1975 and 1981 have positive number of years for both of treatment measures (i.e. overlap). Coefficient estimates suggest each additional year of exposure to war while aged under 6 reduces probability of high school completion by 0.7 percentage points (significant at 10 percent) but exposure above age 6 does not seem to have a significant effect[24].

Column (6) allows for differential linear trends for high school completion of cohorts in war and non-war provinces. Here I revert to binary treatment measures used in columns (1)-(4) and instead of cohort fixed effects include a linear trend and its interaction with War_prov dummy. This specification significantly reduces the estimated war effect and both coefficients are insignificant now. While this might suggest that differential trends are responsible for the estimated war impact, I leave further discussion of this result to the end of this section.

In table 3.5 I carry out several robustness checks. Column (1) reports the benchmark estimation results from the preferred specification in column (3) of table 3.4. In columns (2) to (4) I exclude several cohorts and the estimates remain remarkably the same. In column (5) I extend the control group by including 1930-39 birth cohorts and none of the estimate change.

I can extend the regression in equation (3.2) and look at the whole set of cohort by war province interaction terms. This allows us to look at the evolution of the war non-war gap for all cohorts which could be useful in assessing the significance of differential trends. The regression equation for this is as follows:

$$y_{ijc} = \alpha + \beta_j + \delta_c + \sum_{k=1941}^{1986} (\text{War\_prov}_s \times \text{d}_{ik})\gamma_k + \Psi X_{ics} + \epsilon_{ijc} \qquad (3.3)$$

where d_{ik} is a set of cohort dummies, and γ_k captures the average difference between individuals in cohort k living in war and non-war provinces relative to the 1940

[24]For those exposed for the full 6 years while under age 6, this amounts to 4.2 percentage points reduction. For cohorts exposed to war for 8 years of their education the estimate suggests 0.8 reduction in probability of finishing high school. The former estimate is close to the binary treatment estimates but the latter is much smaller.

birth cohort (1940 is the reference group and dummies are omitted for this cohort). γ_k is expected to be zero for cohorts who finished schooling before the war and should become negative for younger cohorts. Figure 3.4 shows estimates of γs for high school graduation rates. The gray lines show 95 percent confidence intervals (calculated form province-level clustered standard errors). This figure reveals that there is a sharp decline in magnitudes of coefficients after 1972 birth cohorts. In other words, older cohorts seem to have parallel trends and the differential trend seem to appear after 1972.

To more formally test for this, I redefine treatment so that 1972-79 and 1980-86 birth cohorts form the two treatment groups while 1940-1971 birth cohorts form the control group. Regression results for this definition of treatment and control are reported in table 3.6. Columns (1) - (3) combine the two treatment groups. 1972-86 birth cohorts in war provinces are on average 0.04 percentage points less likely to finish high school (significant at 5 percent). Controlling for differential trend makes this coefficient very small and insignificant (column (3)). Perhaps more interestingly when I split the two treatment groups, it turns out the war impact on 1980-86 birth cohorts is robust and remains significant at 10 percent even after controlling for differential trends. Results in column (6) suggest 1980-86 birth cohorts are 4.2 less likely to finish high school in war provinces. 1972-79 cohorts do not seem to have received a significant war impact. This table confirms the intuition that very young cohorts have received a larger impact compared to the older ones. Furthermore, even after controlling for differential trends in column (5) the war impact remains significant at 10 percent for 1980-86 birth cohorts.

3.6 Alternative Explanations

Before interpreting the estimated impacts as causal, I would need to address several concerns. First, I deal with concerns due to the restriction of the sample to non-migrants. I present aggregate migration figures that help alleviate some of the concerns but in the end I cannot fully rule out the possibility that my estimates are driven by war induced migration of higher ability individuals. Second, I discuss potential challenges due to the Iran baby boom during 1979-1986. Here, I use provincial number of students and schools as additional controls to see whether a differential deterioration of educational resources could explain my results. The last set of confounding factors I try to rule out are post-revolution events. Following the

Table 3.4: Main regression results

	Basic	Controls	FE	HH	Continuous	Trend
	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.: Finished high school						
Early×War_prov	-0.061**	-0.050*	-0.048*	-0.047	-0.007*	-0.011
	(0.026)	(0.026)	(0.027)	(0.028)	(0.004)	(0.015)
School×War_prov	-0.023	-0.020	-0.019	-0.003	-0.001	0.005
	(0.014)	(0.013)	(0.012)	(0.017)	(0.001)	(0.010)
Observations	463,552	463,552	463,552	442,521	463,552	463,552
R-squared	0.061	0.178	0.214	0.177	0.214	0.213
Mean Dep. Var.	0.388	0.388	0.388	0.324	0.388	0.388
Controls	N	Yes	Yes	Yes	Yes	Yes
Prov. FE	N	N	Yes	N	Yes	Yes
Cohort FE	N	N	Yes	Yes	Yes	N
Household FE	N	N	N	Yes	N	N

Notes: Table shows coefficient estimates and standard errors from 6 regressions. Dependent variable is a dummy showing whether the individual has finished high school. Early and School are two indicators capturing 1980-1986 and 1963-1979 birth cohorts. I report only the two coefficient of interest, equations (3.1) and (3.2) show full specifications for column (1) and (3) respectively. In columns (2), (3), (5), and (6) I include urban, gender, and family size and their interactions with War_prov as controls. Column (4) includes gender and its interaction with War_prov as controls. Column (5) uses two continuous measures of treatment. Column (6) drops cohort fixed effects and includes dummies for school exposure, and early childhood exposure together with a linear trend and its interaction with War_prov. In all cases standard errors are adjusted for 30 province clusters. *, **, and *** respectively show significance at 10, 5, and 1 percent levels. All regressions use sampling weights but Household fixed regressions are unweighted. Sample restricts to individuals born between 1940 and 1986.

Table 3.5: Robustness regressions

	Benchmark 1940-86	Excl. 1980-86	Excl. 1963-79	Excl. 1940-49	Incl. 1930-39
	(1)	(2)	(3)	(4)	(5)
Dep. Var.: Finished high school					
Early×War_prov	-0.055*	-	-0.056*	-0.053*	-0.056*
	(0.028)	-	(0.028)	(0.030)	(0.028)
School×War_prov	-0.015	-0.013	-	-0.013	-0.017
	(0.014)	(0.014)	-	(0.016)	(0.013)
Observations	463,552	309,902	267,338	430,264	490,018
R-squared	0.132	0.116	0.164	0.115	0.149

Notes: Table shows results of 7 regressions using high school completion as the dependent variable. Different columns use different samples. Column (1) is the same as column (3) in table 3.4, here the sample is non-migrant individuals born between 1940-86. Column (2) excludes school treatment cohorts, column (3) excludes early treated cohorts. Column (4) extend the control cohorts to individuals born between 1930-39. Column (5) restricts the control cohorts to those born between 1950-62. In all specifications I have included cohort and province fixed effects in addition to controls (urban, gender, and family size and their interactions with war province dummy). All regressions use sampling weights. In all cases standard errors are adjusted for 30 province clusters. *, **, and *** respectively show significance at 10, 5, and 1 percent levels.

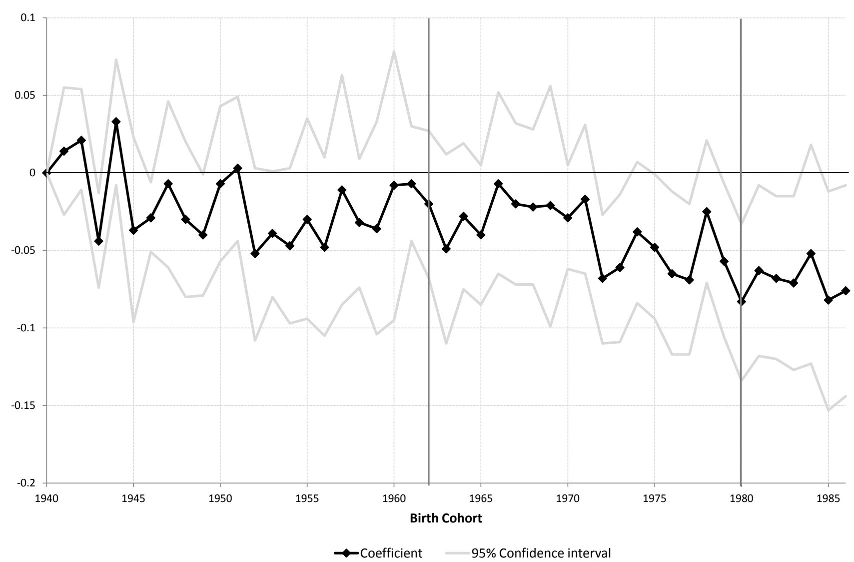

Figure 3.4: Coefficients estimates for interactions of cohort by war province

Notes: Figure plots coefficient estimates and 95 percent confidence intervals for the full set of birth cohort by War_prov interactions as in equation (3.3). Dependent variable is whether the individual has finished high school. 1940 birth cohort is set as the reference. First and second vertical lines mark 1963 and 1980 birth cohorts. Sample used for regressions is individuals born between 1940 and 1986 who are currently living in their birth place. Regressions use sampling weights and standard errors are clustered at province level.

Table 3.6: Regression results for redefined treatment groups

	FE	HH	Trend	FE	HH	Trend
	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.: Finished high school						
I[1972-86]×War_prov	-0.040**	-0.043*	-0.029			
	(0.019)	(0.022)	(0.019)			
I[1980-86]×War_prov				-0.047*	-0.053**	-0.042*
				(0.024)	(0.026)	(0.021)
I[1972-79]×War_prov				-0.030**	-0.024	-0.026
				(0.015)	(0.015)	(0.019)
Observations	463,552	442,521	463,552	463,552	442,521	463,552
R-squared	0.214	0.177	0.210	0.214	0.177	0.212
Mean Dep. Var.	0.388	0.324	0.388	0.388	0.324	0.388
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Prov. FE	Yes	N	Yes	Yes	N	Yes
Cohort FE	Yes	Yes	N	Yes	Yes	N
Household FE	N	Yes	N	N	Yes	N

Notes: Table shows coefficient estimates and standard errors from 5 regressions. Dependent variables is a dummy showing whether the individual has finished high school. I report only the two coefficient of interest. Regression specifications and controls included are similar to the similarly named columns in table 3.4. **I**[1980-86] is a dummy variable that is equal to 1 for 1980-86 birth cohorts and zero otherwise. Similarly, **I**[1972-79] is a dummy variable that is equal to 1 for 1972-79 birth cohorts. **I**[1972-86] is equal to 1 when either **I**[1980-86] or **I**[1972-79] is equal to 1. Sample restricts to individuals born between 1940 and 1986, therefore 1940-71 birth cohorts are used as control cohorts. In all cases standard errors are adjusted for 30 province clusters. *, **, and *** respectively show significance at 10, 5, and 1 percent levels. All regressions use sampling weights but Household fixed regressions are unweighted.

1979 revolution, Iran experienced great instability and a series of events overlapping with early years of war. DD estimates would be confounded by simultaneous events that have a differential impact on war provinces. Ethnic rebellions and terrorist activities are two major post-revolution phenomena that I try to rule out in the last two sub-sections.

3.6.1 Sample selection

Restricting the sample to non-migrants poses a challenge for causal inference. War forced some individuals, who would not have migrated otherwise, to permanently migrate[25]. To the extent that educational attainment of these individuals are different from those who remained in (or returned to) war areas, my treatment effect is biased. If those with better means permanently settled outside war areas, the treatment group defined here captures the set of individuals who would have had lower educational attainment even in the absence of the war causing an overestimate of the war impact[26].

Based on SCI publications, the war resulted in a peak displacement of more than 1.6 million individuals by June 1982 (table 3.7). War migrants were settled in temporary camps, nearby cities, or large cities like Tehran and Esfahan. Khuzestan was the hardest hit province both because it is larger than the other provinces and because large cities like Khorramshahr and Abadan were fully evacuated during the early stages of the war. As table 3.7 shows 76 percent of war migrants are from Khuzestan. The interesting feature of migration patterns is that majority of war migrants were settled in the same province. 49, 92, 98, and 90 percent of war migrants from Khuzestan, Ilam, Kordestan, and Kermanshah provinces were settled in the same province.

Given most of these settlements were temporary it is likely that the majority of migrants have returned to their homes after the war. On the other hand the war lasted for about 8 years and while the large cities were freed in the second year of

[25]Note the term permanent is important here as the sample of temporary migrants would have returned to their birth place and are included in my sample. There is, however, a potential bias even from temporary migration. Households migrated outside war areas during the war might have given birth to children at that time. These children are brought back to household original living place after the war but the children themselves are not living in their birth place and therefore are excluded from my sample. Nonetheless, these children are affected by the war.

[26]Forced migration itself is a mechanism for the impact of war on educational attainment. Interruption of schooling due to forced migration could result in school dropout. The bias discussed above is due to exclusion of war migrants who settled in locations other than their birth place.

Table 3.7: War migrants as of June 1982

Current residence	Residence before war				
	Khuzestan	Ilam	Kordestan	Kermanshah	Other
Khuzestan	**49.2**	1.3	0	0.1	90.2
Ilam	0.6	**91.7**	0	4.7	0
Kordestan	0	0	**98.1**	0.9	4.2
Kermanshah	0.5	0.4	0	**89.9**	4.6
Other	49.7	6.7	1.9	4.3	**0.9**
Total Number	1,253,786	42,501	5,797	144,846	189,817
% of total	76.6	2.6	0.4	8.8	11.6

Notes: The figures are calculated from a publication of SCI. Columns show fraction of war migrants from each war hit province that settled in any of the provinces listed in the rows. The last row shows total number of war migrants in each province.

the war not many residents returned to those cities right away. Living in a place for 8 years increases the chances of permanent settlement. In order to get a broad idea of migration patterns during and after the war, I use aggregate statistics from three rounds of censuses. I use 1986 census for migration numbers between 1976 and 1986, i.e. during the war period. The next round of census in 1996 would suggest how many of the war migrants have returned to their home after the war. Finally, I use 2006 census as a comparison for the two earlier rounds to get an idea of non-war migration patterns (benchmark census).

Figure 3.5a shows net in-migrants entered each province as a fraction of end year province population for the three censuses. I have ordered provinces so the first block shows five war provinces, the second block shows provinces neighboring Khuzestan, and the bottom block shows other provinces. During the war Khuzestan and Kordestan show very high de-population rates. Around 9 percent of Khuzestan's population moved out of the province between 1976 and 1986. Interestingly, neighboring provinces show high in-migration rates, with Boushehr having the highest rate in the country. It is worth noting that Fars and Esfahan are two big provinces and low in-migration rates reflect their large populations whereas Boushehr is fairly small. All neighboring provinces accommodated large numbers of migrants from Khuzestan. These patterns are reversed in the after war census. Khuzestan now shows high in-migration rate while neighboring provinces show de-population with Boushehr having highest de-population rate in the country.

The rightmost panel in the figure 3.5a shows net in-migration for the benchmark census. Generally speaking Tehran, Yazd, and Boushehr show relatively high rates of in-migration while Kermanshah, Kordestan, Hamedan, and E. Azerbaijan have high de-population rates. This benchmark show modest degree of mobility in normal times (between -5 and 5 percent), and corroborates the abnormality of high war time de-population of Khuzestan and its partial re-population after the war.

As table 3.7 showed majority of war migrants were settled in the same province. The earlier figure captured only inter-province migration. Figure 3.5b shows intra-province migration rates for the three rounds of censuses. During the war and after the war, war hit provinces have highest intra-province migration rates which shows higher than average reshuffling within these provinces. In the benchmark census 18 years after the end of the war, intra-province migration rates for war provinces are still quite high but comparable to other provinces. For example, Khuzestan shows an intra-province migration of about 15 and 14 percents in 1986 and 1996 censuses but in 2006 this falls to less than 10 percent. These patterns are consistent with the idea that individuals were displaced at the time of the war but returned to their homes afterward[27].

High intra-province migration rates suggest that while individuals might not be living in their birth place, they might still be in the same province. Since my treatment measure is defined at province level, in table 3.9 column (2) I run the preferred specification on the full sample of individuals born between 1940 and 1986 assuming that anyone who lives in a war (non-war) province in 2006 is in the treatment (control) group. The results show little change for the early childhood impact but the coefficient estimate for school impact is reduced significantly. The fact that the war impact remains stable for the sample of all individuals supports plausibility of the sample restriction. In column (3), I exclude Khuzestan to see if results are driven by the high migration rates in this province. The coefficients are now insignificant but magnitudes remain the same to the benchmark sample.

Another way I can check the plausibility of the sample restriction is to see whether the probability of being included in the sample is affected by the treatment. Figure 3.6a plots average fraction of non-migrant individuals for each cohort in war and non-war provinces. About 55 percent of older cohorts and about 75 percent of youngest cohort in the sample currently reside in their birth place. Interestingly, war and non-war provinces have fairly similar fraction of non-migrant individuals. Running a regression confirms that the same fraction of treated cohorts are non-migrant in

[27] Overall large internal migration rates in Iran are due to high migration rates from rural areas to urban centers.

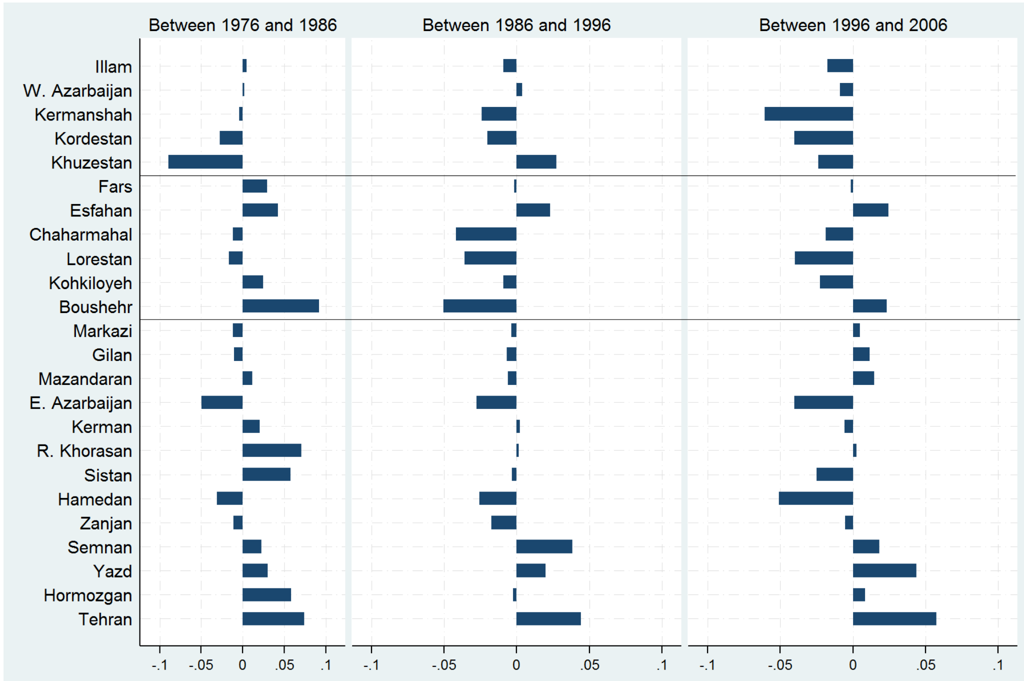

(a) Net inter-province in-migrants (% of province population)

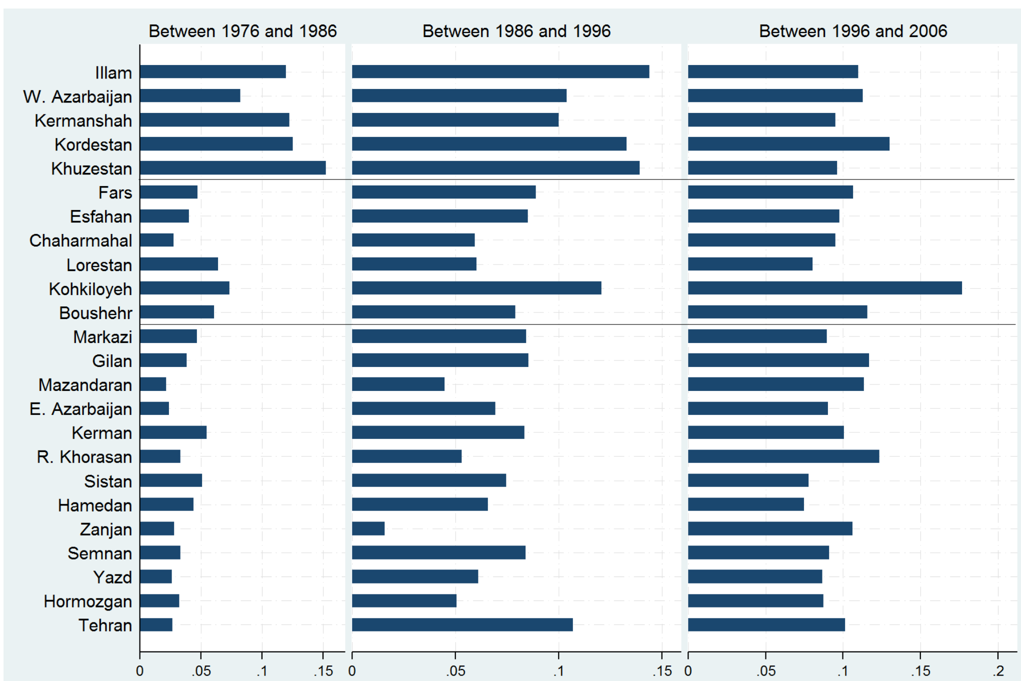

(b) Intra-province migrants (% of province population)

Figure 3.5: Net in-migration into provinces during and after war period

Notes: figures (a) and (b) respectively show inter and intra-province migration rates for three rounds of censuses for 24 provinces. I have defined provinces in a consistent way to make results comparable across rounds of censuses and merged 6 newly formed provinces with their original province. None of the war provinces had split in to further provinces over time. Migration rates are calculated by dividing the relevant migration numbers by total province population at the end date. Source of data is from SCI census publications.

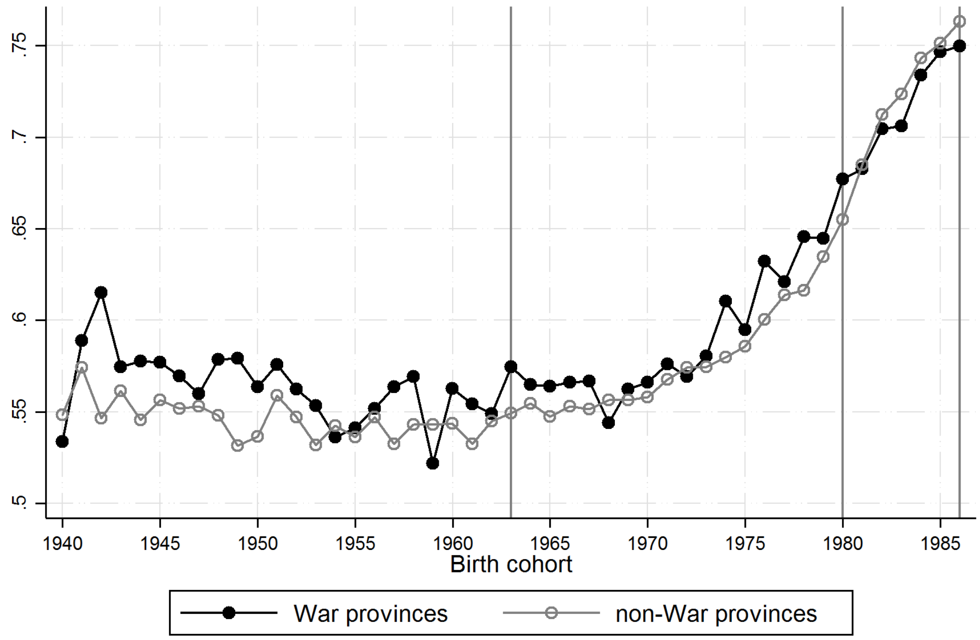

(a) Fraction of non-migrant individuals in each cohort

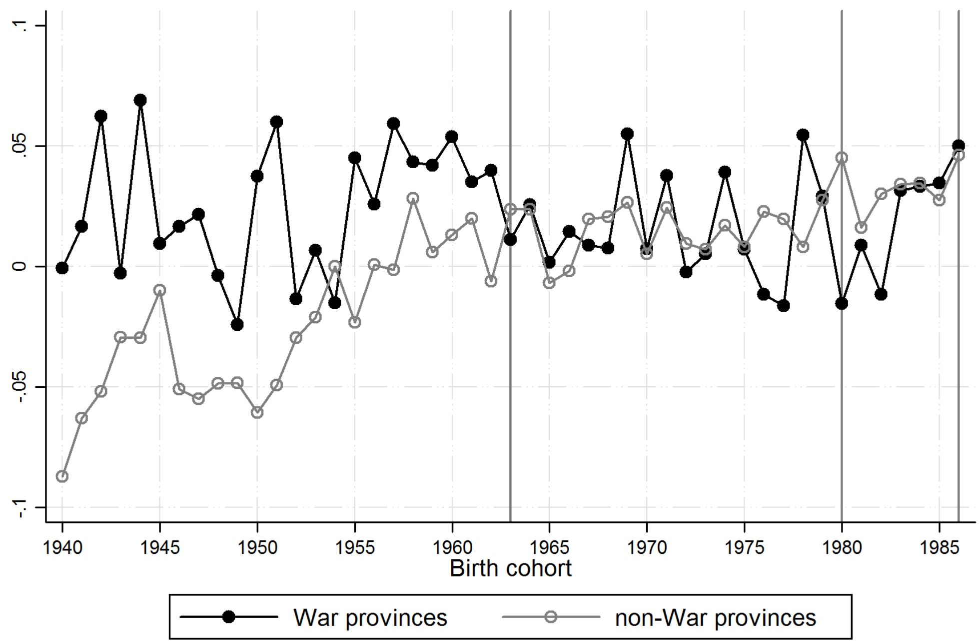

(b) High school completion gap between non-migrants and migrants

Figure 3.6: Impact of non-migrants restriction on war and non-war provinces

Notes: Panel (a) shows fraction of individuals who are living in their birth places in 2006 Census for various birth cohorts in war and non-war provinces. Panel (b) plots the difference between average high school completion rates for non-migrant and migrant individuals. Sample used here is the full sample of individuals born between 1940-1986 and I use sampling weights in calculation of averages.

Table 3.8: Regression results for probability of living in birth place

	Basic	Controls	FE	HH
	(1)	(2)	(3)	(4)
Dep. Var. Individual living in birth place				
Early×War_prov	-0.020	-0.020	-0.019	-0.013
	(0.030)	(0.029)	(0.029)	(0.032)
School×War_prov	-0.003	-0.003	-0.003	-0.001
	(0.019)	(0.018)	(0.017)	(0.020)
Controls	N	Yes	Yes	Yes
Prov. FE	N	N	Yes	N
Cohort FE	N	N	Yes	Yes
Household FE	N	N	N	Yes
Observations	709,219	709,219	709,219	675,240
R-squared	0.021	0.054	0.088	0.072
Mean Dep. Var.	0.612	0.612	0.612	0.614

Notes: Table reports results from regressions of a dummy variable that shows whether the individual is in his/her birth place on covariates. The specifications under each column correspond to specifications estimated in table 3.4. Standard errors are clustered at province level (30 clusters). Clustering at district level (336 districts) does not change significance of any of the coefficients.

war and non-war provinces. Using the same specifications as in equations (3.1) and (3.2) but using an indicator for being in birth place as the dependent variable, table 3.8 shows the interaction terms are not significant in any of the specifications. In terms of magnitude the estimates suggest cohorts born between 1980 and 1986 are 2 percentage points less likely to be in their birth place if they are currently residing in war provinces. Given the mean of the dependent variable this is a 3 percent reduction.

Even if a balanced fraction of individuals are included in the sample across war and non-war provinces, the included sample might be different on characteristics that matter for educational attainment. In figure 3.6b I take a step further and look at the educational gap between non-migrants and migrants across war and non-war provinces to see if in terms of the outcome variable the included individuals are different from excluded ones. In war provinces non-migrants seem to be on average more educated than migrants across most of the cohorts. However for non-war provinces it seems older non-migrants have attained lower levels of education compared to same age migrants, whereas younger non-migrants seem to outperform migrants. Overall it seems the educational gap is broadly similar across war and non-war provinces.

While the above mentioned arguments go some way to relieve concerns, they cannot fully rule out the bias induced from the sample restriction. It is hard to assess this in the absence of micro data on migration patterns during the war[28].

3.6.2 Baby boom

Between 1976 and 1986 Iran had a baby boom with an average yearly population growth rate of 3.9 percent. The population growth rates in the preceding and proceeding decades are respectively 2.7 and 2.5 percent. Figure 3.7a shows number of registered births in the country sharply rises after 1979 from 1.5 million to 2.5 million. The birth rates remain fairly high during 1980-1986 but start to fall after 1986[29]. In this section, I provide two pieces of evidence that suggest the higher war impact on 1980-1986 birth cohorts is not due to the differential impact of the baby boom in war provinces. First figure 3.7b plots average annual registered births for war and non-war provinces for ten birth cohorts around the baby boom period. While

[28]Unfortunately, I did not have access to micro data for 1986 and 1996 Censuses.

[29]The rise and fall in birth rates were mostly due to government campaigns to first increase fertility after the 1979 revolution and then decrease fertility during the second half of the war.

non-war provinces have on average higher number of births, the difference between the two regions is fairly stable.

Even though birth figures move in parallel across provinces, the educational impact of the population expansion could be heterogeneous. For example, war provinces might have built fewer schools to accommodate the baby boom and hence could have overcrowded classes leading to worse educational outcomes. In order to address this concern, I have collected yearly province-level data on number of students, schools and classes from various Iran Statistical Year Books. In table 3.9, column (4) I included log number of schools and students as additional controls in the main specification. While these variables have the right sign both of them are insignificant. Furthermore, their inclusion does not affect the estimated war impact on early childhood and school cohorts. If anything, the estimated effect seems to be larger now.

3.6.3 Ethnic rebellions

The third event that could potentially bias DD estimates is the rebellion movements in West Azerbaijan, Kordestan, and Khuzestan. In the turbulent aftermath of the revolution these rebellions started out mostly as ethnic movements for independence. In fact Iraq expected help from the Arab rebels in Khuzestan but the invasion marked a unification of Arabs and Persians. The most powerful and long lasting rebellion was the Kordestan uprising. In table 3.9, column (5), I have excluded Kordestan to check the robustness of results. The estimated war effects are slightly smaller and the early childhood effect is now insignificant but overall the results are in the same ball park. Furthermore, the rebellions were almost finished in the second half of the war (1984-88) but figures 3.3 and 3.4 suggest a larger impact of the war for the youngest cohorts in my sample (1984-1986).

3.6.4 Other confounding events

Apart from the simultaneous baby boom and ethnic rebellions, two other events warrant some discussion. Right after the 1979 revolution, some factions of the revolutionary groups started to oppose the policies undertaken by the mainstream forces. Soon the opposition moved underground and embarked on assassinations and terrorist bombings in a few major cities between 1979 and 1982. Several observations

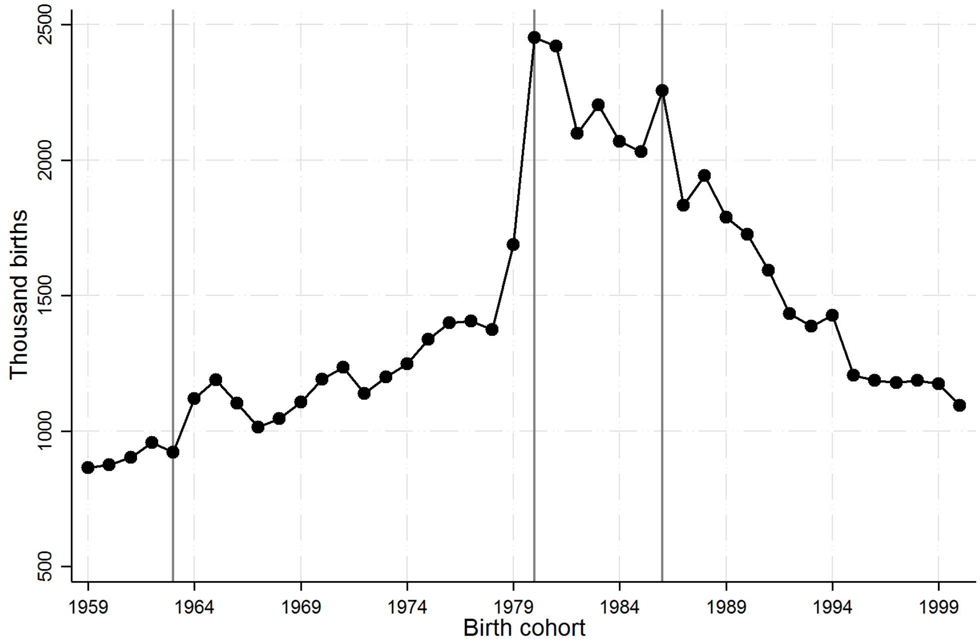

(a) Country-wide registered births

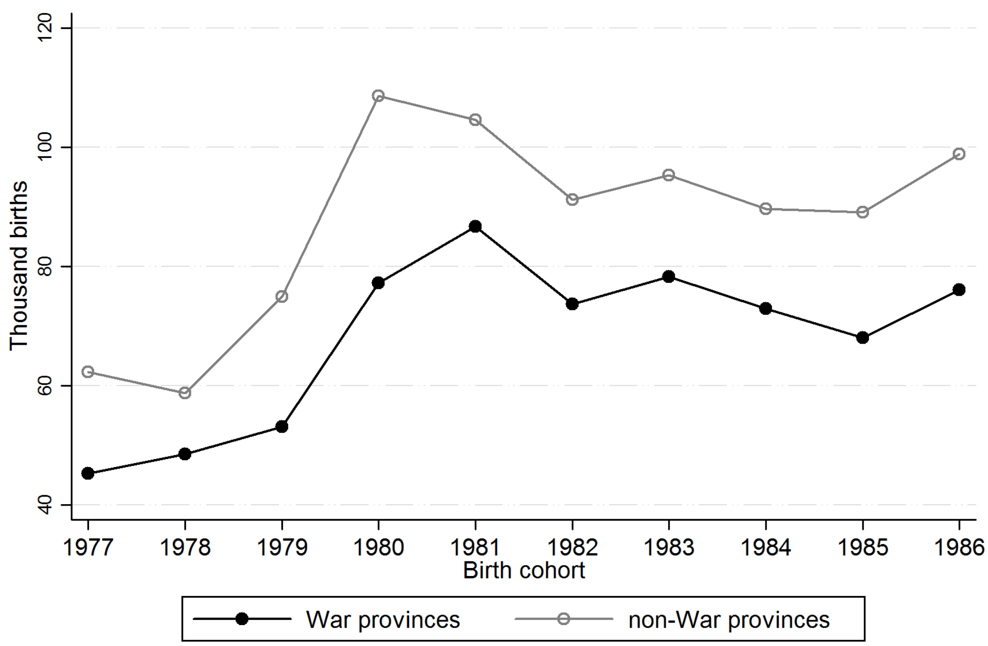

(b) Average provincial registered births

Figure 3.7: Number of registered births over time

Notes: Figure (a) shows number of registered birth in calendar years. First, second, and third vertical lines mark 1963 , 1980, and 1986 birth cohorts. Source of this data is from National Organization for Civil Registrations. Figure (b) plots average number of births in war and non-war provinces. Source of this data is various Statistical Year books from SCI. Registered births are different from actual births during a calendar year because some birth events were registered with delay.

Table 3.9: Regressions for ruling out alternative stories

	Benchmark	All	Excl. Khuzestan	Educ. inputs	Excl. Kordestan	Excl. prov. capitals
	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.: Finished high school						
Early×War_prov.	-0.048*	-0.038	-0.053	-0.058**	-0.040	-0.054*
	(0.027)	(0.023)	(0.036)	(0.025)	(0.028)	(0.032)
School×War_prov.	-0.019	-0.007	-0.017	-0.030*	-0.017	-0.014
	(0.012)	(0.011)	(0.019)	(0.016)	(0.014)	(0.015)
Ln(schools)				0.027		
				(0.028)		
Ln(students)				-0.011		
				(0.047)		
Observations	463,552	711,779	437,350	350,047	454,591	361,773
R-squared	0.214	0.168	0.217	0.182	0.213	0.187

Notes: Table shows several robustness checks for ruling out competing stories. Column (1) replicates estimation results from the preferred specification in table 3.4 for comparison purposes. Column (2) estimates the same specification as column (1) but includes all (i.e. migrant and non-migrant) individuals in the analysis. Here anyone who lives in a war (non-war) province in 2006 is assumed to assigned to treatment (control). The rest of the table focus on non-migrant individuals as in column (1). Column (3) and (5) exclude individuals from respectively Khuzestan and Kordestan provinces. Column (4) includes log number of primary schools and log number of primary students in each province-year as additional controls in the regression. Due to data availability, the sample for this column runs from 1960-1986 and the number of clusters is 22 provinces. Column (6) excludes 30 districts that contain the provincial capitals. Standard errors are clustered at province level (30 provinces). *, **, and *** respectively show significance at 10, 5, and 1 percent levels.

make it less likely that the terrorist activities are responsible for the estimated effects. First, most of terrorist activities took place in major cities (often Tehran). However, when I exclude all provincial capitals the estimated war impact on early childhood cohorts becomes slightly larger (table 3.9, column (6)). Second observation that alleviates concerns is the fact that the treatment effect seems to be stronger for younger cohorts (figures 3.3 and 3.4). This is despite the fact that little terrorist activities happened after 1982.

The other event that requires some explanation is the Cultural Revolution which closed all universities between 1980 and 1982. The stated objective was to bring the tutoring in line with Islamic thought. This event could reduce incentives for finishing high school as the prospect of entering university was unclear. However, it is not entirely obvious that the Cultural Revolution had a heterogeneous impact on war provinces. Furthermore, the strongest impact of the war is on cohorts who started primary or are born during the war. These cohorts are quite far from university education and the universities were expected to open soon.

3.7 Conclusions

In this paper I estimated the reduced form impact of IIW on educational attainment of children. DD estimates suggest probability of finishing high school is reduced by 4.8 percentage points for cohorts born during the war in war provinces, whereas cohorts that spent some years of their schooling during the war saw a reduction of 1.9 percentage points. These estimates suggests a stronger impact for younger cohorts. It seems 1980-86 birth cohorts (born during the war) have received a robust negative war impact which remains significant at 10 percent even after controlling for differential trends. On the other hand, 1972-79 birth cohorts (aged between 8-1 years old when war started) seem to have received a much smaller and insignificant impact. Older cohorts are unaffected.

The main issue with interpreting these estimates as causal is the sample restriction. I have focused on non-migrants in order to identify birth place of individuals. War, however, might influence migration patterns and result in biased DD estimates. Aggregate migration statistics from three rounds of censuses, however, support the idea that at least part of the war migrants returned to their homes. Furthermore, I have shown that a balanced number of individuals are removed from cohorts across war and non-war provinces due to this restriction. But in the end it is hard to address all concerns in the absence of war time micro data on migration patterns. I have also tried to rule out other simultaneous events as potential confounding factors. Most importantly, I have ruled out a baby boom as a candidate for explaining the estimated effects by including number of schools and students in each province directly in the regression.

The results of my analysis show very young and unborn children are more susceptible to adverse shocks. It seems school age children have managed to maintain their education levels in war provinces but children born during the war seem to have suffered a sustained negative impact. It is beyond the scope of this paper to suggest potential remedies for compensation of these effects but spending more resources for education of affected cohorts seems like a reasonable idea. It is not unreasonable to think that some of these results are applicable to contexts beyond IIW. For example, natural disasters (floods and earthquakes) could have negative effects on pregnant women and their prospective babies and could warrant government intervention.

I am planning to extend the analysis in this paper by collecting detailed data on exact timing and location of missile and aerial attacks on cities outside war provinces. Combining this data with birth date and location of individuals would allow a careful investigation of the early childhood and in-utero effects during aerial attacks. This data would deliver cleaner identification of the causal effect because of its precise timing and location. Furthermore, so far I have delivered mostly reduced form estimates of the war effect. It is equally, if not more, important to know the mechanisms that led to these effects. Here I am investigating the use of several rounds of Household Expenditure Surveys during the second half of the war to look at changes in average incomes, and school enrollment rates to shed more light on potential mechanisms.

References

Agarwal, S., C. Liu, and N. S. Souleles (2007). The reaction of consumer spending and debt to tax rebates: Evidence from consumer credit data. *Journal of Political Economy 115*(6), 986–1019.

Ahmadi, K., M. Reshadatjo, N. Sepehrvand, P. Ahmadi, and H. Yaribeygi (2010). Evaluation of vicarious PTSD among children of Sardasht chemical warfare survivors 20 years after Iran-Iraq war. *Journal of Applied Sciences 10*(23), 3111–6.

Ahmadi, K. B., M. Reshadatjou, G. R. Karami, and J. Anisi (2009). Vicarious PTSD In Sardasht Chemical Warfare Victims' wives. *Journal of Behavioral Sciences 3*(3), 195–199.

Akbulut-Yuksel, M. (2010). Children of War: The Long-Run Effects of Large-Scale Physical Destruction and Warfare on Children. *Working Paper*.

Akresh, R. and D. d. Walque (2008). Armed Conflict and Schooling: Evidence from the 1994 Rwandan Genocide. *Policy Research Working Paper 4606, The World Bank*.

Almond, D. (2006). Is the 1918 Influenza Pandemic Over? Long-Term Effects of In Utero Influenza Exposure in the Post-1940 U.S. Population. *Journal of Political Economy 114*(4), 672–712.

Almond, D. and J. Currie (2011a). Chapter 15: Human Capital Development before Age Five. In A. Orley and C. David (Eds.), *Handbook of Labor Economics*, Volume Volume 4, Part B, pp. 1315–1486. Elsevier.

Almond, D. and J. Currie (2011b). Killing Me Softly: The Fetal Origins Hypothesis. *Journal of Economic Perspectives 25*(3), 153–172.

Almond, D., L. Edlund, H. Li, and J. Zhang (2010). Long-Term Effects of Early-Life Development: Evidence from the 1959 to 1961 China Famine. In T. Ito and

A. K. Rose (Eds.), *The Economic Consequences of Demographic Change in East Asia*, pp. 321–345. NBER–East Asia Seminar on Economics, vol. 19. Chicago and London: University of Chicago Press.

Almond, D., L. Edlund, and M. Palme (2009). Chernobyl's Subclinical Legacy: Prenatal Exposure to Radioactive Fallout and School Outcomes in Sweden. *Quarterly Journal of Economics 124*(4), 1729–1772.

Almond, D. and B. Mazumder (2005). The 1918 Influenza Pandemic and Subsequent Health Outcomes: An Analysis of SIPP Data. *American Economic Review 95*(2), 258–262.

Almond, D. and B. Mazumder (2011). Health Capital and the Prenatal Environment: The Effect of Ramadan Observance during Pregnancy. *American Economic Journal: Applied Economics 3*(4), 56–85.

Almond, D., B. Mazumder, and R. van Ewijk (forthcoming). Fasting during pregnancy and children's academic performance. *Economic Journal*.

Arasteh, A. R. (1962). *Education and Social Awakening in Iran*. Brill Archive.

Auerbach, A. J., W. G. Gale, and B. H. Harris (2010). Activist fiscal policy. *Journal of Economic Perspectives 24*(4), 141–164.

Bakhash, S. (2004). Chapter 1: The Troubled Relationship: Iran and Iraq, 1930-80. In L. G. Potter and G. G. Sick (Eds.), *Iran, Iraq, and the Legacies of War*. Palgrave Macmillan.

Barrell, R. and M. Weale (2009). The economics of a reduction in vat. *Fiscal Studies 30*(1), 17–30.

Bhargava, S. and D. Manoli (2013). Why are Benefits Left on the Table? Assessing the Role of Information, Complexity, and Stigma on Take-up with an IRS Field Experiment.

Blattman, C. and J. Annan (2010). The Consequences of Child Soldiering. *Review of Economics and Statistics 92*(4), 882–898.

Blundell, R. (2009). Assessing the temporary vat cut policy in the uk. *Fiscal Studies 30*(1), 31–38.

Bosker, M., S. Brakman, H. Garretsen, and M. Schramm (2007). Looking for multiple equilibria when geography matters: German city growth and the WWII shock. *Journal of Urban Economics 61*(1), 152–169.

Carbonnier, C. (2007). Who pays sales taxes? Evidence from French VAT reforms, 1987–1999. *Journal of Public Economics 91*(5–6), 1219–1229.

Chamarbagwala, R. and H. E. Moran (2010). The Legacy of War: Post-War Schooling Inequality in Guatemala.

Chamarbagwala, R. and H. E. Moran (2011). The human capital consequences of civil war: Evidence from Guatemala. *Journal of Development Economics 94*(1), 41–61.

Chetty, R., J. N. Friedman, T. Olsen, and L. Pistaferri (2011). Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from danish tax records. *The Quarterly Journal of Economics 126*(2), 749–804.

Chetty, R., J. N. Friedman, and E. Saez (2013). Using differences in knowledge across neighborhoods to uncover the impacts of the EITC on earnings. *American Economic Review 103*(7), 2683–2721.

Chetty, R., A. Looney, and K. Kroft (2009). Salience and Taxation: Theory and Evidence. *The American Economic Review 99*(4), 1145–1177.

Chetty, R. and E. Saez (2013). Teaching the tax code: Earnings responses to an experiment with EITC recipients. *American Economic Journal: Applied Economics 5*(1), 1–31.

Chirakijja, J., T. F. Crossley, M. Luhrmann, and C. O'Dea (2009). The stimulus effect of the 2008 uk temporary vat cut. *102nd Annual Conference on Taxation, National Tax Association*.

Choi, J. J., D. Laibson, and B. C. Madrian (2011). $100 bills on the sidewalk: Suboptimal investment in 401 (k) plans. *Review of Economics and Statistics 93*(3), 748–763.

Cordesman, A. H. (1987, October). *The Iran-Iraq war and Western security 1984-87: strategic implications and policy options*. Jane's.

Cordesman, A. H. and A. Wagner (1990, April). *The Lessons Of Modern War, Vol. 2: The Iran-Iraq War* (1ST edition ed.). Boulder, Colo. : London: Westview Press.

Crossley, T. F., H. Low, and C. O'Dea (2013). Household consumption through recent recessions. *Fiscal Studies 34*(2), 203–229.

Crossley, T. F., H. Low, and M. Wakefield (2009). The economics of a temporary vat cut. *Fiscal Studies 30*(1), 3–16.

Davis, D. R. and D. E. Weinstein (2002). Bones, Bombs, and Break Points: The Geography of Economic Activity. *American Economic Review*.

Davis, D. R. and D. E. Weinstein (2008). A Search For Multiple Equilibria In Urban Industrial Structure. *Journal of Regional Science 48*(1), 29–65.

Devereux, M., L. Liu, and S. Loretz (2014). The Elasticity of Corporate Taxable Income: New Evidence from UK Tax Records. *American Economic Journal: Economic Policy 6*(2), 19–53.

Fernandez-de Cordoba, G. and J. L. Torres (2011). The transitory vat cut in the uk: A dynamic general equilibrium analysis. *Economic Issues 16*(1), 1–18.

GfK Business (2008). The extent and nature of the use of computerized accounting by businesses to meet their VAT and corporation tax obligations. *HM Revenue and Customs*.

Harju, J. and T. Kosonen (2013). Restaurant vat cut: Cheaper meal and more service? *Government Institute for Economic Research (VATT), Working Paper 52*.

Hiro, D. (1989). *The Longest War: The Iran-Iraq Military Conflict*. Psychology Press.

HM Customs and Excise (2002). Easing the impact of VAT: Consultation on a Flat Rate Scheme for small firms: Responses to the consultation document issued in June 2001. *HM Treasury and HM Customs and Excise, Parliament Deposited Paper Dep. 02/956*.

HMRC (2006). Survey to establish awareness and take-up of VAT Cash Accounting Scheme (CAS). *HMRC research report No. 27*.

HMRC (2009). VAT Flat Rate Scheme (FRS): Impact assessment of changes to the flat rate percentages in january 2010.

Ichino, A. and R. Winter-Ebmer (2004). The Long-Run Educational Cost of World War II. *Journal of Labor Economics 22*(1), 57–87.

Jappelli, T. and L. Pistaferri (2010). The consumption response to income changes. *Annu. Rev. Econ 2*, 479–506.

Johnson, D. S., J. A. Parker, and N. S. Souleles (2006). Household expenditure and the income tax rebates of 2001. *The American Economic Review 96*(5), 1589–1610.

Jones, D. (2010). Information, preferences, and public benefit participation: Experimental evidence from the advance EITC and 401(k) savings. *American Economic Journal: Applied Economics 2*(2), 147–163.

Jones, D. (2012). Inertia and overwithholding: Explaining the prevalence of income tax refunds. *American Economic Journal: Economic Policy 4*(1), 158–185.

Kadivar, H. and S. C. Adams (1991). Treatment of chemical and biological warfare injuries: insights derived from the 1984 Iraqi attack on Majnoon Island. *Military medicine 156*(4), 171–177.

Karsh, E. (2002). *The Iran-Iraq War, 1980-1988.* Osprey Publishing.

Khateri, S., M. Ghanei, S. Keshavarz, M. Soroush, and D. Haines (2003). Incidence of lung, eye, and skin lesions as late complications in 34,000 Iranians with wartime exposure to mustard agent. *Journal of occupational and environmental medicine 45*(11), 1136–1143.

Kleven, H. J. and M. Waseem (2013). Using notches to uncover optimization frictions and structural elasticities: theory and evidence from Pakistan. *The Quarterly Journal of Economics 128*, 669–723.

KPMG, L. (2006). Administrative burdens - HMRC measurement project report by tax area part 27: Value added tax.

Liebman, J. B. and E. F. Luttmer (2011). Would people behave differently if they better understood social security? evidence from a field experiment. *National Bureau of Economic Research Working Paper 17287.*

Mahvash, J. (2011). Study of divorce in Iran provinces from 1977 to 1998: Emphasis on the role of Iran-Iraq war. *International Journal of Sociology and Anthropology 3*(4), 132–138.

Menashri, D. (1992). *Education and the Making of Modern Iran.* Cornell University Press.

Merrouche, O. (2006). The Human Capital Cost of Landmine Contamination in Cambodia. *HiCN Working Paper 25*.

Mian, A. and A. Sufi (2012). The effects of fiscal stimulus: Evidence from the 2009 cash for clunkers program*. *The Quarterly Journal of Economics 127*(3), 1107–1142.

Miguel, E. and G. Roland (2011). The long-run impact of bombing Vietnam. *Journal of Development Economics 96*(1), 1–15.

Misch, F. and A. Seymen (2013). The effects of countercyclical fiscal policy: Firm level evidence from temporary consumption tax cuts in turkey.

Mofid, K. (1990). *The economic consequences of the Gulf War.* Routledge, Taylor & Francis Group.

National Audit Office (2004). HM Customs and Excise: Tackling VAT Fraud. *London: The Stationery Office*.

National Audit Office (2006). HM Revenue and Customs: Filing VAT and company tax returns. *London: The Stationery Office*.

National Audit Office (2010). HM Revenue and Customs: Engaging with tax agents. *London: The Stationery Office*.

ORC International (2010). HMRC compliance costs and commercial impact of december 2008 vat rate change. *HMRC research report No. 103*.

Parker, J. A., N. S. Souleles, D. S. Johnson, and R. McClelland (2013). Consumer spending and the economic stimulus payments of 2008. *The American Economic Review 103*(6), 2530–2553.

Pike, R., M. Lewis, and D. Turner (2009). Impact of vat reduction on the consumer price indices. *Data and support*, 17.

Potter, L. G. and G. G. Sick (2004, November). *Iran, Iraq, and the Legacies of War*. Palgrave Macmillan.

Saez, E. (2009). Details matter: The impact of presentation and information on the take-up of financial incentives for retirement saving. *American Economic Journal: Economic Policy 1*(1), 204–228.

Sahm, C. R., M. D. Shapiro, and J. Slemrod (2012). Check in the mail or more in the paycheck: Does the effectiveness of fiscal stimulus depend on how it is delivered? *American Economic Journal: Economic Policy 4*(3), 216.

Seely, A. (2009). Vat : the temporary cut in the standard rate. *House of Commons-Standard Note:SN/BT/701*.

Shemyakina, O. (2011). The effect of armed conflict on accumulation of schooling: Results from Tajikistan. *Journal of Development Economics 95*(2), 186–200.

Souresrafil, B. (1989, July). *The Iran-Iraq war*. Guinan Lithographic Co.

Appendix A

Flat rates for FRS categories

Category of Business	24 Apr 02 - 31 Dec 03	1 Jan 04 - 30 Nov 08	1 Dec 08 - 31 Dec 09	1 Jan 10 - 3 Jan 11	4 Jan 11 Onwards
Post offices**	6	2	2	4.5	5
Retailing food, confectionery, tobacco, newspapers or children's clothing	5	2	2	3.5	4
Wholesaling food	7	5.5	5	6.5	7.5
Membership organisation	7	5.5	5.5	7	8
Pubs	6	5.5	5.5	6	6.5
Farming or agriculture that is not listed elsewhere	6.5	6	5.5	6	6.5
Retailing that is not listed elsewhere	7	6	5.5	6.5	7.5
Wholesaling agricultural products	7	6	5.5	7	8
Retailing pharmaceuticals, medical goods, cosmetics or toiletries	8	7	6	7	8
Retailing vehicles or fuel	8	7	5.5	6	6.5
Sport or recreation	8	7	6	7.5	8.5
Wholesaling that is not listed elsewhere	8	7	6	7.5	8.5
Printing	8.5	7.5	6.5	7.5	8.5
Repairing vehicles	8.5	7.5	6.5	7.5	8.5
Agricultural services	9	7.5	7	10	11
Library, archive, museum or other cultural activity	8.5	7.5	7.5	8.5	9.5
Manufacturing food	8.5	7.5	7	8	9
General building or construction services*	9	8.5	7.5	8.5	9.5
Manufacturing yarn, textiles or clothing	9.5	8.5	7.5	8	9
Manufacturing that is not listed elsewhere	10	8.5	7.5	8.5	9.5
Packaging	9	8.5	7.5	8	9
Repairing personal or household goods	10	8.5	7.5	9	10
Hiring or renting goods	9.5	8.5	7.5	8.5	9.5
Social work	9	8.5	8	10	11
Forestry or fishing	10	9	8	9.5	10.5
Mining or quarrying	10	9	8	9	10
Courier Services**	6	9	8	9	10
Transport or storage, including couriers, freight, removals and taxis**	10	9	8	9	10
Travel agency	10	9	8	9.5	10.5
Advertising	11	9.5	8.5	10	11
Hotel or accommodation	10.5	9.5	8.5	9.5	10.5
Photography	10	9.5	8.5	10	11
Publishing	10	9.5	8.5	10	11
Veterinary medicine	11	9.5	8	10	11
Dealing in waste or scrap	11	9.5	8.5	9.5	10.5
Any other activity not listed elsewhere	11	10	9	10.5	12
Investigation or security	11	10	9	10.5	12
Manufacturing fabricated metal products	11	10	8.5	9.5	10.5
Boarding or care of animals	11	10.5	9.5	10.5	12
Film. radio, television or video production	-	10.5	9.5	11.5	13
Business services that are not listed elsewhere	12.5	11	9.5	10.5	12
Entertainment or journalism	12	11	9.5	11	12.5
Estate agency or property management services	11.5	11	9.5	10.5	12
Laundry or dry-cleaning services	12	11	9.5	10.5	12
Secretarial services	11.5	11	9.5	11.5	13
Computer repair services	13.5	11	10	9.5	10.5
Financial services	12	11.5	10.5	12	13.5
Hairdressing or other beauty treatment services	13	12	10.5	11.5	13
Catering services, including restaurants and takeaways	13	12	10.5	11	12.5
Real estate activity not listed elsewhere	13	12	11	12.5	14
Architect, civil and structural engineer or surveyor	13.5	12.5	11	13	14.5
Management consultancy	13.5	12.5	11	12.5	14
Accountancy or book-keeping	13.5	13	11.5	13	14.5
Computer and IT consultancy or data processing	14.5	13	11.5	13	14.5
Lawyer or legal services	13.5	13	12	13	14.5
Labour-only building or construction services*	14.5	13.5	11.5	13.5	14.5
Number of FRS categories	54	56	56	56	56
Number of flat rates	17	16	16	18	17
Range of flat rates	5 - 14.5	2 - 13.5	2 - 12	3.5 - 13.5	4 - 14.5
Standard VAT rate	17.5	17.5	15	17.5	20

Appendix B

Calculation of FRS gains

To calculate counterfactual FRS liability I need to multiply FRS turnover by the applicable flat rate. FRS turnover is total gross business income which should include exempt, zero rated, reduced rated, and standard rated sales as well as any VAT received on sales. Under normal VAT accounting, VAT liability is VAT received on sales minus VAT paid on purchases (subject to certain qualifying rules). VAT traders report net of tax sales and purchases and corresponding VAT on them in VAT returns. Reported sales includes exempt, zero-rated, reduced-rated, and standard-rated sales but doesn't include VAT itself. Therefore, to arrive at FRS turnover I add up reported sales and the corresponding VAT.

In order to calculate FRS gains for VAT traders, I first assigned a flat rate to each trader (base on reported SIC codes) and then calculated FRS turnover from returns data (as above). FRS gains is then derived as the difference between reported VAT liability and calculated counterfactual FRS liability. Assuming the assigned flat rate is τ_F and FRS turnover is S_g I calculate FRS gains as follows

$$
\begin{aligned}
\text{FRS gains} &= T_V - T_F \\
T_F &= \tau_F \times S_g \\
T_V &= T_S - T_P
\end{aligned}
$$

T_F represents FRS liability while T_V shows reported net VAT which itself is the difference between sales VAT (T_S) and purchases VAT (T_P). FRS turnover is basically sum of net of VAT sales and VAT on sales. Both of these values are reported on VAT tax return.

In the next subsection, I explain the details of how I assigned flat rates to VAT traders. Then I present reliability checks I have done to make sure the assigned rates are correct. Finally I discuss several complications in the calculation of gains.

B.1 Assigning flat rates to traders

In principle there are two ways to assign the appropriate flat rate to each firm. In the first method flat rates are set based on observed effective output tax rate for FRS firms within the same SIC2007 code. Two conditions are required for proper functioning of this method: a) non zero mass of FRS traders for most sectors and b) a tight distribution of effective output tax rates for FRS traders in each sector. Out of 719 SIC2007 codes, 304 sectors have less than 30 FRS traders. Ignoring low FRS sectors however removes only about 2% of FRS eligible traders. The more serious issue with this method is the disperse distribution of flat rates within sectors. The scheme requires traders to account for special transactions outside the scheme but report only the sum of all transactions under outputs and output VAT. For example if a trader purchases services (e.g. consultancy) from another EU member state, these are accounted under the reverse charge scheme at the relevant VAT rate (standard, reduced, or zero) but I don't observe each element separately. Therefore the observed effective output tax rate for FRS traders may not reflect the applicable flat rate. Furthermore, some traders might join FRS in the middle of an accounting period, and therefore have a weighted average of standard rate and flat rates as effective output tax rate. The 1 percentage point discount on new VAT registrations further complicates matters.

Therefore, I use traders' reported SIC2007 codes to assign flat rates. HMRC publishes list of applicable flat rates for around 56 "categories of business" and lists several associated "trade names" under each category (332 trade names). I match these trade names to SIC2007 code descriptions from the Office of National Statistics (ONS) to form a mapping between reported SIC2007 codes and published flat rates. For example, ONS describes SIC2007 code of 70229 as "management consultancy activities (other than financial management)". This description matches with the FRS category for "management consultancy" with $\tau_F = 12.5$ percent during 2004-07. Using this manual matching, 78 percent of FRS eligible traders are assigned a flat rate. The largest sectors left out are construction and part of retail sectors because reported SIC2007 codes map to several flat rates. Table B.1 lists the main

Table B.1: Main sectors that are not assigned a flat rate

SIC2007	ONS description	Why unassigned?
41100 to 41202 (3 codes)	Construction of buildings	Both sectors might include “labor-only” or “general” building or construction services based on the share of labor inputs. The former has a flat rate of 8.5 percent while the latter’s 13.5 percent during 2004-2007.
43120 to 43999 (12 codes)	Specialized construction activities	
47190	Other retail sale in non-specialized stores	This includes department stores, general stores (food not predominant), and household stores. Depending on share of sales they could fall in different FRS categories.
47710 47721	Retail sale of clothing in specialised stores Retail sale of footwear and leather goods in specialised stores	Codes combine sale of children and adult clothing but FRS (and VAT) distinguishes between the two.
68100 to 68320 (4 codes)	Real estate activities	Estate agency or property management services

sectors left out of the analysis and the reason why flat rates could not be assigned.

In the last three years of the sample (2008-9 to 2010-11) there were two flat rates in place during a single financial year (due to changes in the standard VAT rate). I use the variable “stagger” that shows the periods returns correspond to, to assign appropriately weighted flat rates to traders during this period. The full weighting used in the assignments are shown in table B.2. For example, during 2008-9 financial year the standard VAT rate was reduced from 17.5 to 15 percent between 1 December 2008 and 31 December 2009. This means there are two sets of flat rates applicable during this time. I denote the pre December 2008 flat rates by $\tau_{F,1}$ and post this time by $\tau_{F,2}$. For a trader submitting annual returns at the end of March 2009 (stagger equal to 0 or 1), I use a weight of $^8/_{12}$ and $^4/_{12}$ on $\tau_{F,1}$ and $\tau_{F,2}$ respectively to arrive at the year-wide flat rates, i.e. $\tau_{F,2008-9} = {}^8/_{12} \times \tau_{F,1} + {}^4/_{12} \times \tau_{F,2}$. HMRC advises traders to use the appropriate rates on sales done before and after 1 December 2008, but I don’t observe the break down of sales. Therefore, the method explained here is equivalent to assuming a uniform distribution of sales across all months. The degree of measurement error depends on the extent that sales differ across months (e.g. December is a high sales volume period for retailers) and the ability of traders

Table B.2: Weights used for assignment of flat rates during the change years

Return period for 2008-9	Weights for 2008-9 (change: 1 December 2008)		Weights for 2009-10 (change: 1 January 2010)		Weights for 2010-11 (change: 4 January 2011)	
	$\tau_{F,1}$	$\tau_{F,2}$	$\tau_{F,2}$	$\tau_{F,3}$	$\tau_{F,3}$	$\tau_{F,4}$
1 April 2008 - 31 March 2009	8/12	4/12	9/12	3/12	8/12	4/12
1 February 2008 - 31 January 2009	10/12	2/12	11/12	1/12	10/12	2/12
1 March 2008 - 28 February 2009	9/12	3/12	10/12	2/12	9/12	3/12
not sure (left out)	-	-	-	-	-	-

Notes: $\tau_{F,1}$ shows the flat rate applicable from January 2004 - 30 November 2008, $\tau_{F,2}$ is flat rate during 1 December 2008 - 31 December 2009, $\tau_{F,3}$ is for 1 January 2010 - 3 January 2011, and $\tau_{F,4}$ is for 4 January 2011 - onwards.

to shift reported sales to favorable tax periods. A look at distribution of effective output and input tax rates for VAT traders confirms there is a significant mass of traders with effective tax rates exactly at the weighted average of standard rates using the weights in table B.2.

B.2 Assignment Reliability

To check the reliability of flat rate assignment I use the observed flat rates for existing FRS traders in the same SIC2007 code. I calculate the observed flat rates, τ_F^o, as the ratio of output VAT over reported gross outputs. To get a clean measure of applicable flat rates, I restrict the sample of FRS traders to those satisfying three conditions: a) on FRS for exactly 12 months, b) passed the FRS discount window, and c) with τ_F^o smaller or equal to the maximum applicable flat rate. The three restrictions help to solve for some of the issues mentioned above about using the observed flat rates.

Figure B.1 shows the histogram of the difference between assigned flat rates and observed ones, $\tau_{F,s}^a - \tau_{F,si}^o$, for the group of FRS traders satisfying the three conditions (subscripts i and s denote traders and sectors, superscripts a and o denote assigned and observed flat rates). The figure shows two encouraging patterns. First, the distribution of the deviation is almost symmetric around zero. This suggests, the difference between observed flat rates and assigned ones is not systematic and reflects

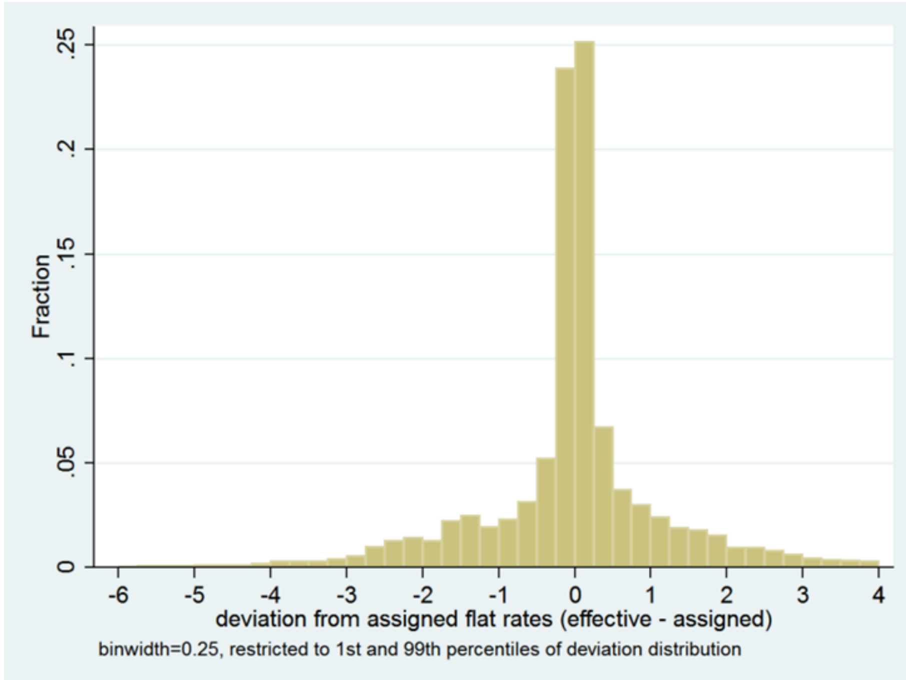

Figure B.1: Histogram of the difference between assigned and observed flat rates

trader specific circumstances and on average the reported number of gainers won't be biased upward or downward. Second, 60 percent of the mass falls in the range of -0.5 to 0.5 percentage points deviation.

To further check whether certain sectors show a high degree of deviation while others don't, I define $\bar{\tau}_{F,s}$ to be average absolute difference between assigned and observed flat rates in sector s:

$$\bar{\tau}_{F,s} = {}^{1}\!/_{N} \sum_{i} \mid \tau^{a}_{F,s} - \tau^{o}_{F,si} \mid \tag{B.1}$$

where N is the number of included FRS traders in sector s and summation is done over the absolute difference for such traders. A large $\bar{\tau}_{F,s}$ signals potential problems with the assignment process. Table B.3 shows the result of this reliability check. 55 percent of eligible VAT traders are in sectors with an average deviation of less than 2 percentage points. These sectors also have higher fraction of FRS traders and gainers.

$\bar{\tau}_{F,s}$ is susceptible to presence of outliers. Therefore, to make sure the assigned flat rates are correct, I investigated the histograms of the observed flat rates for all FRS traders within the sectors with $\bar{\tau}_{F,s} \geq 1$. In all sectors the histograms had a clear mode at the assigned rate. As a final precaution, I re-checked the matching of

Table B.3: Sectoral average absolute difference between assigned and observed flat rates

Average absolute difference ($\bar{\tau}_{F,s}$)	Number of Sectors	Number of observations			% FRS	% gainer
		FRS traders	FRS eligible	FRS gainers		
$[0, 0.5]$	121	210,666	1,224,939	300,134	17%	30%
$(0.5, 1]$	84	46,720	596,268	124,237	8%	23%
$(1, 1.5]$	55	20,792	474,182	73,305	4%	16%
$(1.5, 2]$	48	8,167	344,337	43,046	2%	13%
$(2, \infty)$	254	117,122	1,569,015	59,100	7%	4%
Total	562	403,467	4,208,741	599,822	10%	16%

Notes: The difference between numbers here and numbers in the paper is because some sectors with smaller than 30 FRS trader or FRS gainers are removed from this table.

sectoral descriptions to HMRC trade names for these sectors and found no error or ambiguity.

B.3 Complications in calculation of gains

There are two potential sources of error in calculation of counterfactual FRS liability. First, I use Standard Industry Classification (SIC) codes to assign flat rates but reported SIC codes are usually based on traders declared activities at the time of VAT registration. Some traders might be involved in activities other than those implied by SIC codes leading to measurement error (see appendix C for other errors in SIC codes). While it is not clear whether this causes a systematic over or under estimate of gains, setting the flat rates to the maximum applicable rate in each year shows still 12% of eligible traders benefit from FRS (table 1.5 column (6)). This is a very conservative estimate of FRS gains and still a significant number of traders benefit. Using this method I can estimate gains for categories that I was unable to assign a flat rate. Results show 9% of all eligible traders benefit from FRS under this scenario. This estimate is encouraging and shows the sample of traders left out of the analysis (unassigned flat rate) are not very different.

The second source of error is unobservable complications in the calculation of FRS turnover. Normally FRS turnover is gross turnover, i.e. net sales plus VAT received on sales, but certain transactions are treated differently. Reverse charge transactions are accounted for by purchasing partner as if they are self supplied. VAT on these items appears as output VAT and could be reclaimed as input VAT even under FRS. In FRS liability calculations I can't separate reverse charge transactions and hence overestimate FRS liability because I ignore the possibility of reclaiming input VAT. Similarly provisions for bad debt relief under FRS are ignored leading to an overestimate of FRS liability. Therefore FRS turnover errors are likely to lead to an overestimate of FRS liability and an underestimate of FRS gains.

There are other reasons to believe that the actual number of FRS beneficiaries is higher than what I estimated. First, as mentioned earlier ignoring deductibility of input VAT on certain capital goods results in an underestimation of FRS gains. In my sample 34% of FRS traders claim any input VAT with an average of £1,350. Therefore this could potentially be a large factor working against me. Second, I ignore the 1 percentage point discount on flat rates for new VAT registrations which leads to an underestimate of gains for the population of new entrants. Considering this raises the fraction of gainers by 1 percentage point to 27% of eligible traders. Third, I ignore FRS compliance cost saving which leads to an underestimate of the number of gainers. Finally, I calculate counterfactual liability based on realized sales under VAT accounting. The optimal level of sales however could be different under FRS which leads to higher FRS profits than what I estimate.

Appendix C

Data cleaning procedures for chapter 1

In this appendix I explain all the cleaning and adjustment procedures I have done on the data used in chapter 1 of the thesis.

C.1 SIC2007 corrections

The VAT returns data include a variable that capture the Standard Industry Classification (SIC) code of traders' main activity. HMRC uses descriptions traders declare in question 6 of VAT 1 - Application for Registration form to construct SIC codes but I don't know the exact procedures followed. As SIC codes are used to assign flat rates to traders they hugely influence FRS gains and the analysis in this paper. Therefore it is crucial to make sure this variable is correctly capturing traders' activities.

The main complication in use of SIC codes is the change in the classification system in 2007. Office of National Statistics (ONS), the body responsible for publishing and maintaining of SIC, revised the system in 2007. The SIC codes reported in VAT data should correspond to SIC2003 codes for 2004-5 until 2006-7 financial years and then map to SIC2007 codes for 2007-8 until 2010-11 financial years. To check this, I match SIC2003 and SIC2007 codes from ONS to those reported in the VAT data in the respective periods.

As table C.1 reports, there are very few missing SIC codes in VAT data (column (2)). For firms reporting a correct (in the sense defined below) and constant SIC2007 over

Table C.1: Mis-matches in SIC codes

year	Total observations	Missing SIC in VAT data	Unmatched SIC	Unique SIC codes in VAT
	(1)	(2)	(3)	(4)
2004	1,894,281	<30	2,275	700
2005	2,177,146	8386	13,819	962
2006	2,221,095	<30	2,738	701
2007	2,118,581	<30	114,164	1,365
2008	2,173,988	79	30,684	1,330
2009	2,123,464	<30	15,077	799
2010	2,120,600	<30	16,396	801
Total	14,829,155	8,482	197,144	-

Notes: Column (3) shows the number of observations that had non-missing SIC codes in VAT data but didn't match with SIC codes from ONS. Number of unique SIC codes in ONS data is 699 and 728 respectively in 2003 and 2007 classifications.

the non-missing years, I fill out the missing SIC observations. There is, however, a significant number of mis-matches between ONS and VAT SIC codes in 2007-8 financial year (column (3)). This suggests not all SIC codes reported in 2007 are based on SIC2007 and some of the observations continue to use SIC2003 in this year. Column (4) confirms this idea by showing that in 2007 and 2008 there are significantly more unique codes in the VAT data than the ones exist in ONS classification. Furthermore, when I match the unmatched codes from 2007-8 financial year to SIC2003 codes, 579 unique codes are matched up. This is despite the fact that only two codes remain unchanged moving from 2003 to 2007 classification (ONS tables).

These observations lead me to believe that some traders still report SIC2003 codes in 2007-8 financial year. While the numbers of unmatched observations seem small in table C.1, the problem is deeper. There are around 80 codes that are common in the two classifications but map to different codes. For example "01240" in SIC2003 is "farming of poultry" and maps to "01470" in SIC2007. But the same SIC2003 code of "01240" exists in SIC2007 classification and corresponds to "growing of pome fruits and stone fruits". In other words, not all the matched observations in table C.1 correspond to correct SIC2007 codes. Fortunately, as I said earlier, there are only two SIC2003 codes that map to an identical code in 2007. Therefore I can safely assume that all traders that don't change their SIC codes when moving from

Table C.2: Change of SIC2007 codes across years

Transition years	Number of SIC2007 code switchers	
	Before correction	After correction
From 2004 to 2005	26,821	26,821
From 2005 to 2006	10,989	15,524
From 2006 to 2007	774,983	19,672
From 2007 to 2008	10,197	8,712
From 2008 to 2009	20,383	5,876
From 2009 to 2010	5,106	4,776

Notes: Table shows the number of traders that change their five digit SIC2007 codes moving from one financial year to the following before and after the corrections mentioned in the text are applied.

financial year 2006-7 to 2007-8 are mistakenly reporting SIC2003 codes. If these firms keep on reporting the same SIC code in 2008-9 financial year I still assume they are reporting SIC2003 codes and so on.

The flat rate assignments are based on SIC2007 codes (not SIC2003 codes). Therefore, I need to construct a mapping between SIC2007 and SIC2003 for traders reporting SIC2003 codes in VAT data (majority during financial years before 2007). ONS provides the correspondence between the two classification systems. The difficulty is, however, the multiple to multiple mapping of classifications. 418 SIC2003 codes correspond to a unique SIC2007 code but 281 SIC2003 codes could correspond to up to 15 different SIC2007 codes (136 codes correspond to 2). I randomly pick one of the SIC2007 codes that correspond to the given SIC2003. To partly correct for potential mis-assignments I use the SIC2007 codes reported in VAT data for the same trader from 2007-8 onwards and assign this instead of my random assignment. No corrections are, however, made for traders not observed after 2007-8.

Table C.2 shows the number of traders changing SIC2007 codes from one year to the following. In 2007-8 when the classification system changed, I see an unexpected increase in number of switchers. This is due to the two problems mentioned above: mis-reporting of SIC2003 in place of SIC2007 codes after the change and multiplicity of correspondence between SIC2003 and SIC2007. Carrying out the corrections outlined above, however, results in a much more reasonable number of switchers.

I have replace date of joining FRS with missing if it was prior to 1 April 2002 or after 1 April 2012. Furthermore, a sizable number of traders report FRS date to be missing in 2006-07 financial year. I replace for FRS date using 2005-06 or 2007-08 financial years for these traders. Finally, I use the minimum recorded FRS date for traders that report multiple FRS dates but don't report a change in their FRS condition.

C.2 Deleted observations

In order to increase the reliability of the analysis and as reported in table 1.2 I have dropped several observations. In this section I explain each set of dropped observations and the reason for leaving them out of the analysis.

The first set of observations removed are for traders that are reported to be inactive or deregistered. This is through two variables in the VAT dataset. First, I only keep returns associated with traders reporting as "not deregistered" (dereg_ind equal to 0). I also keep traders reported to be alive (actively trading) at the end of financial year. Deregistration is associated with special treatments and I remove these observations not to confound such special treatments with FRS gains.

The second set of observations removed are based on reported values of sales and purchases. I remove traders that report a zero or missing value for total outputs. These traders either have all tax variables equal to zero (inactive) or have high purchases (e.g. because of start-up costs). I also drop observations that fall above the 99th percentile of the overall distribution of sales or purchases respectively. This is to make sure that outliers don't influence the results. Notice the percentiles of the distributions are calculated after zero sales observations are dropped.

The third set of observations I remove are for traders that show unusual values for effective input and output tax rates. I define effective output tax rate as the ratio of sales VAT to net sales (both are reported in returns). This could vary from zero to the standard VAT rate. For traders in standard rated activities (e.g. retail of household appliances like TV) the effective output tax rate should be equal to the standard VAT rate (equal to 17.5 percent for 2004-5 to 2007-8). Similarly I define the effective input tax rate as the ratio of purchases VAT to net purchases. Based on the distribution of inputs used by each trader the effective input tax rate could vary from zero to the standard VAT rate. Despite this I observe several traders with effective tax rates higher than standard VAT rate. These might be accounting for errors in previous returns, getting bad debt relief, accounting for penalties, and

Table C.3: Number of observations dropped in the cleaning process

Stage	Number of obs	FRS traders
Initial sample	14,829,026	1,084,737
Droppings	1,517,647	2,677
Group 1: Inactive traders	2,873,609	100,926
Group 2: Unusual sales or purchases	837,436	-
Group 3: Unusual effective input/output tax rates	260,078	2,116
Group 4: Other ownership forms	10,460,181	964,356
Cleaned sample	10,460,181	964,356

Notes: Adding individual number of observation for each cleaning step doesn't give total obs dropped because there is overlap between different categories.

other special cases. I drop all traders that show an effective input or output tax rate higher than the standard rate plus 0.5 percentage points (e.g. I drop traders with effective input or output tax rate higher than 18 percent when the standard rate is 17.5 percent).

The fourth set of observations dropped are for traders that report to be registered as clubs, associations, charities, and other organizations. In other words I only include VAT registered traders that report to be a sole proprietor, a partnership, or a limited company (incorporation). Table C.3 shows the number of observations under each of the four categories above and reports the fraction of FRS traders in each sub-sample.