Essays on Political Economy, Inequality and Development.

Marta Schoch

PhD in Economics

University of Sussex

September 2019

Statement

I hereby declare that this thesis has not been and will not be, submitted in whole or in part to another University for the award of any other degree.

Signature:

University of Sussex

Marta Schoch PhD in Economics

Essays on Political Economy, Inequality and Development.

This thesis comprises three chapters.

The first chapter investigates how wealth inequality affects preferences for redistribution and voting behavior. Exploiting house price shocks in England and Wales, it evaluates the effect of wealth inequality on the preferences and voting behavior of homeowners relative to renters. It tests whether homeowners become less supportive of redistribution and in turn more likely to vote for the Conservative Party. The results show that homeowners experiencing positive shocks to their housing wealth are more likely to vote for stated party, but do not decrease their support for the government ownership of public services. Heterogeneity in home value and endowment effects might explain the results.

The second chapter looks at the effect of conscription on political ideology, voting participation and national identity. The motivation of this analysis comes from a renewed policy interest in bringing back conscription to foster civic participation. Considering the costs of conscription identified in the economic literature, this chapter aims at evaluating the effectiveness of this policy in achieving its proposed goal. Using a regression discontinuity design, it uses the introduction of conscription in West Germany to investigate its long-term effects on right-wing ideology and voting participation. The results show no statistically significant effect of conscription on any of the outcomes considered. These results are corroborated when conducting the analysis using a similar empirical methodology and Spanish data.

Lastly, the third chapter analyses the effect of migration on household consumption in rural Ethiopia. Using panel data, it analyses the effect of having at least one migrant in the household on total, food and non-food consumption. Variation in migration over time allows me to unpack heterogeneous effects by duration of migration, possibly explained by changes in the composition of the household.

Acknowledgments

This thesis is the result of a long learning process. Four years ago, I could have not imagined how much I would have grown professionally and, I say this with great satisfaction, how much I would have learned. For this reason, I would like to thank who helped me reaching this goal.

I am grateful to my supervisors, Dr. Julie Litchfield and Dr. Pedro Rosa Dias, for the time they spent guiding me, for their encouragement and for their constant support. They saw the potential of this work from the first day of my doctoral studies and kept reminding me of it, even when I could not see it. I could have not asked for better mentors in these years.

I would like to thank the Economic and Social Research Council for funding my doctoral studies and for allowing me to focus on my research. Similarly, I would like to thank the Migrating out of Poverty programme for giving me access to the data needed to write part of this thesis and for including me in a thrilling research project over the last year. I would also like to thank the department of Economics at the University of Sussex for giving me the opportunity to embark on this journey and for all the feedback received by the faculty over the years. In particular, I would like to thank prof. Barry Reilly for his precious help in developing both my research and my professional experience. Within the department, I cannot thank the PhD community enough for their invaluable help. I came in the office every day knowing that I could count on any of my peers for professional and, mostly, moral support.

I am grateful to Professor Corrado Benassi, who encouraged me to pursue my doctorate. I have many reasons to thank him, from understanding my research interest before I did to guiding me in the choice of a research department, but mostly I would like to thank him for showing me that I could contribute to a wider research community. Finding mentors that give you a perspective on the bigger picture is rare and, in my case, it made all the difference.

I would be grateful to have done this PhD just to have found the friendship of my dearest Graziana. Her support in good and bad periods and the time spent together has often been the highlight of my days in these past years. I thank Silvia and Juanqui, for giving me the comfort zone only siblings can give: where everything is going to be okay even if you do not know what the next research question will be. Edgar, whose contribution to this thesis started with his explanation of what an IV is and continues every day with countless requests of proofreading and conversations about identification strategies. “The PhD is a life-changing experience, where you will likely meet your life partner”, little did I know of how much truth there was in that conference’s presidential address three years ago and that I would turn to you after all.

Lastly, I owe this accomplishment to my mother. As one of the first women to go to high school from her remote village in Sardinia, and the only one to go all the way up to a PhD and a fantastic career, she inspired me since I was born. Her example, her values, her strong belief in the power of education and her constant faith in my capabilities have been the reason behind all my achievements. This thesis is for her.

Contents

Introduction

This thesis comprises three papers.

The first and second paper investigate, respectively, the relationship between inequality and political preferences and the formation of such preferences, national identity and voting participation. In the first paper, the analysis focuses on the effect of wealth inequality in shaping individual preferences for redistribution and the support for political parties. This analysis contributes to the empirical literature on inequality and political preferences, as well as shedding light on the current debate on the consequences of rising wealth inequality.

Following from this interest in understanding political preferences, the second paper analyzes their formation and the effect of a specific policy, rather than inequality, in shaping these outcomes. In particular, the paper investigates the effect of conscription on political ideology, voting participation and national identity. This paper is motivated by the recent interest among European policy-makers in bringing back national service as a way of fostering youth civic engagement and social cohesion, and is aimed at providing causal evidence on the effectiveness of conscription in achieving its proposed goal.

National identity, social cohesion and civic participation are central themes in the current political discourse. Often, policy-makers discuss and promote policies aimed at strengthening these values in response to large migratory flows from lower-income countries. However, this political debate is usually focused on the effects of migration at destination and ignores the discussion on its effects on migrant-sending households. The third paper aims at providing evidence on the relationship between migration and household welfare in low-income countries. This analysis is of particular relevance in contexts where migration might contribute to lift households out of poverty, as in the case of rural Ethiopia.

This thesis contributes in several ways to the existing literature.

The first paper investigates the relationship between wealth inequality and political preferences. In particular, it tests whether higher wealth inequality affects both individual preferences for redistribution and voting for the Conservative party in England and Wales. The motivation for this paper spurs from a recent interest in rising wealth inequality and its consequences on political and economic outcomes.

The paper contributes to the well-known economic literature on inequality and preferences for redistribution. Starting from the median voter theory (Meltzer and Richard, 1981), several scholars have looked at the relationship between rising income inequality and one's preferences for redistribution. The empirical evidence in this field, however, has mostly been focused on wages and income, with only few studies analyzing the distribution of wealth (Di Tella et al., 2007; Caprettini et al., 2019), traditionally more difficult to measure. Thus, this paper contributes to the literature by providing causal evidence of the effect of wealth inequality on political preferences.

Of particular relevance for this paper is the work by Barth and Moene (2016), which investigates the effect of inequality on both preferences for redistribution and on the support for a party with a political agenda in line with such preferences. Barth and Moene (2016) also test these relationships empirically and provide correlations showing that as individuals become more likely to support redistribution, they are more likely to support more left-wing parties. This paper innovates by testing whether these relationships are causal and whether they apply to changes in the distribution of wealth. Studies looking at wealth show that land redistribution affects both individual preferences (Di Tella et al., 2007) and voting decisions (Caprettini et al., 2019). However, and to the best of my knowledge, none of these studies looks at both preferences and political party support.

To carry out the analysis, I identify an increase in wealth inequality between homeowners and renters using the house price boom in England and Wales for the period between 1995 and 2007. Exploiting county-level variation in house price shocks and using individual level data from the British Household Panel Survey, I evaluate the effect of house price shocks on homeowners preferences and voting behavior relative to renters. Following the literature, I test whether richer homeowners become less supportive of redistribution, and in turn are more likely to vote for the Conservative Party. I find that homeowners experiencing positive house price shocks are more likely to support the Conservative party. However, I do not find a decrease in their support for redistribution.

The results show that homeowners preferences vary depending on the value of the property owned. I also find that when using self-reported value of one's property the expected mechanisms are fully corroborated: higher self-reported house values are correlated with a higher probability of voting for the Conservative party, and lower preferences for redistribution. This second result might suggest a link between endowment effects and homeowners preferences.

Following from this interest in the analysis of political preferences, the second paper evaluates whether conscription affects the formation of political ideology, voting participation and national identity. This analysis contributes both to the policy discussion aimed at bringing back conscription in some European countries, as well as to the economic literature on the effects of conscription. In recent years, several policy makers in Europe have discussed the possibility of bringing back conscription. Among these, some have motivated such reform as a way to foster civic engagement and social cohesion among the youth (for media coverage of this political discussion, see here and here).

However, bringing back conscription would have non trivial costs both in terms of government budget (being one of the main reasons why conscription has been abolished in Europe) and in terms of economic outcomes. These effects have long been investigated in the economic literature. Angrist (1990) first showed the existence of a wage gap between Vietnam-era draftees and non-draftees, finding negative long-term effects on wages for the former. Following from that seminal work, several studies investigated the effect of conscription on wages, showing mixed results (Imbens and Klaauw, 1995; Grenet et al., 2011; Bauer et al., 2012). More recently, Card and Cardoso (2012) find that peacetime conscription has positive effects on wages for men with lower levels of education, suggesting that conscription might have positive, though small, effects on wages as it allows men to learn new skills. The literature has further investigated the effects of conscription on other outcomes. It finds that conscription has negative effects on health (Autor et al., 2011; Bedard and Deschênes, 2006), it might increase or decrease one's investment in higher education (Cipollone and Rosolia, 2007) and it increases crime rates (Galiani et al., 2011).

Following from these evidence on the costs of conscription, this paper contributes to the policy discussion by presenting causal evidence of the effect of conscription on political ideology, voting participation and national identity. Thus, the main contribution of the paper is to inform the policy discussion aimed at bringing back conscription as a way to affect such values. This evidence is crucial to evaluate the effectiveness of this policy in light of the costs that come with it.

In this paper, I conduct an empirical analysis using a regression discontinuity design and exploiting the introduction of conscription in West Germany to investigate its long-term effects on right-wing ideology, and voting participation. Using this empirical method, I am able to exploit an exogenous shock to the probability of being drafted for men in the treatment group based on their date of birth, and compare them to individuals in the control group who were too old to be eligible for the draft.

The results show no statistically significant effect of conscription on any of the outcomes considered. These results are corroborated by a series of robustness checks. As part of these checks, I conduct the analysis using a similar reform and Spanish data. The advantage of doing so are several. First, I find no short-term effects of conscription on political attitudes. Second, I do not find statistical significant evidence on the effect of conscription on Spanish national identity. Third, I am able to provide evidence of the external validity of the results. Overall, the analysis does not find convincing evidence that conscription affects political attitudes, suggesting that the intended benefits of this policy might not offset its costs.

Lastly, the third paper investigates the relationship between migration and origin household consumption. The motivation for this paper spurs from the central role of migration in the current political, economic and development agenda (see Sustainable Development Goals 2030). While the political discourse is often focused on the consequences of migration in host countries, the design of effective migratory policies will also depend on its effects on origin countries.

The economic literature has long studied the link between migration and development, identifying people's relocation as a way to increase individual and household welfare (Harris and Todaro, 1970; Stark and Bloom, 1985). This is particularly important for poorer households living under financial constraints and income variability (Clemens et al., 2014; Clemens and Ogden, 2014), which are usually higher for poorer households strongly depending on agriculture (Rosenzweig, 1988). Empirical evidence has shown that migration can have large positive effects on household welfare (Yang, 2008; Gibson and McKenzie, 2014; Clemens and Tiongson, 2017), but that this effect might become negative depending on the context of analysis. In fact, Gibson et al. (2011) find that in the short-run the absence of the migrant has a negative effect on household per capita consumption. These mixed results further motivate the analysis of this paper.

The paper carries out the analysis using panel data from rural Ethiopia. This context is of particular interest both for its high levels of poverty, among the highest in the world, and for the relevance of migration in the country. Several studies have looked at the relationship between migration and land rights (De Brauw and Mueller, 2012), youth agricultural employment (Bezu and Holden, 2014; Mueller et al., 2018), as well as migrant's consumption and income (de Brauw et al., 2017). Following from these findings, this paper contributes to the literature by investigating the effect of having at least one migrant in the household on household consumption.

The paper contributes further by conducting the analysis using newly available panel data (Gibson and McKenzie, 2014; Beegle et al., 2011). I find that migration has a negative effect on per capita consumption, at least in the short term. I investigate two possible mechanisms behind this result, finding that migration affects household composition and showing qualitative evidence on the change in migration experience for the period of analysis. Following from these results, the paper contributes further by showing empirical evidence on the effect of migration in Sub-Saharan countries. This is particularly important to evaluate the welfare effects of migration in one of the poorest areas in the world, especially given the relatively scarcer evidence focused on this region (Lucas, 2006).

The thesis is organized as follows. Section 1 presents the analysis on the effects of wealth inequality on political preferences. Section 2 presents the analysis on the effects of conscription on political ideology, voting participation and national identity. Section 3 investigates the effects of migration on household consumption in rural Ethiopia. The general conclusions, bibliography and the appendices to each of the papers follow.

1 The effect of wealth inequality on political preferences: evidence from the house price boom in England and Wales.

1.1 Introduction

Rising wealth inequality is at the center of the political and economic debate (Alvaredo et al., 2017a; Piketty et al. 2014;for a political discussion on the topic see here). Yet, analyzing the effects of wealth inequality separately from income and wages has been traditionally more difficult. In particular, while there is a large literature showing the salience of income and wage inequality in affecting one's political preferences (Alesina and Giuliano, 2009; Giuliano and Spilimbergo 2014; Fetzer 2019), fewer studies focus on changes to the distribution of wealth (Di Tella et al., 2007, Caprettini et al., 2019). Recent research (Barth and Moene, 2016) has provided theoretical evidence showing that increasing wage inequality not only affects individual preferences for redistribution, but leads to a change in voting for one party or another, according to the party's stand on redistributive policies.

The aim of this paper is to contribute to this literature and investigate how wealth inequality affects both preferences for redistribution and political party support.

This relates to a long standing literature in economics on preferences for redistribution. Seminal work in this field is based on the predictions set out in the median voter theory (Meltzer and Richard, 1981), stating that higher inequality will result in higher redistribution as the median voter will be worse off. Such prediction has been extensively tested empirically and the evidence shows that inequality matters in determining one's preferences (a review of this literature can be found in Alesina and Giuliano, 2009). Among these studies, Giuliano and Spilimbergo (2014) show that local-level macroeconomic shocks affect both preferences for redistribution and the probability of supporting one party or another in the U.S., suggesting that preferences for redistribution are coherent with party support. This empirical evidence, however, has rarely focused on the link between wealth inequality and political preferences, with the exception of few studies looking at the effect of wealth redistribution on individual preferences and voting behavior. When looking at land, Di Tella et al. (2007) find that redistribution reforms make individuals more in favor of free-market policies, while Caprettini et al. (2019) find that they increase the probability of winning the elections for the incumbent party.

The closest study to this paper is Barth and Moene (2016), who use a model to formalize the effect of inequality on both preferences for redistribution and voting behavior and test it using a correlation analysis. The authors demonstrate that increasing wage inequality will result in individuals voting for political parties offering less generous redistributive policies and this will result in lower welfare spending, leading to even higher inequality. They also provide cross-country evidence showing correlations in support of their theory. In particular, they show a positive association between higher wage inequality and the probability of a party less favorable to redistribution holding the majority in the parliament. This is then positively associated with lower levels of public spending. Following from these findings, this paper investigates whether these predictions apply to wealth inequality and whether the relationships they hypothesize are causal.

The first contribution of this paper is its focus on wealth inequality rather than wages or income. To capture the effect of a change in the distribution of wealth, the analysis focuses on changes in the distribution of housing wealth in England and Wales. It does so for two main reasons. First, housing wealth represents one of the larger shares of private wealth in this context (Atkinson and Piketty, 2007; Atkinson et al., 2011), making it a good proxy for individual total wealth. Moreover, and differently from other assets, changes in housing wealth inequality affect a large share of the population in the context of analysis, where the home-ownership rate is of around 70 percent. Second, the effects of housing wealth inequality between homeowners and renters are particularly relevant in this context. A large literature investigates the effects of the housing price boom and its different distributional effects on homeowners and renters in the U.K..

Increasing housing prices affect homeowners differently in terms of labor market decisions, making them more likely to anticipate their retirement (Disney and Gathergood, 2018), they affect positively homeowners health and increase their demand for private healthcare (Fichera and Gathergood, 2016). Higher house prices increase segregation in British cities (Sá, 2015), and other studies provide evidence of the effect of house price shocks on consumption smoothing over the life-cycle (Disney et al., 2010). The relevance of these effects brings housing policy at the center of the British economic and political discussion (Hood and Joyce, 2017), as well as at the center of the political program of the major parties in the country. Proposals as “Housing for the many” (https://labour.org.uk/issues/housing-for-the-many/) and “Housing the nation” (https://www.conservatives.com/sharethefacts/2018/01/conservatives-housing-the-nation) are only the most recent political stands on a decade-long debate on the housing market. Following from this discussion, this paper contributes further in showing whether rising wealth inequality also affects voting and political preferences.

The second and main contribution of this paper is to carry out a single-country empirical analysis on the causal effect of wealth inequality on political preferences. In order to move beyond the correlations presented in Barth and Moene (2016), I exploit an exogenous source of variation in wealth inequality to estimate its effect on a change in individual political preferences, addressing both issues of reverse causality and endogeneity that are typical of cross-country correlation studies.

The identification strategy of this study relies on the presence of unanticipated county-level house price shocks experienced in England and Wales between 1995 and 2007. These shocks measure an exogenous increase in wealth inequality and will affect differently homeowners and renters, who share the same geographical area but hold or not housing wealth assets (Disney et al., 2010). If the predictions set out in the literature hold, this increase in wealth inequality should cause a decrease in homeowners preferences for redistributive policies such as taxation and welfare spending (Barth and Moene, 2016). This is because, for example, homeowners experiencing an appreciation of their property would prefer to pay lower taxes. At the same time, homeowners experiencing a positive house price shock should also become more likely to support the Conservative party, traditionally more averse to include welfare policies in their political agenda.

The results of the paper corroborate these mechanisms only partly. On one hand, I find that a positive shock in house prices at the local level leads homeowners to increase their support for the Conservative party. On the other, I do not find evidence of a decrease in pro-public sector attitudes nor in the support for government intervention in the economy.

The third contribution of the paper is to propose two potential mechanisms to explain the main results. First, the paper examines whether the effect on preferences is heterogeneous among homeowners. I do this by using variation in the value of property types as a measure of higher or lower wealth endowment among homeowners. In line with the theory and with empirical evidence from the U.K., homeowners of more expensive properties might benefit more from house price shocks and need to rely less on public services. I find evidence of such heterogeneity.

First, homeowners of more expensive properties (detached and semi-detached houses) are more likely to support the Conservative party than homeowners of cheaper ones (terraced houses and flats in the sample of analysis) . Second, the former are less likely to support redistribution than the latter. These results are in line with the predictions made by the literature and are crucial to interpret the average statistically insignificant result on preferences for redistribution.

Second, I use self-reported value of one's property as an alternative measure of wealth inequality. I find that the hypothesis tested is fully corroborated when conducting such estimation. In fact, an increase in self-reported house value both increases the probability that homeowners support the Conservative party and decreases their support for redistributive measures. This result suggests two things. First, it confirms the presence of endowment effects in the housing market, with homeowners valuing their property more than its market value. Second, it shows that individual perceptions of the property's value strongly correlate with political preferences and voting choices. To the best of this author's knowledge, this constitutes novel evidence on the determinants of political preferences.

The remaining of the paper is articulated as follows. Section 1.2 overviews the literature. Section 1.3 describes the context of interest and section 1.4 the data used in the analysis. Section 1.5 discusses the identification strategy. Section 1.6 illustrates the empirical specification used in the analysis. Section 1.7 provides the summary statistics for the sample of interest. Section 1.8 contains the results, which are investigated further in section 1.9, where I present evidence supporting possible mechanisms. It concludes with the robustness checks used to validate the main results and with the overall conclusions (Section 1.10 and 1.11). Supplementary materials to this paper can be found in Appendix 1.

1.2 Literature Review

This paper contributes to the literature on inequality and political preferences, and to a growing literature on the consequences of wealth inequality.

Both theoretical and empirical work has addressed the limitations of the median voter theory proposed by Meltzer and Richard (1981). In their seminal work, the authors prove that individuals form their preferences according to their distance from the mean level of income. The model predicts that for higher levels of inequality, the median voter will be worse off and thus demand more redistribution. This core version of model has been extended by Piketty (1995) to account for what individuals believe to be driving inequality, i.e. luck or effort. Bénabou and Tirole (2006), then proposed a further extension of the model to account for social mobility.

Aiming at testing these models, empirical work has investigated whether inequality affects preferences for redistribution according to the theoretical predictions (Alesina and Giuliano, 2009). Among the studies more relevant for this paper, Giuliano and Spilimbergo (2014) find that exposure to macroeconomics shocks, such as recessions, changes permanently individual preferences for redistribution. Giuliano and Spilimbergo (2014) also show that an increase in preferences for redistribution is in line with an increase in the probability of having more left-wing political preferences and higher probability of supporting a Democratic candidate in the presidential elections. Giuliano and Spilimbergo (2014) use the probability that an individual supports a Democratic candidate to examine the validity of self-reported political preferences.

Using a similar rationale but focusing on wage inequality rather than exposure to recessions, Barth and Moene (2016) develop a theoretical model showing how higher wage inequality affects negatively both preferences for redistribution and the probability of supporting left-wing parties, traditionally more generous in terms of redistributive policies. In their model, voters support one political party or the other according to the welfare spending proposed, which will then constitute the level of welfare spending implemented by the government (Barth and Moene, 2016). The main assumption is that a reduction in inequality will increase demand for redistribution. The authors stress the social insurance motive behind this increase in demand, i.e. poorer voters becoming richer experience a higher risk of income loss and thus demand for more redistribution for insurance motives.

This translates into higher levels of welfare spending as a consequence of political competition, i.e. higher demand for welfare spending will increase parties' competition in terms of welfare spending promises introduced in their electoral agenda. This relation works in the opposite way for higher levels of inequality: richer individuals demand for less redistribution and so do individuals moving from the middle of the distribution to the left tail. These predictions are confirmed by the correlation analysis conducted by the authors, finding that an increase in the country's Gini coefficient is associated with higher probability of having a right-wing government and lower levels of welfare state. Following from this model, this paper contributes by testing empirically its hypotheses and by focusing on the distribution of wealth, providing novel causal evidence beyond the cross-country correlations provided in Barth and Moene (2016).

The choice of focusing on wealth inequality relates to a growing and recent literature documenting an increase in top wealth inequality (Atkinson et al., 2011; Piketty et al., 2014). The new evidence collected in the World Inequality Database (Alvaredo et al., 2017b) shows that not only private wealth constitutes an increasing share of national income, but that wealth is highly concentrated at the top 10 percent (and even more at the top 1 percent) of the distribution. This levels of wealth inequality could have consequences in terms of limiting inclusive growth, increase inter-generational inequality and increase the role of inheritance (Piketty et al., 2014).

Notwithstanding this increasing interest in high levels of wealth inequality, there are few studies assessing the causal effect of wealth inequality on political preferences. An exception are some studies providing causal evidence on the the effect of land redistribution reforms. Di Tella et al. (2007) finds that landless individuals in Argentina acquiring property rights are more likely to support free market and individualism. Caprettini et al. (2019) show that land redistribution in Italy increases the support for the incumbent party and this effect lasts for as long as for forty years after the reform implementation. However, to the best of this author's knowledge no empirical work focuses on the effects of rising wealth inequality on both preferences for redistribution and political party support.

This paper contributes further to the literature on the consequences of housing wealth inequality between homeowners and renters in England and Wales. Focusing on this context has two main advantages. First, real estate constitutes the larger share of private assets in the U.K.. In fact, the wealth went up to be 4.5 times national income in 1995 to more than 6.3 times in 2007 (Alvaredo et al., 2017b), with rising value of real estate contributing the most to this increase. The World Inequality Database (Alvaredo et al., 2017b) reports that private wealth in the form of housing assets was three times the national income of the U.K. in 2013, increasing steadily from around 1.5 times in 1995. Moreover, and similarly to other developed countries covered in the database, public wealth in the form of housing wealth has been decreasing substantially in the last decades and constitutes less than 10 percent of national income. This plausibly reflects the decrease in council owned houses and the rise in private homeownership rates (Disney and Luo, 2016).

Second, the housing wealth distribution experienced unprecedented changes during the period of analysis (Hills, 2013). Due to both deregulation policies and an unprecedented housing price boom, the distribution of housing between homeowners and renters have changed to the advantage of the former. In particular, the increasing rate of home-ownership, caused by the Right to Buy legislation in 1980 (Disney and Luo, 2016), and the exponential rise in housing assets value following from the 1990s-2000s housing price boom, have motivated several studies looking at the consequences of the resulting wealth effects favoring homeowners. Hills (2013) finds that marketable wealth increased remarkably between 1990 and 2005, as did inequality (independently of cohort effects). Using data from the BHPS the author shows that the increase in median wealth experienced in this period has been mainly explained by an increase in housing wealth and in house prices. Following from their relevance in affecting the wealth distribution, average house prices have been often used by the literature to proxy for house wealth shocks (Attanasio et al., 2011; Attanasio et al., 2009;Campbell and Cocco, 2007).

In particular, Disney et al. (2010) look at how homeowners experiencing unexpected increases in their house wealth (proxied with house price shocks) adjust their consumption over the life cycle. Using similar data and a similar measure of wealth effects, Disney and Gathergood (2018) find that older men homeowners and young female homeowners spouses decrease their supply of labor as a results of local level house wealth effects. The increase in average house prices and the resulting wealth effect have been found to affect also health outcomes. In particular, homeowners experiencing an increase in the value of their home have better health and are more likely to demand for private healthcare (Fichera and Gathergood, 2016). Moreover, the same study finds that an increase in housing wealth value can offset negative effects on health resulting from worsening labor market conditions.

These studies confirm both the relevance of housing wealth for individuals in the U.K. and the presence of differential effects favoring homeowners relative to renters as a result of housing wealth effects. However, to the best of this author's knowledge, no evidence exists on the effect of this increase in wealth inequality between homeowners and renters on political preferences. This paper uses similar data and the same period of analysis of previous studies (Disney et al., 2010; Disney and Gathergood, 2018; Fichera and Gathergood, 2016) to identify whether the above mentioned effects were accompanied by changes in political beliefs and party support.

Lastly, this paper relates to the scarce empirical literature on the consequences of wealth polarization. This is a different concept from inequality, more focused on the appearance and disappearance of groups in the population, such as those of homeowners and renters (Chakravarty, 2015). Most of the work done after Esteban and Ray (1994) seminal work on the definition of polarization, has been focused on income and on the different ways to measure polarization. In general, polarization measures were developed to explain the disappearance of the U.S. middle class during the 1980s (Wolfson, 1994).

D'Ambrosio and Wolff (2001) find that homeownership status is positively correlated with wealth polarization. Following from this literature, the recent increase in political polarization (Autor et al., 2016) has lead researchers to focus on the relationship between rising economic inequality and political polarization (Bonica et al., 2013). The empirical evidence on this link remains mostly limited to correlations between economic and political polarization. Although this paper does not look at measures of polarization itself, it provides evidence on the effect of an increasing gap in housing wealth between homeowners and renters on their political preferences.

1.3 Context

This section aims at providing information on the context of England and Wales in the period of analysis. It focuses on two main aspects: one is the relevance of housing as a measure of individual wealth, the other is its relationship to the political agenda of the Conservative party during that period.

1.3.1 Housing wealth

Real estate is the main component of private wealth in the U.K. (Piketty et al., 2014). The homeownership rate during the early 2000 was of around seventy percent, showing that a consistent share of the British population is affected directly by changes in housing wealth value. In this analysis, I focus on England and Wales as the house price information is available at the local level for the longer period of time and starting from 1995. This year coincided with the start of a housing price boom unprecedented in the post-war period. As a result of large deregulation reforms to the housing market (Disney and Luo, 2016), housing prices more than doubled in all of the regions of the country (with Greater London registering a threefold increase).

Figure 1 shows regional average house prices in England and Wales for the period 1995-2007. This increase in housing prices allows me to identify a clear increase in inequality between homeowners and renters. Evidence of this increase in inequality can also be seen in figure 6 (in Appendix 1), plotting a significant shift to the right in the distribution of homeowners self-reported house value between 1995 and 2007[1].

[1]Data for renters is missing, as by definition renters do not own the property where they reside in (information is available for second or additional property, but it is not exploited in this paper). Self-reported house value will be used in the empirical analysis and explained in Section 1.9

Figure 1: Regional-level average house prices for a typical semi-detached house in England and Wales in the period between 1995 and 2008.

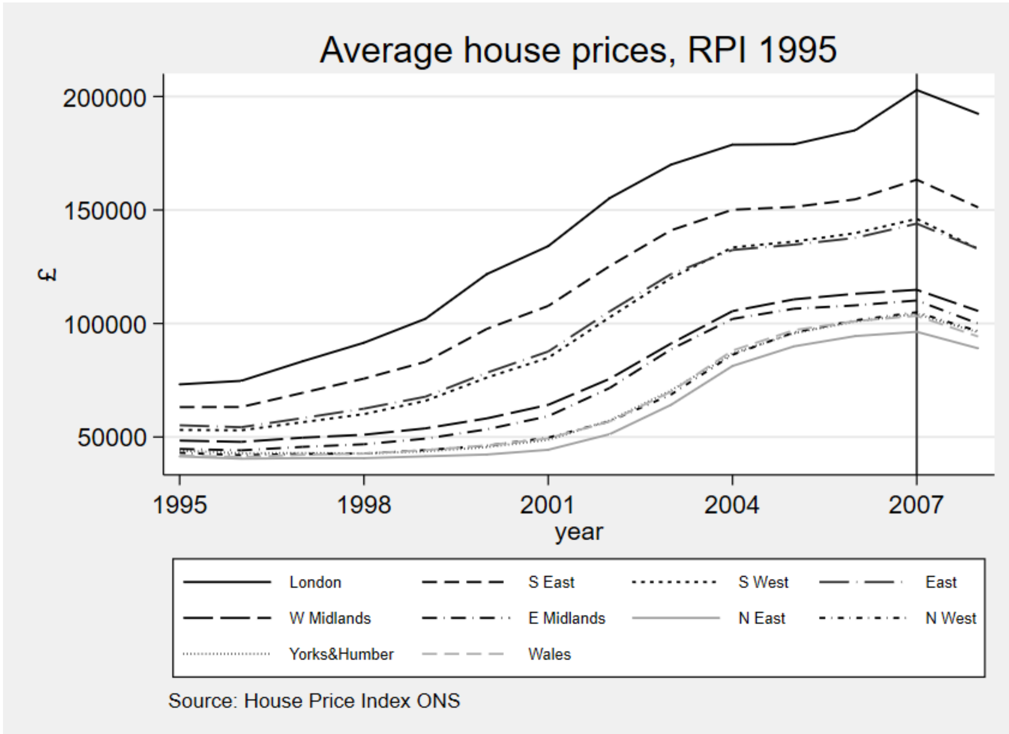

Note: Real average house prices at the regional level for the period 1995 to 2008, vertical line at 2007. From top to bottom indicatevely: Greater London, South East, South West, East, West Midlands, East Midlands, Wales, North West, Yorkshire and the Humber, North East.

1.3.2 Housing and political agenda

Following Barth and Moene (2016), I identify right-wing parties to be less generous in terms of welfare spending and, thus, representing the interests of individuals with lower preferences for redistribution. In the context of the U.K., and during the period of analysis, this party is identified to be the Conservative party.

The Conservative party has traditionally promoted less redistributive policies and lower welfare spending in its political agenda. If the hypothesis of Barth and Moene (2016) holds, richer individuals will be more likely to support the Conservative party as their preferences for welfare spending and for the public role of the State in the economy decrease. A possible rationale is that homeowners favor lower taxes on their property or that rely less on public services relative to renters (Fichera and Gathergood, 2016). Evidence supporting the stand of the Conservative party on redistributive measures is provided in several political discussions and platforms.

The Conservative party has always adopted a clear stand on these matters in its political agenda. In the period 1995-2007, one of the main points of the Conservative manifesto has been to support and facilitate home ownership via tax cuts and advantageous buying possibilities. Similarly, the Conservative party has had, traditionally, a less generous political program in terms of welfare spending and public services provision. The following quotes show some examples of the party's stand on this topic extracted from their own political manifestos.

"*It's time to [...] cut taxes and regulation; time for our schools and hospitals to benefit from choice and freedom; time to show respect to our pensioners; time for real savings not welfare dependency; time to endow our universities; [...]*" Conservative Party Manifesto 2001

"*Lower taxes promote enterprise and growth. But they also promote the right values. Hard working families have suffered from Labour's tax raids on mortgages and marriage, pensions and petrol, buying a home and having a job.*" Conservative Party Manifesto 2005

Alongside the increase in housing wealth inequality seen in figure 6, the support for the Conservative party has been rising in the same period of time (figure 7 in Appendix 1). These two stylized facts point towards a correlation between increasing housing wealth for homeowners and an increase in support for the Conservative party.

In line with this correlation, evidence has shown that an increase in housing prices at the local level affected positively British homeowners relative to renters. In particular, it increased the health status of homeowners and their demand for private healthcare (Fichera and Gathergood, 2016), and it changed their labor supply decisions (Disney and Gathergood, 2018). These effects might point towards a change in homeowners preferences for public services provision, corroborating the change in beliefs suggested by Barth and Moene (2016), as well as a change in political party supported.

1.4 Data

I use two main data sources for this analysis.

First, I use the house price index published monthly by the Office of National Statistics (ONS). This contains county-level information on house prices for England and Wales from

1995 until present. Following from the limitations of using housing price data from different sources (Chandler and Disney, 2014), the ONS house price index represents the latest attempt to harmonize house price information from a multitude of sources to have a more comprehensive measure. It does so by incorporating historical house price information from both the Land Registry data on transactions of sold houses, and from administrative data looking at house characteristics, more useful when capturing the evolution of housing quality. This means that the the house price index also captures increases in prices due to improvements in housing quality, which could mean an underestimation of house price values relative to other indices. The ONS house price index adopts a rolling definition of a typical house, periodically updating transaction data of completed mortgages by "combining the price of all house types in proportion to the frequency with which properties with those characteristics are actually sold." (Chandler and Disney, 2014).

Moreover, because weights are determined using transaction data, these indices are more likely to reflect the prices of the subset of houses that are transacted, rather than of the entire housing stock. The index is estimated using hedonic regression techniques aiming at controlling for variations in terms of different mortgage policies and different sample of mortgage institutions. In fact, the ONS data contains information relying on a survey of most mortgage lenders (70-80 percent of mortgage market). One of the main advantages of this house price index is the availability of housing prices broken down by property types: detached house, semi-detached house, flat, terraced house. This will allow to differentiate the analysis in order to pinpoint the heterogeneity of the effects for different types of homeowners. All the house price data are adjusted to 1995 level prices using the yearly Retail Price Index published by the ONS.

Second, I match county-level house price information to individual data from the British Household Panel Survey (BHPS)[2]. The BHPS is a panel survey covering the period from 1991 to 2008. For each year it covers a sample of approximately 5500 households and 10,000 individuals. Previous literature has used the same data to explore whether house price shocks improve the health status of individual (Fichera and Gathergood, 2016) or change their labor

[2]Due to some data limitations the geographical match of house price data (based on 2011 administrative boundaries) to individual data in the BHPS (based on 1991 administrative boundaries) is not perfect. Thus, the housing price data used in this analysis covers fifty-three counties in England and Wales.

supply (Disney and Gathergood, 2018). The main advantage of the dataset is its multi-purpose design. This allows me to have individual-level information on values, attitudes and political preferences, as well as information on home-ownership status and income.

The sample used in the analysis is restricted to household heads, being the person financially responsible for the household, aged between 15 and 80. The sample only includes individuals that state to be either owners or renters, dropping around two percent of the original sample who reports to live in the house rent-free or to share the accommodation. In order to match this data to the house price index, the original BHPS sample is limited to to head of households living in England and Wales in the period 1995-2007.

Additionally, the main sample does not include individuals switching between the renter and homeowner status between each wave. This is done in order to avoid biased estimates, especially in terms of redistribution preferences of owners downgrading to renting. The estimates presented in this paper might also be affected by individuals moving across counties as a response to house price shocks or as a response to increasing support of the Conservative party. In the robustness checks the estimation excludes individuals moving across counties from one year to the next.

The BHPS provides information on all the dependent variables of interest. In particular, I use two variable in the database to measure political support for the Conservative party. The first one is a binary variable equal to one if the individual identifies the Conservative party as the party supported, and the second one equals one if the individual reports to feel closer to that party compared to any other party. In an alternative specification, I measure party support as the probability of voting for said party. To do so, I use a binary variable equal to one if the individual reports to have voted the Conservative party in the general elections of 1997, 2001 or 2005.

To measure a change in preferences I use data on political and social beliefs. In particular, I measure pro-redistribution attitudes creating a binary variable equal to one if the individual reports to agree or strongly agree with the following statement "Major public services ought to be in state ownership", and zero otherwise. Other studies have used similar questions to measure pro-redistribution attitudes (Clark et al., 2010; Ashok et al., 2015)[3] as it is clearly

[3]Barth and Moene (2016) refer to a decrease in individual preferences for welfare spending which correlates

related to one's preferences over the public provision of services. I use similar binary variables to measure other attitudes related to the role of the government in the economy, using three more statements on values included in the BHPS: "Private enterprise is the best way to solve Britain's economic problems", "The government ought to impose a maximum level of money one can make" and "It is the government's responsibility to provide a job for everyone who wants one". For all these statements the individual can choose among strongly agree, agree, neither agree or disagree, disagree and strongly disagree.

1.5 Identification strategy

I am interested in the effect of wealth inequality on political attitudes. Before explaining the main challenges to the identification of shocks to the wealth distribution, I motivate the use of housing wealth and house price data to measure wealth inequality.

First, I use housing wealth as a proxy for individual total wealth. This is a credible approximation in the context of England and Wales, where housing constitutes the larger asset in individual's portfolio and has often been used in the literature to measure individual total wealth (Campbell and Cocco, 2007; Disney et al., 2010; Disney and Gathergood, 2018). By focusing on housing wealth, I am able to identify an increase in inequality between two well-defined groups of the population: homeowners and renters.

Second, I follow the literature in using average county-level house price data to measure housing wealth. Several studies use this approach and use county-level house price data in England and Wales for the period between 1995 and 2007 in their analysis (Disney et al., 2010; Fichera and Gathergood, 2016; Disney and Gathergood, 2018). They do this for two reasons. On one hand, this period of analysis saw average house prices more than doubling in these two countries (Fichera and Gathergood, 2016). On the other, geographical and temporal variation in prices at the county-level allows to study a more dis-aggregated effect than previous studies using regional data. This allows to control for similar geographical characteristics affecting homeowners and renters. This allows me to use house prices to measure rising wealth inequality as an increase in house prices will increase the value of one's property, widening the gap between

with voting for right-wing parties. In their work the level of welfare spending, measured at the national level using the welfare generosity index (Scruggs et al., 2014), implemented by right-wing governments is then the outcome of these change in individual preferences.

homeowners and renters, i.e. our measure of wealth inequality.

Although including county and year fixed effects will capture time and space heterogeneity and looking at the effect on homeowners relative to renters will control for local labor market conditions (Fichera and Gathergood, 2016), identifying wealth inequality using only average local house prices would produce biased estimates. This is because of unobserved heterogeneity that might affect both local house prices and political preferences, e.g. quality of public goods (Gibbons et al., 2013), as well as issues of reverse causality. To overcome such limitations, I identify an exogenous shock to the distribution of housing wealth to evaluate the causal effect of a change in wealth inequality on political preferences.

Following Disney et al. (2010) and Disney and Gathergood (2018), I estimate a year-on-year shock to house prices as a proxy for wealth shocks. The intuition behind this strategy is to use only the unanticipated change in local average house prices as a measure of exogenous changes to the distribution of housing wealth. This unanticipated change, or surprise, will be uncorrelated with other local level characteristics once year and county fixed effects are included in the estimation. Thus, house price shocks represent an exogenous source of variation that affects differently homeowners and renters living in the same county and experiencing similar geographical conditions, but differing in terms of housing wealth endowment (Disney et al., 2010).

This exogenous shock is estimated using the residuals of an auto-regressive model of order two on average house prices using county-level fixed effects. Including county fixed effects allows to isolate the effect of local unobservable characteristics. This is estimated, as follows:

$$p_{t,c} = \alpha + \beta_1 p_{t-1,c} + \beta_2 p_{t-2,c} + County_c + \varepsilon_{t,c} \tag{1}$$

$$\varepsilon_{t,c} = \Delta(p_{t,c} - p_{t-1,c}) - E[\Delta(p_{t,c} - p_{t-1,c})] = PriceShock_{t,c} \tag{2}$$

Where p is the natural log of the average house price in county c at year t as defined by the ONS house price index. $PriceShock_{t,c}$ is the predicted unanticipated change in average house prices at the county level for each year t of the period of analysis. Table 1 shows the results of the auto-regressive model with county fixed effects. The R-squared in the model is 98 percent. Hence 98 percent of the variation in prices is explained by the auto-regressive structure. The shocks are identified off the remaining 2 percent.

Figure 8, in Appendix 1, aims at providing graphical and intuitive evidence of the price shocks. It shows the deviation from the predicted trend of average prices at the national level. In the next section I discuss how I use the same yearly shocks calculated at the local level to estimate the effect on political preferences. Following from the discussion presented in this section, using the exogenous house price shocks will provide causal evidence of the effect of an increase in wealth inequality on the outcomes of interest.

Table 1: Auto-regressive model of order two of county level average house prices in England and Wales between 1995 and 2007.

	(1) Avg Price
Avg Price (t-1)	1.683***
	(0.0205)
Avg Price (t-2)	-0.756***
	(0.0223)
Constant	0.846***
	(0.0217)
CountyFE	Y
N	759
Clusters	69
R-squared	0.981

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Clustered Standard Errors in parenthesis

Note: All models estimated using OLS with county fixed effects for 69 counties in England and Wales and for 12 years for the period 1995-2007. House price data are from the house price index published monthly by the ONS expressing the average price of all property types. Variable Price all is the natural log of the CPI-deflated average price at county level of all house types.

1.6 Empirical model

The empirical analysis presented in this paper looks at the effect of wealth inequality on two different outcomes: preferences for redistribution and support for the Conservative party.

In the previous section, I have discussed how house price shocks can be used to identify an exogenous change to the housing wealth distribution. Following from that discussion, this section explains the empirical models used to evaluate the effect of these price shocks on the two outcomes of interest.

1.6.1 Political party

Firstly, I estimate the effect of house price shocks on individual support for the Conservative party, using the following specification:

$$y_{i,c,t} = \alpha + \beta_1 PriceShock_{c,t} + \beta_2 Owner + \beta_3 Price\hat{S}hock_{c,t} * Owner + \beta_4 Z_i + \gamma_c + \delta_t + \varepsilon_{ict} \quad (3)$$

where $y_{i,c,t}$ is a binary variable equal to one if the individual states to support the Conservative party ($SupportTory$). I use an alternative measure to capture support for the Conservative party, this is a binary variable equal to one if the individual states to feel closer to the Conservative party relative to all other parties ($CloserConservative$).

Owner is a binary variable equal to one if the individual is a homeowner and equal to zero if the individual is a renter. Thus, the coefficient $\beta_3$will give the effect of an unexpected change in average prices at the county level ($PriceShock_{c,t}$), estimated following the model in section 1.5, for homeowners relative to renters. The net effect for owner will be given by $\beta_1 + \beta_3$. Z_i is a vector of individual characteristics, in line with what is done in the literature on preferences for redistribution (Giuliano and Spilimbergo, 2014): marital status, employment status, type of employment, gender, highest educational degree, and the log of real equivalized net household income.

The same model is estimated to predict the probability of having voted for the Conservative party in the previous year's general election. The following model is estimated only for the three years post election (1998, 2002, 2006) using the lag of the house price shock as follows:

$$\begin{aligned} y_{i,c,t-1} = \alpha + \beta_1 PriceShock_{c,t-1} + \beta_2 Owner_{i,c,t-1} + \beta_3 Price\hat{S}hock_{c,t-1} * Owner_{i,c,t-1} + \quad (4) \\ + \beta_4 Z_i + \gamma_c + \delta_t + \varepsilon_{ict} \end{aligned}$$

where $y_{i,c,t-1,}$ is the probability of having voted for the Conservative party in the previous year's general election. $Price\hat{S}hock_{c,t-1}$ is the house price experienced at the county level in the year of the general election, and all other variables are defined as in equation 3.

Estimating equation 4 allows me to check whether self-reported support for the Conservative party is accompanied by an actual vote during elections. Although the correlation coefficient between these two measures is 84 percent, I still perform this test to provide evidence of the robustness of the results.

The next step is to evaluate whether the change in political party supported is accompanied by a change in preferences for redistribution. The next section discusses the empirical model used in the estimation of the latter.

1.6.2 Preferences for redistribution

I estimate a similar linear probability model to evaluate the effect of house price shocks on preferences for redistribution. This second part of the estimation relies on the same empirical specification used to estimate party support, but it is estimated separately using different outcome variables for each of the values analyzed. In particular, I estimate:

$$y_{i,c,t,} = \alpha + \beta_1 Price\hat{S}hock_{c,t} + \beta_2 Owner + \beta_3 Price\hat{S}hock_{c,t} * Owner + \beta_4 Z_i + \gamma_c + \delta_t + \varepsilon_{i,c,t} \quad (5)$$

Where $y_{i,c,t}$ is a binary variable indicating if the respondent agrees or strongly agrees with one of following the statement:

1. Private enterprise is the best way to solve Britain's problem ($Private$)
2. Major public services and industries ought to be in state ownership ($Public$)
3. The government should provide a job for anyone who wants one ($GovJob$)
4. The government ought to impose a maximum level of money one can make ($MaxMoney$).

Statements 1, 2 and 3 are included in the BHPS for the years 1995, 1997 and 2004, while question 4 was asked in 1996, 1998 and 2003. β_3 captures the net effect of a house price shock for homeowners relative to renters. In the main analysis, I will focus on the change in preferences for the state ownership of public services ($Public$). Following from previous literature (Clark et al., 2010), this constitutes a plausible measure of preferences for a larger role of the government in public service provision and better captures preferences for redistribution through government spending.

I will provide evidence of the effect on the other preferences to check the consistency of the results. In particular, following the hypothesis tested, one would expect that rising wealth inequality will make homeowners more likely to agree with statement 1 ($Private$), and less likely to agree with statement 3 ($GovJob$) and 4 ($MaxMoney$). All the estimations include individual controls, as well as year, age and county dummies. Standard errors are bootstrapped with 200 repetitions.

1.7 Summary Statistics

This section provides information on the context of England and Wales between 1995 and 2007 and on the characteristics of the sample used in the analysis.

First, this section documents the increase in house prices and its variation by region and property type. The period between 1995 and 2007 was characterized by a unprecedented increase in house prices in England and Wales. Figure 1 shows the increase in average price broken down at the regional level. Moreover, this increase interested all types of properties (figure 2 in Appendix 1), with the average price of a detached houses being the highest and doubling over the period 1995-2007. I will use the ranking of properties by average price in figure 2 to evaluate the impact of house price shocks for richer homeowners relative to the poorest (i.e. owners of flats).

Figure 2: Average house prices by property type in England and Wales for the period between 1995 and 2007.

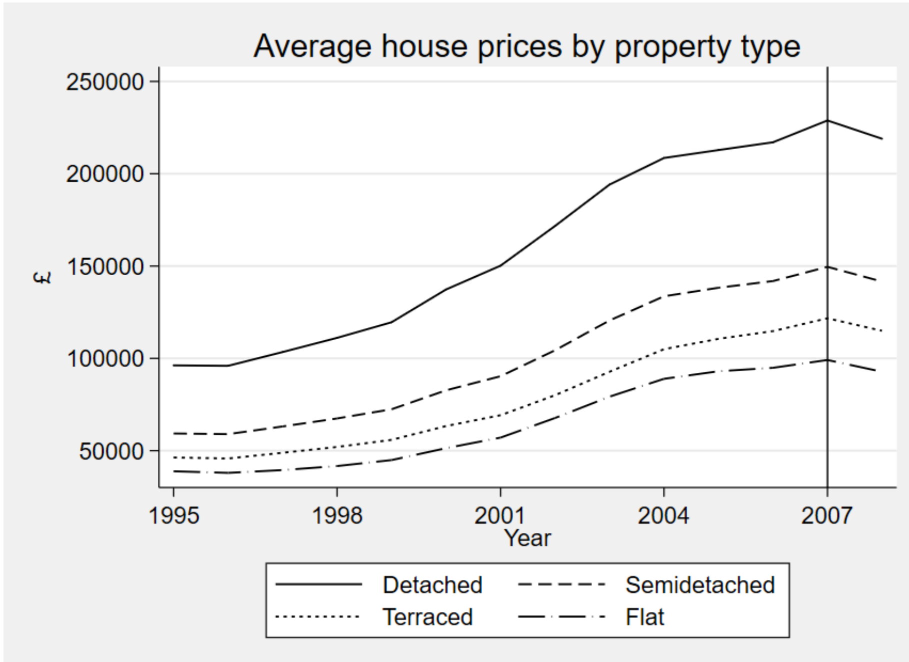

Second, table 2 shows the summary statistics for the sample used in the empirical analysis. The homeownership rate in the country reached a peak of 72 percent, showing that changes in housing wealth affect a big part of the British population. Among these, 2 percent became owners in each year and 7 percent became renters from one year to the next. These switchers between home ownership status are dropped from all the estimations[4]. Another consistent part of the sample moved across counties in each year. In the robustness checks these groups are excluded from the estimation. Table 2 shows that 62 percent of the owners in the sample own their house with mortgage.

Overall, table 2 shows that the two groups are fairly homogeneous in terms of household size and age. There are some consistent differences in terms of household composition with married or cohabiting couples (73 percent for owners and only 43 percent for renters). Then, home-owners are more educated, with three percent of the them having a higher education degree (versus one percent of renters), and are more likely to be male, 63 percent versus 47

[4]Notice that I only exclude switchers in home-ownership status between waves, rather than anyone who ever switched from being owner to renter in the whole period between 1995 to 2007. Thus, it is possible that one individual is dropped from the analysis in the year when she switched status, but then kept in the following year if the home-ownership status did not change again.

percent of renters. Renters are also more likely to be unemployed than owners.

Table 2: Summary statistics of the sample used in the analysis of individual head of household characteristics by homeownership status: BHPS data.

	Renter	Owner
Owner		0.72
Renter to owner		0.02
Owner to renter	0.07	
Moved across counties	0.16	0.07
Age	45	50
Male	0.46	0.63
Higher Degree	0.01	0.03
Mortgage		0.62
Couple	0.43	0.73
Household size	2.41	2.60
Unemployed	0.08	0.02
Retired	0.21	0.23
Self Employed	0.04	0.10
Employed	0.38	0.57
Outcomes		
Voted in Last General Election	0.65	0.80
Support Tories	0.14	0.30
Voted Tories	0.16	0.32
Closer Tories	0.17	0.37
Public Services	0.41	0.39
Private Enterprise	0.19	0.32
Max Wage	0.30	0.20
Govmt Job	0.56	0.38
Property types		
Detached	0.04	0.31
Semidetached	0.27	0.36
Flat	0.22	0.03
Terraced	0.34	0.26
Observations	18288	47487

Note: Summary statistics for the main sample of analysis. Sample includes head of households only. Data from the BHPS wave 5-17 (1995-2007). Binary variables "Owner to renter" and "Renter to owner" indicate individuals who switched between homeownership status during the period of analysis, these will be dropped in the main analysis. Binary variables "Public Services" "Private enterprise" "Max wage" "Govmt Job" equal 1 if the individual agrees or strongly agrees with statements on political and social beliefs. Variables "Detached" "Semidetached" "Flat" "Terraced" indicate self-reported type of house of residence from the BHPS.

Now, I provide information on the outcome variables of interest. First, we see that there are some differences in terms of voting turnout, with 80 percent of owners and only 65 percent of renters stating to have voted in the last general election. This might suggest that homeowners are more involved in civic life and have higher social capital (DiPasquale and Glaeser, 1999). Looking at the outcome variables of interest for this analysis, table 2 shows that there are some differences in terms of party support and voting behavior between homeowners and renters.

Homeowners are more likely to support the Conservative party (30 percent of homeowners support that party against only 14 percent of renters); to have voted for the Conservative party in the previous election (32 percent of homeowners); and to feel closer to that party (37 percent). Over time the support for the Conservative party has been stable during this period, but increased steadily after 2005.

Looking at the differences between the two groups in terms of political preferences and beliefs, table 2 shows that there are not consistent differences in terms of agreement over public ownership of services. Both group support this statement in 40 percent of the cases ($Public$). Clearer differences arise when looking at different types of beliefs. Owners are more likely to have pro-private enterprises attitudes ($Private$). Oppositely, renters agree more both with government provision of jobs ($Govjobs$) and with government intervention on the amount of money one can make ($Maxwage$).

Figure 9, in Appendix 1, shows the proportion of individuals agreeing with each of these statements by home ownership status. Overall, the analysis of summary statistics and correlations suggests that the hypothesis theorized in the literature holds in this specification (Barth and Moene, 2016). In fact, homeowners are more likely than renters to support the Conservative party and have less favorable preferences towards the government intervention in the economy. However, so far I am not able to see a clear cut difference between the two groups in terms of support for public services being in state ownership.

1.8 Results

Before proceeding to the results of the estimation of equations 3 to 5, this section presents the results of a naive regression estimating the correlation between county-level average house prices and the outcomes of interest. The aim of this initial exercise is twofold. First, I investigate further the correlations emerging from the summary statistics. Second, presenting this estimation shows whether the results change when using my preferred specification and house price shocks to estimate the causal effects of wealth inequality on party support and preferences.

Table 3 presents the results of a naive OLS regression of average house prices on political preferences. These correlations show mixed results. A small, positive and statistically significant correlation between the interaction term of home-ownership status and the probability of supporting or feeling closer to the Conservative party (columns 1 and 2). On average, there is no statistically significant correlation between the interaction term and political beliefs, with the only exception of a positive correlation with supporting the government provision of jobs.

Table 3: Regression analysis of the correlation between county-level average house prices, party support and political preferences for individuals living in England and Wales 1995-2007.

	(1) Support Conservative	(2) Closer Conservative	(3) Public	(4) Govmt Job	(5) MaxWage	(6) Private
Avg Price	-0.027	0.018	-0.053	-0.032	-0.046	0.003
	(0.027)	(0.034)	(0.056)	(0.054)	(0.040)	(0.050)
Avg Price x Owner	0.027***	0.028**	0.017	0.028*	-0.001	0.018
	(0.009)	(0.011)	(0.017)	(0.017)	(0.015)	(0.014)
Owner	-0.181*	-0.169	-0.224	-0.398**	-0.043	-0.129
	(0.097)	(0.127)	(0.194)	(0.187)	(0.174)	(0.163)
Couple	-0.001	-0.005	0.008	0.014	0.009	-0.014
	(0.005)	(0.007)	(0.010)	(0.010)	(0.008)	(0.009)
Household income	0.049***	0.056***	-0.029***	-0.057***	-0.038***	0.049***
	(0.004)	(0.005)	(0.007)	(0.007)	(0.005)	(0.006)
Higher Degree	-0.043***	-0.095***	0.011	-0.271***	-0.061***	0.129***
	(0.012)	(0.015)	(0.024)	(0.021)	(0.018)	(0.022)
Undergraduate	0.002	-0.036***	0.008	-0.240***	-0.074***	0.125***
	(0.008)	(0.009)	(0.014)	(0.013)	(0.011)	(0.013)
Secondary Edu	0.079***	0.089***	-0.017*	-0.170***	-0.061***	0.056***
	(0.005)	(0.006)	(0.009)	(0.009)	(0.008)	(0.008)
Male	0.003	-0.002	0.070***	-0.055***	-0.008	0.149***
	(0.004)	(0.005)	(0.008)	(0.008)	(0.007)	(0.007)
Employed	0.003	0.009	-0.002	0.003	-0.008	-0.039***
	(0.006)	(0.008)	(0.012)	(0.012)	(0.011)	(0.010)
Household size	-0.001	-0.006**	-0.005	0.003	0.001	-0.001
	(0.002)	(0.003)	(0.004)	(0.004)	(0.003)	(0.003)
Self Employed	0.101***	0.135***	-0.071***	-0.098***	-0.066***	0.081***
	(0.009)	(0.012)	(0.017)	(0.016)	(0.013)	(0.015)
Unemployed	0.004	0.010	0.015	0.033	0.028	-0.029
	(0.010)	(0.015)	(0.024)	(0.022)	(0.021)	(0.019)
Retired	0.034***	0.043***	-0.010	-0.032*	-0.033**	0.028
	(0.010)	(0.012)	(0.018)	(0.017)	(0.016)	(0.017)
N	47015	33524	19001	19537	19634	18824

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust Standard Errors in parenthesis

Note: All models estimated using OLS. All outcome variables are binary variables equal to 1 if : (1) the individual supports the Conservative party, (2) the individual supports the Labour party, (3) the individual agrees with state ownership of public services, (4) the individual agrees with the government providing jobs, (5) the individual agrees with the government imposing a maximum wage, (6) the individual agrees with a stronger role of the private sector in the economy. All types is a continuous variable indicating the natural log of the average price at the county level for 53 counties in England and Wales between 1995 and 2007.Owner is a binary variable equal to 1 if the individual is a homeowner and equal to zero if she is a renter. The regressions are estimated using a linear probability model. Each regression includes controls for marital status, education level, household income, gender, employment status, household size and binary variables for county, age and survey year.

These correlations change when estimating the causal effect of house price shocks on political preferences, suggesting that the estimates of the previous specifications might be biased.

Table 4 shows the main results of this analysis. Homeowners are on average 1.2 percentage points more likely than renters to support the Conservative party. This effect is larger and statistically significant when I look at the coefficient of interest, i.e. the interaction between home-ownership status and house price shocks. The results show that homeowners experiencing a ten percent unexpected increase in average prices at the county level become on average 2.1 percentage points more likely to support the Conservative party (and 2.4 percentage points more likely to feel closer to it). The net effect for homeowners compared to renters is of an increase of 1 percentage points in the probability of supporting the Conservative party. Evaluating the magnitude of this effects at the average, a total house price shock of 10 percent increases the probability of supporting the Conservative party of approximately 2.8 percent for homeowners.

In line with this, column 2 in table 4 shows a similar 1 percentage point effect when using an alternative measure of support for the Conservative party compared to renters experiencing house price shocks in the same county. Column 3 shows that the support for the Labour party decreases, in line with the main finding for the Conservative party, but in this case the coefficient is not statistically different from zero. Notice that the baseline category is supporting all other parties in the U.K. including the Conservative party.This result corroborates the hypothesis suggesting that richer groups of the population becoming *richer*, as a results of house price shocks, are more likely to support more right-wing leaning party, when using the definition for the Conservative party used in Barth and Moene (2016).

Table 4 includes all the controls used in the estimation. Note that, household income is positively associated with voting for the Conservative party, and more educated individuals are less likely to vote for the Conservative party.

Table 4: The effect of wealth inequality on support for the Conservative party and preferences for government ownership of public sector in England and Wales 1995-2007: main results.

	(1) Support Conservative	(2) Closer Conservative	(3) Support Labour	(4) Public
Price shock	-0.109	-0.155	0.150	-0.116
	(0.079)	(0.097)	(0.120)	(0.162)
Price shock x Owner	0.213***	0.240**	-0.094	-0.030
	(0.072)	(0.097)	(0.113)	(0.165)
Owner	0.120***	0.145***	-0.069***	-0.032***
	(0.005)	(0.006)	(0.007)	(0.012)
Couple	0.004	-0.000	0.000	0.006
	(0.006)	(0.007)	(0.007)	(0.012)
Household income	0.046***	0.053***	-0.014***	-0.028***
	(0.004)	(0.005)	(0.004)	(0.008)
Higher Degree	-0.041***	-0.095***	-0.064***	0.006
	(0.013)	(0.014)	(0.016)	(0.027)
Undergraduate	0.004	-0.036***	-0.039***	-0.001
	(0.008)	(0.010)	(0.008)	(0.015)
Secondary Edu	0.076***	0.087***	-0.096***	-0.018**
	(0.006)	(0.007)	(0.006)	(0.008)
Male	0.003	0.001	0.017***	0.075***
	(0.004)	(0.006)	(0.005)	(0.008)
Employed	-0.002	-0.001	0.004	0.010
	(0.006)	(0.009)	(0.009)	(0.013)
Household size	-0.001	-0.006**	0.012***	-0.006
	(0.002)	(0.003)	(0.003)	(0.004)
Self Employed	0.091***	0.121***	-0.076***	-0.054***
	(0.010)	(0.014)	(0.011)	(0.018)
Unemployed	0.008	0.008	-0.009	0.019
	(0.011)	(0.016)	(0.015)	(0.029)
Retired	0.030**	0.037***	-0.008	0.001
	(0.012)	(0.013)	(0.012)	(0.019)
Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes
N	40708	28679	40708	15793

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Bootstrapped Standard Errors in parenthesis

Note: All models estimated using OLS. Standard errors are bootstrapped (200 reps).The dependent variables are binary variables equal to 1 if :(1) the individual supports the Conservative Party, (2) the individual feels closer to the Conservative party, (3) the individual supports the Labour party, (4) the individual agrees with state ownership of public services. "Price shock all" is estimated as the residuals of an autoregressive model of order 2 on average house prices at the county level for 53 counties in England and Wales in the period 1995-2007. Owner is a binary variable equal to 1 if the individual is a homeowner and zero if a renter. Sample includes head of households only. Data from the BHPS wave 5-17 (1995-2007) and from the ONS-HPI. The estimation includes a set of controls: binary variables for marital status, gender, education level, retirement status, employment status, log hh equivalised income, county, year of survey and age.

I corroborate these findings further by estimating the model presented in equation 4. This exploits the information on the party for which individuals voted in the last general election.

These results are in line with what presented above. In particular, table 5 shows that individuals experiencing a house price shock in the year of a general election are more likely to have voted for the Conservative party and less like to have voted for Labour. The sign of the coefficient for the interaction term does not change if using a probit model to estimate the average marginal effects following the same equation. However, the effect becomes smaller and not statistically significant.

Table 5: The effect of wealth inequality on the probability of voting for the Conservative and Labour party in England and Wales 1995-2007: post-election years only

	(1)	(2)	(3)	(4)
	OLS		Probit	
	Voted Conservative	Voted Labour	Voted Conservative	Voted Labour
Price shock	-0.503**	0.319	-0.323	0.232
	(0.252)	(0.286)	(0.375)	(0.396)
Price shock x Owner	0.497**	-0.449*	0.215	-0.382
	(0.233)	(0.236)	(0.325)	(0.397)
Owner	0.135***	-0.130***	0.152***	-0.144***
	(0.014)	(0.013)	(0.016)	(0.016)
N	8499	8499	8499	8499

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Bootstrapped Standard Errors in parenthesis

Note:.Standard errors are bootstrapped (200 reps). Columns 1 and 2 are estimated using a linear probability model, while columns 3 and 4 are estimated using probit regressions and evaluating the effect at the margin. The dependent variables are binary variables equal to 1 if :(1) (3) the individual voted for the Conservative Party, (2) (4) the individual voted for the Labour party. "Price shock all(t-1)" is estimated as the residuals of an autoregressive model of order 2 on average house prices at the county level for 53 counties in England and Wales in the period 1995-2007. t-1 refers to the period before the individual is asked about the voting casted in the last general election. Owner is a binary variable equal to 1 if the individual is a homeowner and zero if a renter. Sample includes head of households only. Data from the BHPS wave 5-17 (1995-2007) and from the ONS-HPI. The estimation includes a set of controls: binary variables for marital status, gender, education level, retirement status, employment status, log hh equivalised income, county, year of survey and age.

These results show that the first part of the hypothesis being tested is corroborated in my empirical analysis. In fact, homeowners experiencing an increase in housing wealth become more likely to support the Conservative party relative to renters.

The next step is to investigate whether this increase in support for the Conservative party support is accompanied by a change preferences for redistribution. The main results are reported in the last column of table 4. They shows that homeowners do not become less likely to support state ownership of public services. In fact, the coefficient for the interaction term of price shocks and home-ownership status is negative, but not statistically different from zero. On average, for a ten percent house price positive shock homeowners are 0.3 percentage points less likely to have pro-public attitudes relative to renters. However, the net effect of a house price shock for homeowners, of 0.6 percentage points, is not statistically different from zero.

I test whether this effect changes when looking at the other variables capturing political beliefs in table 6. None of the coefficients for the interaction term are statistically different from zero. In line with the literature, homeowners are more likely to have favorable views towards the role of private enterprises in the economy (column 1), and less likely to have favorable views towards the government provision of jobs and regulation of wages (columns 2 and 3).

Table 6: The effect of wealth inequality on redistributive attitudes in England and Wales 1995-2007: government's ownership of public services, role of government and private sector in the economy.

	(1) Private Ent	(2) Max wage	(3) Govmt Job
Price shock	-0.297**	0.096	-0.266*
	(0.133)	(0.156)	(0.157)
Price shock x Owner	0.146	0.024	0.217
	(0.154)	(0.165)	(0.159)
Owner	0.061***	-0.057***	-0.086***
	(0.009)	(0.009)	(0.011)
Controls	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Age FE	Yes	Yes	Yes
County FE	Yes	Yes	Yes
Observations	15640	16196	16244

Note: All models estimated using OLS. SE bootstrapped(200 reps). The outcome variables are binary variables:(1) equal to 1 if individual agrees with private ownership of public services, (2) if agrees with government ought to cap maximum amount one can earn, (3) if agrees with government having to provide jobs. Price shock is estimated as the residuals of an autoregressive model of order 2 at the county level for England and Wales in the period 1995-2007. Owner is a binary variable equal to 1 if the individual is a homeowner and zero if a renter. Sample includes head of households only. Data from the BHPS wave 5-17 (1995-2007) and from the ONS-HPI. All regressions includes a set of controls: binary variables for marital status, gender, education level, retirement status, employment status, ln of household equivalised income, county, age and survey year).

1.9 Mechanisms

I explore two mechanisms that might explain the results presented so far. The first one is that homeowners might have heterogeneous responses to a local house price shock. This can be due to the fact that richer homeowners will react differently from poorer homeowners. Thus, the average insignificant effect on preferences for redistribution might hide important heterogeneous effects. I explore this possibility by using house price data for four different property types. The second possible explanation is that homeowners value their property differently from what captured with county-level house price data. I explore this using self-reported valuation of one's property.

1.9.1 Heterogeneity in property value

One possible reason driving the results presented in the previous section is the heterogeneity in political preferences among homeowners. Homeowners of different property types might be affected differently by an increase in prices. Thus, homeowners of cheaper properties, i.e. owners of flats, could benefit less from the increase in the value of their housing wealth, and support redistributive measures more than owners of other properties.

Figure 10, in Appendix 1, shows the probability of supporting the Conservative party by type of the property owned. On average, owners of detached houses (43 percent) are more likely to support the Conservative party than owners of other properties (27 percent semidetached, 21 percent terraced, 33 percent flats). Figure 10.a shows that the increase in support for the Conservative party grew steadily for owners of detached houses starting from 2000. This descriptive evidence is confirming an heterogeneity among homeowners in terms of support for the Conservative party. Owners of detached properties agree with State ownership of public services in 34 percent of the cases, owners of semi-detached properties in 40 percent, owners of terraced houses in 43 percent and owners of flats in 39 percent. This suggest that owners of more expensive properties are less likely to favor redistribution compared to owners of less expensive properties.

Figure 11, in Appendix 1, shows the probability of having pro-public sector attitudes by type of property owned during the period of interest. This shows that the trend in terms of support for the public sector has been decreasing for all types of homeowners in the period of analysis. I estimate whether this heterogeneity is confirmed when using a regression analysis. I restrict the sample of observations to homeowners only, and use a simple linear probability model to estimate the correlation between increase in average prices of different properties and the change in probability of supporting the Conservative party and having pro-public sector preferences. I do this estimating the following model:

$$y_{i,c,t} = \alpha + \beta_1 pricedetached + \beta_2 pricesemidetached_{c,t} + \beta_3 priceterraced_{c.t} + \beta_4 Z_i + \gamma_c + \delta_t + \varepsilon_{i,c,t} \tag{6}$$

Then I use my preferred specification to estimate the causal effect of a house price shock at the county level[5] on the change in probability of supporting the Conservative party and having pro-public sector attitudes for homeowners of different properties.

Following equation 6, I estimate the causal effect of an exogenous increase in house prices at the county level on the political preferences of homeowners for four different property types.

$$y_{i,c,t} = \alpha + \beta_1 shockdetached_{c,t} + \beta_2 shocksemidetached_{c,t} + \beta_3 shockterraced_{c.t} + \beta_4 Z_i + \gamma_c + \delta_t + \varepsilon_{i,c,t} \tag{7}$$

Where $y_{i,c,t}$ is a binary variable equal to one if the individual supports the Conservative party (Support Conservative) or if agrees with state ownership of public services (*Public*), as in the main analysis. I use as a reference category for this estimation homeowners of flats experiencing an unexpected increase in house prices in their county. β_1 measures the effect of a house prices shock for owners of detached properties, β_2 measures the effect for owners of semi-detached houses, and β_3 the effect for owners of terraced houses relative to the baseline category.

The regression includes the same individual controls used in the rest of the analysis, as well as age, year and county fixed effects.

[5] Notice that the house price shock is estimated using an auto-regressive model of order two as in section 3. In this specification, however, the autoregressive model is estimated separately for each property type. In fact, the ONS house price index provides information on house prices by property types. I then merge the house price shocks by property type and county to the individual level data in the BHPS. The BHPS contains information on type of property where each respondents lives. I categorize as detached houses those who report detached house (31.45 percent), as semidetached those who report semidetached or bungalow (36.52 percent, as terraced house those who report terraced and end-terrace (26.35 percent), as flats those who report purpose built flat, converted flat, bedsit (5 percent).

Table 7: Heterogeneity analysis of the effect of wealth inequality on political preferences in England and Wales 1995-2007: homeowners only by type of property owned.

	Avg Price		Price Shocks	
	(1)	(2)	(3)	(4)
	Support Conservative	Public	Support Conservative	Public
Price Detached	0.103***	-0.053		
	(0.020)	(0.033)		
Price Semidetached	0.090***	-0.034		
	(0.019)	(0.032)		
Price Terraced	0.063***	-0.078**		
	(0.018)	(0.031)		
Detached	-1.086***	0.599	0.140***	-0.014
	(0.234)	(0.393)	(0.013)	(0.022)
Semidetached	-1.024***	0.413	0.006	0.048**
	(0.221)	(0.362)	(0.011)	(0.022)
Terraced	-0.745***	0.920***	-0.040***	0.072***
	(0.200)	(0.352)	(0.011)	(0.021)
Shock Detached			0.218**	-0.114
			(0.098)	(0.202)
Shock Semidetached			0.185**	-0.208
			(0.083)	(0.180)
Shock Terraced			0.034	-0.396**
			(0.089)	(0.181)
Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes
N	34154	13834	29689	11504

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Bootstrapped Standard Errors in parenthesis

Note: Bootstrapped standard errors in parenthesis (200 repetitions). All outcome variables are binary variables equal to 1 if : (1) (3) the individual supports the Conservative party,(2) (4) the individual agrees with state ownership of public services. Price detached, price semidetached, price terraced are continuous variable indicating the natural log of the average price at the county level for 53 counties in England and Wales, by property type.Detached, semidetached, terraced are binary variables equal to 1 if the individual is a homeowner of each different property (baseline category: flat owners). Shock detached, shock semidetached, shock terraced are continuous variables measuring the shock in house prices for each property type.The regressions are estimated using a linear probability model. Each regression includes controls for marital status, education level, household income, gender, employment status, household size and binary variables for county, age and survey year.

Table 7 shows the results. Columns 1 and 2 show the results of a OLS regression of average county-level house prices and political preferences, while columns 3 and 4 show the results for house price shocks. Fist, I find that average house prices are positively correlated with support for the Conservative party. The results showed in column 1 of table 7 show that homeowners of more expensive type of properties are progressively more likely to support the Conservative party than homeowners of flats. In particular, a one percent increase in house prices increases the probability of supporting the Conservative party by 0.1 percentage points for owners of detached houses relative to owners of flats. This correlation becomes smaller in magnitude for owners of semidetached properties (0.09 percentage points) and terraced houses (0.06 percentage points). Looking at column 2, I find a negative correlation between average house prices of more expensive properties and pro-public sector preferences relative to preferences of flat owners. The only statistically significant effect is for homeowners of terraced houses, becoming 0.08 percentage points less likely to favor redistribution than owners of flats.

The results of my preferred specification using exogenous price shocks are reported in columns 3 and 4 of table 7. These results confirm that owners of detached properties are 2 percentage points more likely to support the Conservative party as a result of a 10 percent positive shocks in local house prices relative to owners of flats. Similarly, the effect is of 1.8 percentage points for owners of semi-detached properties, while the effect is smaller for owners of terraced houses. This finding suggests that there is an heterogeneity in the support for the Conservative party among homeowners. In fact richer homeowners, i.e. owners of more expensive properties, are more likely to support the Conservative party compared to poorer ones, i.e. owners of cheaper properties, who experience a positive house price shock at the county level.

Important heterogeneity can be found when looking at preferences for the state ownership of public services. Column 4 shows that homeowners of all properties are less likely to favor redistribution compared to owners of flats. However, the effect becomes significant, in magnitude and statistically, for owners of terraced houses. In particular, owners of terraced houses are 4 percentage points less likely than owners of flats to agree with the state ownership of public services as a consequence of a 10 percent positive shock to house prices, while the reduction is smaller and not statistically significant for owners of semidetached (-2 percentage points) and detached properties (-1.1percentage points). These results show that there is heterogeneity across homeowners in terms of preferences for redistribution, but this is statistically significant only for owners of slightly more expensive properties. This might suggest that the change in preferences for redistribution interests homeowners at the bottom of the property value distribution and not home-owners of semi-detached and detached houses. A possible explanation for this result is that owners of more expensive properties are less dependent on public provision of services and a shock to their property's value might have a lower impact compared to home-owners starting from a lower starting point (in the distribution of property value).

To conclude, I find evidence in line with the theory tested when exploring heterogeneity among homeowners, by exploiting variation in the value of property types. I find that political preferences of richer homeowners are affected differently than those of poorer homeowners. In particular, the support for the Conservative party becomes stronger for homeowners of more expensive properties, corroborating the hypothesis tested. Moreover, once these heterogeneity are explored, I find evidence in support of the hypothesis that richer individuals are less likely to support redistribution. In fact, I find that homeowners of cheaper properties, here of flats, are less likely to support the Conservative party and are more likely to agree with state ownership of public services than homeowners of more expensive properties.

1.9.2 Endowment effects

A second mechanism explaining why the main results do not corroborate the predictions made in the literature is that homeowners perceive the value of their housing wealth to be different than the one it has on the market. In this section, I show that when using self-reported value of one's property of residence as a measure of housing wealth, the results show a statistically significant decrease in the probability of having pro-redistribution attitudes. A plausible explanation might be the existence of an endowment effect for homeowners.

Figure 12, in Appendix 1, shows that on average homeowners overestimate the value of their property relative to the house prices provided by the ONS house price index. It is also likely that one's perceptions on own house value affect political beliefs differently from official house price measures. Thus, homeowners might perceive to be richer than they are, and change their political preferences accordingly. This hypothesis is in line with the literature on endowment effects and on perceptions of inequality.

To test this mechanism I exploit the following question asked yearly in the BHPS only to homeowners: "About how much would you expect to get for your home if you sold it today?". The results following in this section present estimation of the following equation:

$$y_{i,c,t} = \alpha + \beta hsval_{i,c,t} + \beta_4 X_{i,c,t} + \gamma_c + \delta_t + \varepsilon_{i,c,t} \tag{8}$$

where $y_{i,c,t}$ is a binary variable equal to one if the individual i in county c in year t supports the Conservative party or if agrees or strongly agrees to each of the statements on values used in the main analysis, $hsval$ is the log of self-reported house value, X includes a series of individual characteristics as in previous estimations. β is the coefficient of interest and measures the effect of a 1 percent increase in house value on the outcome of interest. All the regressions are estimated using OLS, standard errors are clustered at county level and include year, age and county fixed-effects.

Table 8: Correlation between self-reported house value and political preferences: homeowners only

	(1) Support Conservative	(2) Public	(3) Govmt Job	(4) Max Wage	(5) Private
House Value	0.179***	-0.0835***	-0.104***	-0.0703***	0.116***
	(13.26)	(-6.75)	(-9.14)	(-6.35)	(8.96)
Controls	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes
N	37303	10778	10970	11173	10717

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Clustered Standard Errors in parenthesis

Note: All models estimated using OLS. SE clustered at county level The outcome variable is a binary variable equal to 1 if: (1) the individual supports the Conservative Party, (2) equal to 1 if individual agrees with state ownership of public services,(3) if agrees with government provision of jobs (4) if agrees with government ought to cap maximum amount one can make, (5)private enterprises solving economic problems. Data from the BHPS wave 5-17 (1995-2008). The estimation includes a set of controls: binary variables for marital status, gender, education level, retirement status, employment status,log hh equivalised income. Binary variables for age, survey year and county.

Table 8 shows that homeowners reporting higher house values are more likely to support the Conservative party, in line with the main results. However, differently from what estimated previously, homeowners reporting higher house values are less likely to be in favor of the public provision of goods and of the government's intervention in the economy. In line with this, they are more likely to support private enterprises as being the main solution to economic problems. All of the coefficients of interest in table 8 are statistically significant and indicate a precisely estimated correlation between self-reported house value and preferences.

The contribution of this additional finding is twofold. First, it shows that the perception of wealth affects beliefs differently than market-value house price shocks. This contributes to the literature showing that perception of inequality and of the income distribution vary largely relative to aggregate measures (Gimpelson and Treisman, 2015). It also confirms large experimental evidence on endowment effects showing that individuals value a good more just because they own it.

Figure 13, in Appendix 1, corroborates further this idea by showing that homeowners are remarkably more likely than renters to think that owning a house is important. This finding suggests that perceived value of one's assets might be a relevant predictor of individuals beliefs and political preferences.

1.10 Robustness

I conduct several robustness checks to validate the results of the analysis.

First, I exclude from the sample individuals moving across counties between each period of time. This isolates possible endogeneity between the choice of new location, and house prices or political preferences. Results for the estimation excluding movers from the sample are reported in table 25 (in Appendix 1). These results validate the main findings.

Second, I perform the main analysis excluding the county of Greater London. I do this to check whether the results hold once the biggest outlier in terms of house prices is excluded from the analysis. The results in table 26, in Appendix 1, remain robust in their sign. However, the coefficient of the interaction term between house price shocks and homeownership status is only significant at the ten percent level, and the coefficient on feeling closer to the Conservative party becomes statistically insignificant. This is plausibly due to a lower variation in house price shocks away from the area of Greater London. It also suggests important geographical variation in the probability of supporting the Conservative party. This would be in line with the predicted hypothesis: homeowners in the area of Greater London are richer than homeowners in other regions. Hence, for a positive house price shock the probability of supporting the Conservative party should increase more.

I then estimate the results using data from 1995 to 2008, to include the year of the financial crisis. Table 27, in Appendix 1, shows that the results are in line with those presented in the main analysis. However, the magnitude of the coefficient for supporting the Conservative party is now smaller. A ten percent house price shock increases of 1.5 percentage points in the probability of supporting Conservative for homeowners relative to renters. The net effect is then smaller and of around 0.3 percentage points increase for owners. This suggests that the financial crisis might have an effect on other factors that influence the support for political parties.

Moreover, the financial crisis and the austerity measures implemented in the years right after 2008 have radically changed the political scenario in the United Kingdom (Fetzer, 2019). As a result of this change in the political scenario and the rise in support for the UKIP party, it is less clear whether looking at the Conservative party would be a convincing outcome to capture preferences less favorable towards redistribution. A possible extension of this analysis would examine how the support for parties changes as a result of the crisis. However, such analysis goes beyond the scope of this paper of looking at both support for parties and preferences. Moreover, this analysis cannot be implemented as the questions on preferences for redistribution used in this paper have been discontinued in the Understanding Society dataset covering the period after 2008 [6].

1.11 Conclusions

In this paper I estimate whether an increase in wealth inequality changes individual preferences for redistribution in line with one's support for political parties.

The findings of this paper contribute to a growing literature on the effect of inequality on political outcomes, innovating by focusing on wealth inequality, rather than income and wages. In particular, I test empirically whether increasing wealth inequality affects both the support for the Conservative party and pro-public sector attitudes, as predicted in the model on wage inequality formalized by Barth and Moene (2016).

In this analysis, I focus on housing wealth as a measure of individual total wealth. Previous literature has shown that this constitutes the largest share of individual asset in the context of analysis (Bastagli and Hills, 2012). Moreover, previous studies identify housing wealth effects

[6]Understanding Society continues to follow part of the BHPS sample together with a new sample of households entering the survey in 2009. Additional problems in extending the analysis are linked with the geographical matching of the house price index with individual level data.

to have different effects on homeowners relative to renters (Disney et al., 2010; Fichera and Gathergood, 2016; Disney and Gathergood, 2018). Given the relevance of the effects on labor supply decisions, consumption smoothing and health care demand, there is room to believe that these wealth effects also affect political preferences.

Using the house price boom in England and Wales between 1995 and 2007, I identify the causal effect of wealth inequality on political preferences, moving beyond the cross-country correlations presented in Barth and Moene (2016). Following Disney et al. (2010), I estimate an unanticipated change in house prices at the county-level to capture an exogenous shock to the housing wealth distribution. Then, I estimate whether it affects differently the political preferences of homeowners relative to renters. The results corroborate only partly the mechanisms of interest. In fact, homeowners become more likely than renters to support the Conservative party as their property unexpectedly increases in value. However, they do not change their support for government ownership of public services. These results hold after conducting a series of robustness checks.

I explore two possible mechanisms explaining these results.

First, I show that the effect on support for the public sector is driven by homeowners of cheaper properties. This is in line with the literature on inequality and preferences for redistribution, and shows that there is a large heterogeneity across homeowners preferences in the context of analysis. This is plausibly driven by the large heterogeneity in house value among homeowners. Second, I provide novel evidence on the relevance of one's own housing value perception on the formation of political preferences and on political party support. In particular, when using self-reported house value as a measure of wealth inequality, I find that the hypothesis tested (Barth and Moene, 2016) is fully corroborated. In fact, homeowners self-reporting higher values of own property are more likely to support the Conservative party and less likely to have pro-public sector preferences. This provides additional evidence of the presence of an endowment effect in the British housing market by showing that the valuation of homeowners is on average higher than the market value of houses, as well as showing that one's beliefs are affected by individual perceptions of socioeconomic status rather than aggregate measures.

Before concluding, I would like to draw the attention of the reader to possible limitations of this analysis. In particular, the measure of preferences for redistribution used in this paper might not fully capture one's redistributive attitudes. Measuring preferences for redistribution is a challenging task when the empirical method and data collection are not specifically targeted to understanding individual preferences over higher or lower levels of income and wealth redistribution. To overcome such limitation, several papers in the literature on preferences for redistribution use survey questions related to the government's intervention in the economy to capture one's attitude towards public spending and tax pressure. This choice lies on the assumption that as one prefers a stronger role of the government in the economy, especially in terms of assistance to poorer shares of the population, she also favours higher transfers of welfare from richer to poorer individuals via an increase in taxes for the former.

Alesina and Fuchs-Schündeln (2007) measure preferences for redistribution using survey questions on the role of the State in providing financial security to those in need (e.g. unemployed, poor etc.), Giuliano and Spilimbergo (2014)use questions on whether the State should assist the poor and on the role of hard work versus luck in determining one's economic opportunities. In line with what done in this paper, Clark et al. (2010) use the same questions from the BHPS to measure preferences for redistribution, i.e. major services and industries ought to be owned by the State and the government should put a limit to the maximum amount one can make. Combined with other questions on the role of private enterprises in solving economic problems and on the role of the government in providing a job to anyone who needs one, I believe that the set of questions used in the paper captures individual attitudes more or less in favour of redistribution.

However, I acknowledge that this measure can have some limitations. In particular, it would be ideal to have a measure of how much income and wealth one would be willing to pay in taxes to help poorer sectors of the population or to improve the quality of public services. In fact, it might well be that one favours redistribution, but opposes State ownership of public services if she associates this with mismanagement or misallocation of public funding resulting in poor-quality of the service provided. Although the other questions measuring one's preferences over the State's assistance to those who lose a job and role in capping wages help providing a clearer picture, having information that specifically measures individual preferences on tax increases[7] and wealth transfers would give a more comprehensive and precise idea on one's preferences over redistribution.

To conclude, this paper shows that wealth inequality has a role in affecting both political party support and preferences for redistribution. Moreover, it shows that heterogeneity in the value of one's owned wealth matters in explaining the effects, as it does one's perception of own property's value.

[7]For the purpose of this paper, information over preferences for changes in the council tax for more expensive properties would be particularly useful to measure one's preferences for, specifically, wealth redistribution. This information is not available in the BHPS, but future work might be carried out using information on changes to the council tax and political preferences.

2 The effect of conscription on political ideology, voting participation and national identity.

2.1 Introduction

This paper aims at evaluating the effect of conscription in shaping individual political ideology, voting participation and national identity.

The idea that conscription is associated with civic duty and national cohesion is not new. It was with this aim that conscription was introduced in France in 1798 (Jourdan law). Since then, conscription has been present in Europe and around the world in one form or another, until the end of the Cold War brought most European countries to move towards a fully-volunteer professional army. In recent years, however, an active policy discussion has culminated with reforms aimed at bringing back national mandatory service. In 2018, Sweden passed a law re-introducing conscription for eighteen years old men, and in the same year a law introducing a one-month long national service was approved by parliament in France.

The motivations behind such reforms are varied, from pure defense strategy to the promotion of civic engagement and national cohesion (this was the case in E. Macron presidential campaign, see: https://www.bbc.co.uk/news/world-europe-44625625). Aside from Emmanuel Macron, other policy-makers share the idea that re-introducing conscription will positively affect civic values. Matteo Salvini, Italy's former Interior Minister, has declared that conscription would be "good for democracy" and "educational" for young men and women (for M. Salvini's declaration see: http://www.ansa.it/english/news/politics/2018/02/07/military-service-should-return-salvini-2_39b7af17-1fa9-4a05-b2dd-24ba443769b5.html). Germany's current Ministry of Defense, Annegret Kramp-Karrenbauer (CDU), has been open to discussing the possibility of bringing back conscription in Germany seven years after the country decided to abolish it (for more information on the German political discussion, see https://www.telegraph.co.uk/new debates-return-national-service-amid-serious-military/). Additionally, the re-introduction of conscription is part of the electoral program of the far right-wing party Alternative für Deutschland.

Notwithstanding this renewed interest in the benefits of conscription, the costs associated with such policy are non-trivial.

First, a conscript-based army is less cost-efficient than a fully-professional army (Poutvaara and Wagener, 2007). Although the stipend paid to conscript is lower, having a conscript-based army still constitutes a considerable burden on the government's budget as it is the State who stipends conscripts, provides food, accommodation and transportation back home during service. Cutting down on these costs was one of the drivers of the abolition of conscription in Germany, for example, where it resulted in saving around 8.3bn euros (as reported: https://www.dw.com/en/german-military-cuts-to-put-effective-end-to-conscription/a-5909841). Second, there is a large economic literature assessing the direct and indirect costs of conscription. Military service delays entry in the labor market and reduces wages in the short-run (Angrist, 1990; Angrist and Krueger, 1994; Imbens and Klaauw, 1995), increases crime (Galiani et al., 2011) and has mixed effects on the demand for higher education (Cipollone and Rosolia, 2007; Card and Lemieux, 2001; Bauer et al., 2014). In light of such costs, a rigorous analysis of the effectiveness of conscription in shaping citizenship values is needed to motivate policies aimed at bringing conscription back. The main contribution of this paper is to provide such analysis and evidence on the causal effect of conscription on political preferences.

Conscription might affect political ideology, voting participation and national identity in opposite ways. By teaching discipline and obedience to duty it might increase voting participation, but its mandatory nature might reduce civic engagement if it generates antagonism against the State. Its strong hierarchical and authoritative structure might recall values embodied in more right-wing parties and affect political ideology (Roghmann and Sodeur, 1972). Lastly, by relying on a strong sense of identity with one's nation it might make individuals more likely to identify with their country (Akerlof and Kranton, 2005). Evaluating the direction of these mechanisms is an empirical question, which I address in this analysis using a regression discontinuity design and data for West Germany and Spain. The main result of this paper is that I do not find statistically significant evidence of an effect of conscription in shaping voting participation, political ideology nor national identity.

I start by evaluating the long-run effect of the introduction of conscription in West Germany on individual political preferences. Using data from the German Socio-Economic Panel, I estimate the effect of peacetime conscription on the probability of having more right-wing political ideology than the median, and on the probability of going to vote in parliamentary elections. Following previous literature (Bauer et al., 2012, Bauer et al., 2014), I exploit the introduction of conscription in 1955 to identify an exogenous change in the probability of being drafted for men turning eighteen in that year. Following from the de-militarization sanctions imposed to Germany after WWII, the probability of being conscripted for men turning eighteen before the reform is close to zero. Compared to men too young to have fought in the second World War and too old to be eligible for conscription, men born right after the reform are between 25 and 40 percentage points more likely to having been conscripted.

After showing that the treated and control groups are comparable, I estimate the treatment effect of the reform by looking at the difference in political attitudes between the two groups. The introduction of conscription coincided with several political transformations in the context of analysis: the admission of the Federal Republic of Germany in the NATO, for example, and other policies that where carried out to mark a political detachment from the past decades of Nazi dictatorship. Although, these structural transformations should not threaten the validity of the analysis as they are likely to affect both the treated and the control, this paper contributes further by conducting the analysis in a different context. As a robustness check, I evaluate the effect of conscription in Spain on the political preferences of men turning eighteen just before and just after 1997. Announced in 1996 by Prime Minister J.M.Aznár, this reform affected men born in or after 1979 and turning eighteen at the time of the reform, who were no longer eligible to be drafted. The advantages of conducting the analysis using Spanish data are threefold.

First, it allows me to show evidence of the external validity of the results, absence of which is often considered as one of the main limitations of using regression discontinuity design. Second, it allows me to study short-term effects of conscription as data is available for a period much closer to the implementation of the reform compared to the German case. Third, data availability on national identity allows me to explore a salient component of the current policy debate around the re-introduction of conscription. I find that there is no statistically significant evidence corroborating the hypothesis that conscription affects political ideology, voting participation or national identity. These results are robust to a series of specification and are comparable in sign and magnitude in the two countries of analysis.

This paper contributes to several branches of literature.

First, it contributes to a long standing literature in economics on the effects of wartime and peacetime conscription. Angrist (1990) seminal work first showed evidence of lower wages for Vietnam era draftees compared to non-draftees. Similar findings of lower earnings for conscripts have been backed up by analysis conducted in other countries, e.g. the Netherlands (Imbens and Klaauw, 1995). Later research on the effect on wages, shows mixed evidence. Updated results on Vietnam veterans find that earnings differentials converge to zero in the long run (Angrist et al., 2011), Albrecht et al. (1999) find positive effects of service on earnings in Sweden, while Grenet et al. (2011) find zero effect for British conscripts.

Moreover, the literature finds strong heterogeneity in the effects of peacetime conscription on earnings based on conscripts socio-economic background. Card and Cardoso (2012) find that in Portugal men with lower levels of education benefited from peacetime conscription, as their wages are 4 to 5 percent higher than comparable non-conscripts. This finding validates the hypothesis that conscription might have positive effects driven by the acquisition of new skills and the possibility to access a better social network during service.

The opposite is also possible. The literature on the effect of military service and crime shows that conscription during peacetime increases crime rates (Galiani et al., 2011). Hjalmarsson and Lindquist (2016) find similar results and strong heterogeneous effects, showing that conscripts from lower socio-economic backgrounds (proxied with father's education) are more likely to commit post-service crime, even more so if they had a criminal history before being drafted. These results validate the hypothesis that during conscription men might be exposed to the use of violence and weapons, as well as to worse social-networks.

Following from these results, the question of whether conscription might have positive effects on citizenship values is an empirical one. By teaching discipline and civic duty conscription might increase civic engagement and voting participation post-service. At the same time, it might also contribute to foster more authoritarian political views. Although the political science literature does not find evidence in support of this hypothesis (Campbell and McCormack, 1957; Roghmann and Sodeur, 1972), it has been hypothesized that the isolation from civil society and the strict disciplinary approach can contribute to shape political views (Akerlof and Kranton, 2005). Moreover, the fact that more right-wing parties support the re-introduction of conscription might suggest a right-wing partisanship of such reform.

Second, this paper contributes to the literature on the formation of political ideology, political preferences and their effect on voting behavior. This literature has shown that different institutions, economic growth and inequality affect the formation of political preferences (Alesina and Giuliano, 2009). In particular, Giuliano and Spilimbergo (2014) provide evidence that being exposed to negative shocks during the impressionable years affects the formation of preferences for redistribution more than when in other age periods, and that these effects last over time. The impressionable years hypothesis is a well-documented psychological theory (Krosnick and Alwin, 1989) stating that individuals form their political and social values in the period between 18 and 25 years of age. Being this the age period when individuals experience conscription, there is room to believe that if conscription has an effect in shaping political preferences, this effect will be long-lasting.

Additionally, Giuliano and Spilimbergo (2014) show that a change in preferences for redistribution is associated with changing support for a party or another and with changing political values as more or less liberal, and Caprettini et al. (2019) find evidence that individuals stick to their voting preferences for a period of up to four decades. Following a similar rationale Barth and Moene (2016) show that a change in beliefs affects the probability to vote for a more or less right-wing party. In light of the recent electoral success of populist/far-right/anti-immigrant parties, a growing literature has been moving beyond the analysis of preferences for redistribution and has started to investigate the determinants of voting for those parties (Acemoglu et al., 2013; Rodrik, 2018; Fetzer, 2019). This paper contributes to this literature by investigating whether conscription plays a role in shaping such preferences.

Lastly, this paper contributes to the literature on the formation of identity (Kranton, 2016 for a review). Stronger identity with one's group has been proved to be relevant for the efficient functioning of economic organizations (Akerlof and Kranton, 2005) as it increases the probability that an individual exerts effort on the workplace. This is particularly true in the military context, where an individual efficiency on the battlefield depends on how much he identifies with the organization he is fighting for, i.e. the State. Empirical evidence supports the relevance of identity formation. Clots-Figueras and Masella (2013) provide evidence of the effect of being taught in Catalan language on one's national identity. In particular, individuals exposed to mandatory education in Catalan language are less likely to identify with Spain as a nation, more likely to want more independence from Spain, and more likely to support regionalist parties. This paper provides novel evidence in this field by using Spanish data on national identity as an alternative outcome affected by conscription.

2.2 Literature Review

This paper relates to three strands of the literature. First, it contributes to the economic literature on the effects of conscription. This literature has mostly looked at the effects of conscription on labor market outcomes. Angrist (1990)seminal work on the effect of wartime conscription on Vietnam draftees wages finds that, ten years after the war ended, wages are 15 percent lower for white veterans compared to those of non-veterans. Similarly, Imbens and Klaauw (1995) find that conscripts in the Netherlands earn 5 percent less than non-conscript and Angrist and Krueger (1994) find that World War II veterans do not earn more than comparable non-veterans in the long-run.

More recent studies show mixed evidence of the effect of conscription on wages. Angrist et al. (2011) provide an update of the effect on Vietnam veterans wages, showing that the differential with non-veterans tends to zero in the longer run. Card and Cardoso (2012) show that, although, the average difference in wages between Portuguese peacetime conscript and non-conscript is not different from zero, this difference turns positive when looking at less-educated men. This finding shows that the military service might provide valuable skills in the workplace for men coming from poorer socio-economic backgrounds. Siminski (2013) finds strong negative effects of military service on labor force participation for Australian veterans. Bauer et al. (2012) find that, once selection bias is accounted for, the wages and number of days of employment of German conscript do not differ from those of non-conscript.

A second set of papers looks at the effect of conscription on human capital formation.

Usually, conscription is carried out during early adulthood, and in a period of time coinciding with the decision to invest in tertiary education. The findings on the effect of conscription on education are also mixed. Card and Lemieux (2001) find that, thanks to the possibility of avoiding conscription if enrolled in college, men were 4-6 percentage points more likely to be enrolled. Bauer et al. (2014) find that men born after the introduction of conscription in

Germany are 15 percentage points more likely to hold a university degree, corroborating the finding that men might decide to enroll to university to avoid the draft. Oppositely, Cipollone and Rosolia (2007) find that men exempted from compulsory military service were 2 percentage points more likely to graduate from high-school in Italy, while Di Pietro (2013) finds that later on the abolition of conscription had no effect on university enrollment in the same country.

More unanimous results on the direction of the effects of conscription come from the literature on health and crime. Bedard and Deschênes (2006) find that early mortality is higher among World War II and Korean War veterans. These findings on mortality are corroborated by Johnston et al. (2016), who find strong negative effects of military service on mental and physical dysfunction for Australian Vietnam veterans. Autor et al. (2011) find that Vietnam veterans are more likely to obtain disability-related transfer income, and that the recipient of these transfers are more likely to be affected by Post-Traumatic Stress Disorder.

Galiani et al. (2011) find strong causal evidence of the effect of conscription on the probability of men having post-service criminal records. Differently from most of the findings on health outcomes, these effects are robust for conscripts serving both during wartime and during peacetime. Similarly, Hjalmarsson and Lindquist (2016) find that peacetime military conscription increases the probability of committing crime after service for men between 23 and 30. Overall, this first literature shows mixed effects on the costs of conscription. However, these evidence should be taken into account when advancing policies aiming at keeping or returning to a conscription based army. In particular, benefits coming from the hypothesized increase in civic engagement should offset the costs of conscription in terms of human capital and crime.

The second contribution of this paper is to the literature on the formation of political ideology and voting participation. Being conscripted might affect individual political preferences in several ways. Conscription is carried out during early adulthood, when political and social values are formed. Being trained to discipline, to respect the authority and to identify with the military might make one more likely to engage in civic life and to obey to civic duties, e.g. by going to vote. At the same time, this experience might shape preferences and ideology towards more authoritarian positions.

The economic literature on political preferences shows that political ideology matters in shaping preferences for redistribution and taxation (for a review see Alesina and Giuliano, 2009). Different institutional contexts (Alesina and Fuchs-Schündeln, 2007; Alesina et al., 2004), macro-economic shocks during early adulthood (Giuliano and Spilimbergo, 2014) are only some of the factors that might make one more favorable towards higher levels of redistribution and lower levels of inequality. In fact, Giuliano and Spilimbergo (2014) find that if an individual is exposed to a macroeconomic recession during the age period between eighteen and twenty-five, she is more likely to support redistribution and to vote for left-wing parties. For their identification, Giuliano and Spilimbergo (2014) rely on the impressionable years theory (Krosnick and Alwin, 1989) stating that individuals form their political and social values during this period of time, and that this values change little over time.

In line with the impressionable years hypothesis, empirical evidence finds that voting patters tend to be sticky over time. Caprettini et al. (2019) find that individuals in Italy became more likely to vote for the Democristian party as a result of land redistribution reforms, and that this effect persists for forty years after the reform is implemented. Differences in political ideology and preferences for redistribution determine the support for political parties (Giuliano and Spilimbergo, 2014; Barth and Moene, 2016). Traditionally, more left-wing party would support political agendas with higher levels of public spending and more progressive taxation than more right-wing parties.

The literature on the formation of political preferences and voting behavior has recently gained new interest by economists. Mounting empirical evidence and the recent electoral success of extremist parties suggest that other factors might determine the formation of political preferences and the support for parties in elections. Education is correlated with supporting more left-wing parties embodying more tolerant views towards cultural and economic internationalization (Piketty, 2018), globalization and the pressure put on wages exposed to higher competition because of trade lead to an increasing preference for Republicans in the U.S. (Rodrik, 2018; Autor et al., 2016), austerity policies determined the growing success of UKIP in the U.K. (Fetzer, 2019). The findings presented in this paper investigate both possible effects of conscription on political preferences, which might matter for voting behavior, and additional evidence on the long-run effects of policies affecting individuals in their impressionable years.

An additional contribution of the paper is to the political science literature. Campbell and McCormack (1957) and Roghmann and Sodeur (1972) put forward the hypothesis that exposure to the use of weapons, and to a discipline-based training might lead conscripts to having more right-wing ideology. However, the findings suggest no correlations between conscription and ideology. The contribution of this paper is to test this mechanism using a rigorous econometric analysis and exploiting a regression discontinuity design to identify a causal effect. Moreover, this paper looks at the effects of conscription directly on voting participation as a measure of civic engagement. Thus, it contributes to the social capital literature often using voting participation as a determinant of social cohesion, trust, political accountability and well-being (Putnam, 2000; Nannicini et al., 2013; Butler et al., 2016).

Thirdly, section 2.7.1 shows evidence of the effect of conscription on the formation of national identity. Kranton (2016) reviews the literature on the formation of identity and its effect on education, labor supply, work effort and consumption. In particular, Akerlof and Kranton (2005) provide a good example of how conscription and being a member of the armed forces can shape identity. During this experience, individuals are isolated from the rest of society, they share the goal of maximizing their efficiency on the battlefield by incurring potentially the highest cost. Thus, the role of identity with the armed forces and, more in general, with the State is crucial for individuals to exert the effort needed to maximize their efficiency. This model is obviously more related to professionals being part of the armed forces. However, the type and objective of military training are similar to those of conscription.

Empirical evidence on what shapes identity, and in particular national identity, is still scarce (Clots-Figueras and Masella, 2013). In this paper, I present evidence on the effect of conscription on the formation of Spanish identity. The context of Spain is of particular interest when studying what shapes identity. The presence of strong separatist movements has always been present in Spanish politics and has culminated in more (ETA terrorist group attacks) or less (turmoil post- Catalan referendum) violent episodes of conflict, with negative effects on the economic performance in the country (Abadie and Gardeazabal, 2003). The literature on this presents evidence of the effect of different reforms on the probability of having more separatist attitudes. Clots-Figueras and Masella (2013) show that being taught in Catalan language makes one's identity as Catalan stronger and increases the probability to vote for Catalan parties. However, additional empirical studies are rare. Thus, the contribution of this paper is to show whether conscription might contribute to strengthen national identity.

2.3 Context and data

Conscription was introduced in the Federal Republic of Germany (hereafter, West Germany or Germany) in 1955. This reform coincided with the admission of West Germany in NATO and with the creation of the Bundeswehr (national army).

The law regulating the formation of the national army is the Wehrpflichtgesetz. Differently from what happened 6 months later in the German Democratic Republic, this law rules that the national army should be mostly formed by conscripted men. The law ruled that fit-to-serve men born in or after the 1st July 1937 had to serve in the armed forces for a period of twelve months when turning eighteen, and that the first cohort to be drafted would start serving in May 1955. Women were exempted from the draft until 2001. An alternative type of service was available for conscientious objectors since the introduction of conscription. However, the sample size is negligible for the cohorts used in this analysis, and remained so until the 1970s (Bauer et al., 2012). The introduction of conscription followed from the post-WWII sanction of de-militarization imposed to Germany. Thus, the introduction of this reform lends itself to be evaluated using regression discontinuity design thanks to the clear-cut jump in the probability of having been conscripted for men eligible to be drafted after the reform was enacted. In fact, the probability of having being drafted increases exogenously for men born after the cutoff date compared to men who never served because younger than 18 at the end of WWII and born before July 1937.

The main analysis is conducted using data from the German Socio-Economic Panel (hereafter, SOEP). The SOEP runs yearly since 1984 and collects information at the individual level on biographic characteristics, on political preferences and voting participation in federal parliamentary elections. The key information needed to perform the analysis of this paper is on individual date of birth of the individuals. Individual level information on date of birth measured in months is collected in the biographic questionnaire. This was first run in 2001 and asked once and retrospectively to all individuals surveyed in that year and in following waves. The same questionnaire can be used to derive information on the probability that an individual was conscripted. In fact, among other information on socialization experiences during youth this questionnaire includes the question “Did you do the military/civil service?”. Information on military or civil service is available for 15,534 individuals. The sample is restricted to only men and to men residing in West Germany in 1989. Information on whether the individual was living in West Germany before re-unification is asked yearly in the individual level questionnaire of the SOEP.

The final sample is of 11,137 men for which there is information on month of birth, military service and place of residence before 1989. Finally, the sample used in the analysis is composed of men born 48 months before and after the cutoff date of interest. The size of the sample decreases further when it is matched to information on the outcome variables of interest. In fact, to measure right-wing ideology I use the question on political ideology that was asked only in three waves of the survey (2005, 2009, 2014).

The survey asks the respondent “In politics, people often talk about "left" and "right" when describing different political views. When you think about your own political views, how would you rate them on the scale below?”. The individual can answer reporting political views to be on a scale from zero “Far left” to ten “Far right”. The median value in the sample is five. Thus, the analysis on right-wing ideology is conducted on a binary variable, “Rightwing” equal to 1 if the individual reports values above the median, i.e. values from six to ten inclusive. As a robustness check I present results for a trimmed version of this variable, by dropping observations in the the top and bottom ten percent of the original political attitudes variable, i.e. excluding values lower than three and higher than seven. I also present the results obtained by keeping the outcome variable in its original format. The question on voting participation was asked in the two waves following the general elections of 2009 and 2013. The variable of interest for the analysis is a binary variable equal to one if the individual reports to have voted in either of the two elections.

As robustness checks to the main analysis, I explore whether conscription affected other activities that can measure civic engagement, namely participation in local politics and volunteering activity. In waves 2003, 2008 and 2013 the survey included a section on participation in local activities and organizations. Following the literature on social capital (Knack and Keefer, 1997), I construct two additional binary variables. The original question used asks “Now some questions about your free-time. Please indicate how often you take part in each activity: daily, at least once a week, at least once a month, seldom or never" for a series of activities, including "Participation in public initiatives, in political parties, local government" and "perform volunteer work". The two variables derived from this question, i.e. "local politics" and "volunteers", take the value of zero if the individual reports never and one otherwise.

Lastly, an additional analysis is run on the individual's stated support for the two main political parties in Germany: SPD and the coalition of CDU (Christian Democratic Union) and CSU (Christian Social Union in Bavaria). Table 9 contains information on the support for these parties. The SOEP survey asks yearly "Which party do you lean toward?" and presents a list of the major parties in Germany. Given that the CSU is a regional party for Bavaria, the analysis on support for CDU/CSU is conducted on an outcome variable equal to one if the individual states to support either of these parties and zero otherwise. Similarly, the "SPD" variable is equal to one if the individual reports to support the Social Democratic Party and zero otherwise.

2.4 Identification Strategy and methodology

I use the introduction of mandatory military service as an exogenous shock to the probability of having been conscripted for men born right after its introduction compared to men born right before. Provided that these two groups are comparable, I can identify the causal effect of conscription on political attitudes by looking at the difference in the outcome for these two groups.

I use a regression discontinuity design to analyze the effect of conscription on the political attitudes of individuals residing in West Germany at the time of the reform. The fundamental identifying assumption is that the probability of being drafted for men born before and after the reform date depends only on one's date of birth, and not from individual characteristics, i.e. I can rule out selection into treatment. Thus, using a regression discontinuity design I can estimate the probability of having been drafted as a function of date of birth for a small number of cohorts born just before and just after the change in reform. Restricting the sample to cohorts born just before and just after the reform allows to claim that the difference between the treatment group (affected by the reform) and the control group (those not affected) are not statistically different from zero. If these two groups are comparable in the observable and unobservable characteristics, if the reform was binding and if the density of the running variable is continuous around the cutoff date, the only difference between the treated and control groups is the treatment effect.

First, I estimate the probability of having been conscripted for men born in or after the 1st July 1937 compared to men born before this cutoff date. Thus, the treated group affected by the reform are men eligible for conscription, while the control group are men older than eighteen at the time of the reform (and who did not fight in World War II) and no longer eligible to be drafted. I estimate the following set of regressions using parametric and non-parametric techniques (Lee and Lemieux, 2010; Calonico et al., 2014a). I indicate the treatment status using a binary variable which is equal to one if the individual was born in or after the cutoff date. I define the cutoff $\bar{c}$ to be the date at which the reform was enacted. Thus :

$$Treated_i \begin{cases} Born_i \geq \bar{c} & D_i = 1 \\ Born_i < \bar{c} & D_i = 0 \end{cases}$$

Then, I use the following equation to estimate the probability of having done military service for men born before or after the cutoff date:

$$conscription_{i,t} = \alpha_0 + \alpha_1 Treated_i + \alpha_2 AgeAtReform_{it} + \alpha_3 X_{i,t} + \varepsilon_{it} \tag{9}$$

where $conscription_{i,t}$ is the probability of having been conscripted. $Treated_i$ is a dummy variable equal to one if the individual was born after the change in reform and zero otherwise. Hence, the sign of α_1 is expected to be positive, indicating an increase in the probability of having been conscripted, and α_1 measures the probability of having been conscripted for cohorts born after the change in reform of military service. The reduced form equation is:

$$y_{it} = \gamma_0 + \gamma_1 Treated_i + \gamma_2 AgeAtReform_{it} + \gamma_3 X_{it} + \eta_{it} \tag{10}$$

While the second stage equation estimating the causal effect of conscription, instrumented with the treatment variable, on the outcome of interest is as follows:

$$y_{it} = \beta_0 + \beta_1 con\hat{scription}_{i,t} + \beta_2 AgeAtReform_{it} + \beta_3 X_{it} + \theta_{it} \tag{11}$$

where y_{it} is a binary variable indicating political attitudes and voting participation. $con\hat{scription}_{i,t}$ is the probability of having been conscripted instrumented with individual's date of birth estimated using equation 9. All regressions include the same set of individual level variables $X_{it}$to control for: year and region of residence, marital status, labor force status, years of schooling, place of residence during childhood and parental education. β_1 is the coefficient of interest and measures the causal effect of conscription on the outcome of interest.

I restrict the sample of analysis to men born in just before and just after the cutoff data. I use data driven techniques as in Calonico et al. (2014a) to define the optimal bandwidth around the cutoff, which allows to optimize the trade-off between sample size and precision of the estimates. I then use the same optimal bandwidth in the estimation of parametric regressions. For the parametric estimation I estimate equations 9,10 and 11 using interacted polynomials of the running variable, measured in month of birth. Using interactions of the running variable function allows to control for a change in the slope of such function on the left-hand side and on the right-hand side of the cutoff.

I conduct the estimation using both a quadratic and a linear specification to test the sensitivity of the results to higher order polynomials (Gelman and Imbens, 2019; Lee and Lemieux, 2010). It is worth mentioning at this point that the main analysis is conducted on West Germany because the information on individual date of birth is available at the month of birth level, while the data for Spain only contains information on year of birth. The advantages of having information on the month of birth are twofold. First it allows to conduct the analysis using non-parametric data driven techniques. These methods, as detailed in Lee and Lemieux (2010), use data driven techniques and data points close to the cutoff to estimate the coefficients without knowing the functional form of the running variable. Using an incorrect functional form is particularly problematic in the case of Regression Discontinuity Design, as the jump in the function might vary substantially depending on the functional form used in the regression, leading to biased estimates. Thus, running both parametric and non-parametric regressions is key to test whether the results depend on the functional form of the running variable. In order to estimate the regressions using non-parametric techniques for a small bandwidth around the cutoff, researchers face a trade-off between sample size and precision of the estimates (see pp. 314-318 in Lee and Lemieux, 2010).

It is, then, necessary that the number of observation close to the cutoff is high enough to produce precise estimates. Moreover, the finer the unit of measurement of the running variable the more similar units on the left and on the right of the cutoff can be claimed to be. For these reasons, when having only information on year of birth, the optimal bandwidth needed to estimate the non-parametric regressions might be too large and the treatment effect for that bandwidth will be biased. Following from this discussion, when estimating the regressions using Spanish data I will only implement parametric techniques.

The remaining of this section aims at presenting graphical evidence of the internal validity of the identification strategy. To do this I provide evidence of the first stage discontinuity by showing a jump in the probability of having been conscripted for the treatment group compared to the control. Figure 3 shows the discontinuity at the cutoff for men born before and after the introduction of conscription in Germany. Additional evidence is needed to show the continuity of the running variable density around the cutoff. Figure 14, in Appendix 2, shows graphical evidence of the Mc Crary test for the discontinuity of the running variable (McCrary, 2008). This shows that there is not a statistically significant jump in the probability of being born after the reform.

To test for manipulation around the cutoff, figure 15 (in Appendix 2) shows graphical evidence that the discontinuity in the probability of having been conscripted is not present one year before or one year after the implementation of the reform used in the analysis. This suggests that the reform was effective and there were not significant anticipation or delay effects in the implementation of the reform. These evidence is corroborated by the regression analysis presented in section 2.7.

Figure 3: Discontinuity in the probability of having been conscripted by month of birth for men born within 48 months before and 48 after the introduction of conscription: West Germany.

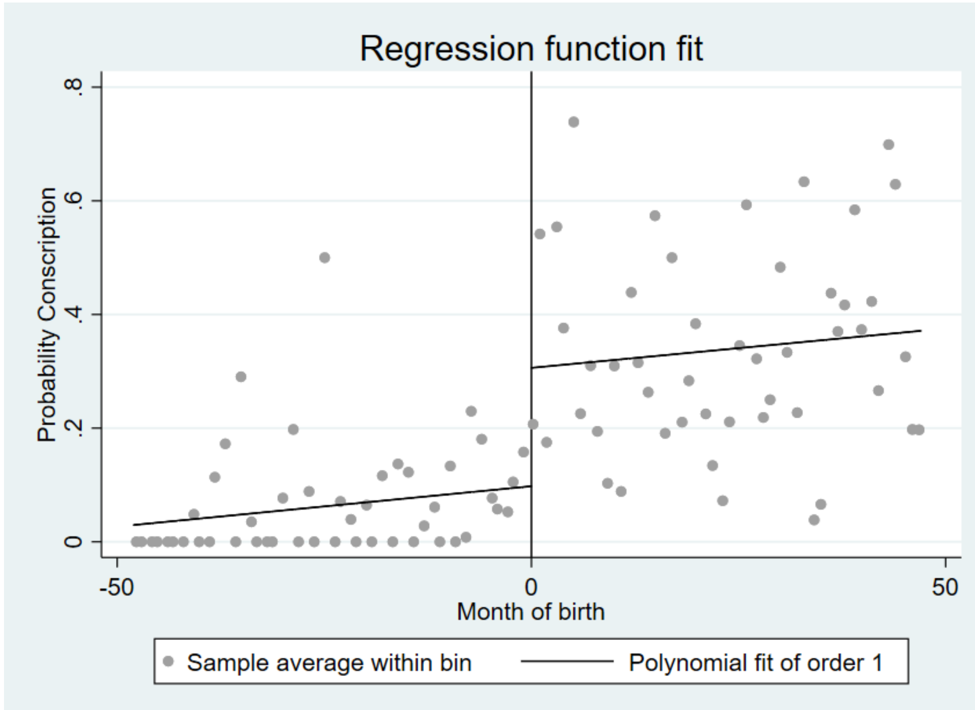

2.5 Summary statistics

The main objective of this section is to outline the characteristics of the treatment and control group. This information allows to verify whether the treatment and control groups are comparable in terms of observable characteristics. It also informs on whether the sample of analysis is representative of the whole population.

Table 9 summarizes the characteristics of the full sample (column 1), of the sample used in the analysis and born 48 months before and after the introduction of conscription (column 2), of the treated group born after the cutoff date (column 3) and of the control group born before the cutoff date (column 4). Column 5 presents the difference in means for each variable by treatment status. First, table 9 shows the difference in the probability of having been conscripted for the treatment and the control group. In particular, men born within 48 months after the introduction of conscription are 27 percent more likely to have been conscripted compared to men born 48 months before. The table also shows that the proportion of individuals carrying out civil service is marginal in the sample of analysis. This is in line with findings in the literature stating that conscientious objector started constituting a consistent group of conscripts only since the early 1960s (Bauer et al., 2012). Second, table 9 shows the difference in age for men in the treatment and control group. This difference is the mechanical consequence of the identification assumption, i.e. one's date of birth, resulting in men in the treatment group to be 4 years younger on average. As discussed in the previous section the bandwidth of 48 months is estimated using non-parametric techniques.

Then, column 5 shows that the treated and control group are comparable on average in terms of individual characteristics, with the exception of home-ownership status being higher for the treated group. In terms of comparability with the full sample. I observe that the sample used in the analysis is around 20 years older than the average, less likely to have been conscripted, and more likely to be out of the labor force. However, it is representative of the population in terms of average household income, and schooling.

One of the advantages of the SOEP data is the availability of retrospective information in the biographic questionnaire on pre-treatment individual characteristics. Table 9 presents this information. I use questions on place of residence during childhood and parental education. The variable on place of residence during childhood takes four values for: outside urban areas (baseline category), small city, medium-sized city and big city. I recode the original variable to obtain four binary variables to be included in the regressions as controls. This is to control for potential unobservables that determine the probability of being conscripted based on the place of residence. Similarly, including controls for parental education helps controlling for socio-economic background of conscripts and non-conscripts.

The literature on conscription finds that pre-treatment socio-economic characteristics are one of the key determinants of heterogeneous effects of service on wages (Card and Cardoso, 2012) and crime (Hjalmarsson and Lindquist, 2016). Hence, I include parental education to both account for a direct effect of this variable on the probability of being conscripted (more educated parents might decide to send their children to university, for example, or transmit different political values), and to use education as a proxy of parental income. Since almost the full sample has German parents, I exclude this control from the analysis. Although controlling for these characteristics should account for unobservable and observable pre-treatment differences, it is worth commenting on the statistical significant differences in pre-treatment characteristics.

I notice that there is a statistical significant difference in the probability of having lived in a medium or big city. However, if there was an issue of selection such that men residing in bigger urban areas (plausibly characterized by different civic and social values compared to smaller cities), were less likely to be conscripted the sign of this difference would be negative and increasing in magnitude for the binary variables *SmallCity*, *MidCity* and *BigCity*, but this is not the case for the second. Similarly, I fail to detect a pattern suggesting selection into treatment when looking at parental education. Moreover, the differences statistically different from zero are of very small magnitude throughout column 5 in table 9.

The bottom part of table 9 shows the summary statistics for the outcome variables included in the main analysis and in the robustness checks. It is worth noting that the proportion of individuals that voted in the federal parliamentary elections of 2009 or in those of 2013 is very high in the whole sample, with a turnout of 89 percent. This proportion is even higher among the sample used in the analysis, plausibly showing a higher tendency for older generations to go to vote. This high level of turnout should be kept in mind when interpreting the results and when proposing policy implications, as there is limited scope to increase turnout further. The sample sizes in the last line of table 9 are indicative. They correspond to the cross-sectional sample of individuals for which information on month of birth and conscription is available. However, the sample size will vary in the regression analysis depending on the outcome variable analyzed, and on whether the panel component is used in the analysis.

Table 9: Summary statistics of the sample individual characteristics for men only in Germany: full sample, born 48 months before and after the cutoff, and by treatment status.

		RD sample			
	(1)	(2)	(3)	(4)	(5)
	All	+/-48 Months	Treated	Control	Diff
Individual characteristics					
Conscription	0.52	0.24	0.35	0.08	0.272 ***
Treated	0.89	0.59	1.00	0.00	1.000
Civil service	0.20	0.01	0.01	0.00	0.011
Age	51.83	71.44	69.96	73.60	-3.643 ***
Couple	0.70	0.79	0.80	0.79	0.009
Working	0.67	0.11	0.12	0.09	0.029 *
Homeowner	0.63	0.73	0.74	0.70	0.045 *
HH Income	7.97	7.81	7.81	7.80	0.014
No School cert.	0.01	0.01	0.00	0.01	-0.007
Elementary	0.05	0.05	0.05	0.05	0.000
Higher	0.31	0.30	0.30	0.30	-0.001
Years Schooling	12.88	12.46	12.45	12.48	-0.026
Baseline characteristics					
Residence in childhood					
- Big city	0.24	0.27	0.25	0.30	-0.050 **
-Mid city	0.18	0.14	0.16	0.12	0.034 *
-Small city	0.21	0.20	0.19	0.21	-0.020
-Out city	0.38	0.39	0.41	0.37	0.036
Father's education:					
- No School	0.02	0.01	0.01	0.01	0.002
- Secondary School	0.68	0.75	0.74	0.77	-0.035 ***
- Middle School	0.14	0.09	0.10	0.08	0.014
- Higher	0.15	0.14	0.14	0.13	0.013
Mother's education:					
- No School	0.03	0.02	0.03	0.02	0.010
- Secondary School	0.72	0.79	0.79	0.80	-0.016
- Middle School	0.17	0.13	0.13	0.14	-0.015
- Higher	0.07	0.05	0.06	0.03	0.022 *
Father German citiz.	0.94	0.96	0.97	0.96	0.004
Mother German citiz.	0.95	0.98	0.97	0.98	-0.003
Outcome variables					
Rightwing	0.31	0.37	0.38	0.34	0.044
Voted	0.89	0.95	0.95	0.95	0.001
CDU/CSU	0.44	0.52	0.52	0.52	0.003
SPD	0.33	0.36	0.36	0.36	-0.005
Volunteers	0.33	0.37	0.36	0.37	0.009
Local politics	0.16	0.19	0.18	0.19	0.010
N	11,137	1,312	533	779	

Note: Summary statistics for the sample used in the main analysis. Column 1 presents the mean value of each variable for the whole population. Column 2 presents the mean value of each variable for the sample used in the analysis. Column 3 and 4 provide information for the treated and control group. Treated is a binary variable equal to one if the individual is eligible to be drafted in 1955. Column 5 presents the difference in means between the treated and control average values.

2.6 Results

This section presents the results of the analysis.

Before moving to the results of the estimation of the causal effect of conscription on political attitudes, this section shows qualitative evidence of the correlation between conscription and the outcomes of interest. This proves useful as these correlations differ consistently from the estimated causal effects, which are then validated by a series of robustness checks. Following from the results presented in this section, I conclude two things. First, that looking at simple correlations produces biased results and misleading policy implications. Second, that I cannot find statistical significant evidence of a causal effect of conscription on either political ideology or voting participation.

Table 10 shows the results of a simple OLS regression of the binary variable, indicating whether the individual was conscripted[8], and the outcomes of interest. This naive regression estimates the correlation between having been conscripted and all the outcome variables later used in the analysis. The results presented in table 10 include all the controls used in the analysis: year of survey, region of residence, age, marital status, homeownership status and labor force status, and the baseline covariates controlling for place of residence during childhood and parental education. The regressions are run on the whole sample, regardless of the year of birth, and show that conscripts are around 2 percentage points less likely to have more right-wing political attitudes (column 1), and this is insensitive to trimming the outcome variable at the top and bottom ten percent (column 3). Conscripts are also 2 percentage points more likely to vote. However, there is no correlation between conscription and participating in local politics (column 4), volunteering (column 5) or supporting the main political parties in Germany (columns 6 and 7). Thus, by just looking at these correlation I might conclude that conscription is effective, although this effect is small, in increasing civic engagement measured with voting participation, while it does not increase the probability of being more right-wing.

I explore this second correlation further by looking at the different correlation between mandatory military service and political attitudes, relative to civil service. Table 28, in Appendix 2, shows the correlation between having done military service as opposed to civil service

[8]In this case, since I am using the whole sample, individuals might have opted both for the military and for the civil service.

(around 20 percent of the full sample opted for civil service, see table 9). Alternative types of service were available to conscientious objectors, mainly in the form of health-care sector activities. However, the first conscientious objectors started service in 1961 (Bauer et al., 2012). Thus, looking at heterogeneity by type of service is only possible in the correlation analysis. The results in table 28, in Appendix 2, show that men who did the military service are 15 percentage points more likely to have right-wing ideology and 14 percentage points more likely to support the coalition of CDU/CSU relative to men who did the civil service, suggesting that the average negative result of 2 percentage points reported in table 10 hides an important heterogeneity between these two groups.

Although these correlations should be interpreted with care, they show a relationship between the type of service and political attitude. Relationship that might be present before the treatment and hide selection issues, further undermining the correlations presented in table 10. Following from this correlation analysis, I conclude two things. First, the correlation between conscription and political outcomes differs by type of service, suggesting that the correlations presented in table 10 might be affected by selection issues. Second, and more relevantly, these correlations suggest a positive and statistical significant relationship with voting participation and a negative relationship with right-wing ideology. I show in the remaining of this section that such relationship are not causal and cannot be used to inform policy recommendations.

Table 10: Regression analysis of the correlations between conscription and political ideology, party support and civic engagement measures in West Germany: full sample of men.

	Rightwing (1)	Voted (2)	Rightwing: trim (3)	Local politics (4)	Volunteers (5)	CDU/CSU (6)	SPD (7)
Conscription	-0.025**	0.026***	-0.025**	-0.008	-0.011	0.002	-0.018
	(0.011)	(0.008)	(0.012)	(0.009)	(0.013)	(0.015)	(0.016)
Year + Age	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y
Baseline	Y	Y	Y	Y	Y	Y	Y
N	6115	4518	5262	5213	5226	3831	3831

Note: Each coefficient is estimated using a separate regression. Rightwing is a binary variable equal to 1 if the individual reports to have a political ideology more rightwing than the median. Voted is a binary variable equal to 1 if the individual voted in the last general election. Conscription is a binary variable equal to 1 if the individual was conscripted. All regressions include controls for: year of the survey, age and region of residence at the time of the survey, marital status, labor force participation, years of schooling, place of residence during childhood and parental education. SE clustered at month of birth

There are several reasons to move beyond these correlation analysis and explore the causal estimates of conscription on political attitudes.

Conscription might not be randomly allocated across individuals, and this would produce biased estimates of the effects of conscription on the political outcomes considered. Individuals might be more likely to be conscripted based on pre-treatment socio-economic characteristics. This would mean that the correlations presented in table 10 are biased and might be driven by the correlation between pre-treatment characteristics (e.g. parental income), rather than conscription, and political attitudes. Other observable and unobservable factors might be correlated with a higher probability of been conscripted and with the outcome variables of interest. To overcome these limitations, I estimate the effect of conscription on political attitudes using the methodology outlined in section 2.4.

As standard in the regression discontinuity literature (Lee and Lemieux, 2010), I first explore the tested hypotheses using graphical evidence. Figure 3 shows the jump in the probability of being conscripted for men born right before and right after the reform. I provide similar evidence for the main outcome variables, figure 4 shows the discontinuity in right-wing ideology and voting participation for men affected by the reform compared to men who were not. In both cases the jump in the probability is positive, confirming the correlation analysis and the hypothesis put forward by policy-makers for which conscription increases civic engagement, and by political scientists for which it also increases the probability of having right-wing ideology.

Figure 4: Discontinuity in the outcome variables by month of birth for men born within the 48 months before and the 48 months after the introduction of conscription in West Germany.

(a) Right-wing ideology

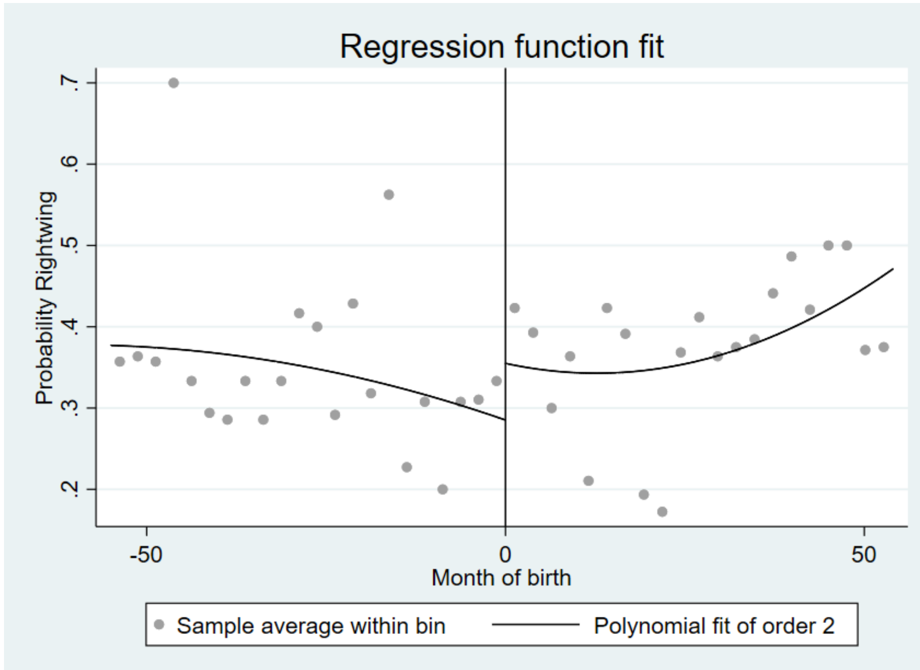

(b) Voting participation

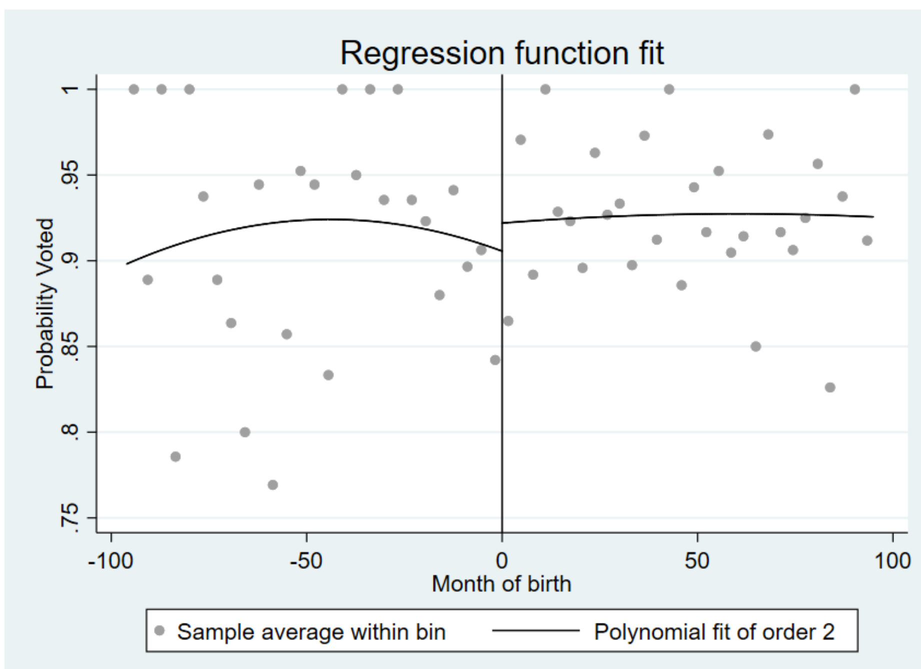

Table 11: The effect of conscription on right-wing ideology in West Germany: main results estimated using non parametric and parametric methods.

	(1) FS	(2) RF SS	(3) FS	(4) RF SS	(5) FS	(6) RF SS
Panel A: Non-Parametric						
Treated	0.224 ***	0.045	0.217 ***	0.047	0.196 ***	0.038
	(0.060)	(0.055)	(0.059)	(0.055)	(0.054)	(0.055)
Conscription		0.203		0.216		0.216
		(0.245)		(0.252)		(0.275)
Year+ Age	N	N	Y	Y	Y	Y
Controls	N	N	N	N	Y	Y
Bandwidth		50.5		49.1		60.2
N		7397		7397		6104
N Eff. obs.		867		849		936
Panel B: Parametric with interactions						
Treated	0.198 ***	0.010	0.195 ***	0.004	0.189 ***	0.015
	(0.055)	(0.059)	(0.053)	(0.058)	(0.054)	(0.064)
Conscription		0.048		0.021		0.021
		(0.293)		(0.294)		(0.330)
Year+ Age	N	N	Y	Y	Y	Y
Controls	N	N	N	N	Y	Y
Window		48 months		48 months		48 months
Pol. Order		1		1		1
F-Stat		12.7		13.7		12.1
N		829		829		756
Treated	0.272 ***	0.114	0.267 ***	0.110	0.258 ***	0.090
	(0.083)	(0.074)	(0.080)	(0.076)	(0.081)	(0.084)
Conscription		0.419		0.411		0.411
		(0.286)		(0.296)		(0.324)
Year+ Age	N	N	Y	Y	Y	Y
Controls	N	N	N	N	Y	Y
Window		48 months		48 months		48 months
Pol. Order		2		2		2
F-Stat		10.9		11.1		10.2
N		829		829		756

Note: Each coefficient is estimated using a separate regression. Rightwing is a binary variable equal to 1 if the individual reports to have a political ideology more rightwing than the median. Treated is a binary equal to 1 if the individual was born after the introduction of mandatory conscription, i.e. in or after July 1937. Conscription is the predicted probability of having done military service estimated from the FS regression. Panel A presents the coefficients of the estimation with non-parametric methods around the optimal bandwidth (Calonico et al. 2014). Panebl B presents the coefficients of the parametric estimation with interactions of the running variable before and after the cutoff date. All regressions include controls for: year of the survey, age and region of residence at the time of the survey, marital status, labor force participation, years of schooling, place of residence during childhood and parental education. Regressions include interacted second order polynomials of the running variable (month of birth). Standard errors are clustered at the month of birth.

Table 11 shows the results of the estimation of the effect of conscription on the probability of having more right-wing political ideology than the median. The results are presented as follows. Column 1 and 2 present the results when no additional controls are included in the regression. Columns 3 and 4 include controls for year of the survey and age. Columns 5 and 6, show the preferred specification and includes the full set of controls outlined in section 2.5.

Standard errors are clustered at the month of birth.

Panel A in table 11 shows the results of the estimation of equations 9, 10 and 11 using non-parametric techniques for the optimal bandwidth around the cutoff as in Calonico et al. (2014b). The odd-numbered columns in the top line show estimates for the first stage equations. These show that individuals born after the introduction of conscription are around 21 percentage points (with estimates ranging from 18.9 to 22.4 percentage points) more likely to have been conscripted compared to men born before the change in reform. These results are calculated for an optimal bandwidth of 49 to 60 months[9] before and after the cutoff date, and remain stable across the different specifications. The reduced form shows that the effect of the reform on right-wing ideology is of around 4 percentage points and not statistically different from zero. The second row of Panel A shows the results of interest, i.e. the effect of conscription on right-wing ideology. Here *Conscription* is the predicted probability of having been conscripted instrumented with the treatment binary variable being equal to 1 if the individual was born in or after the 1st July 1937. The effect is positive and of around 21 percentage points, but not statistically different from zero.

Panel B in table 11 shows the results of the estimation of equations 9,10 and 11 using parametric techniques for the optimal bandwidth around the cutoff date. To ensure comparability and consistency throughout the estimations of this analysis, the bandwidth of the parametric estimation is approximated to 48 months before and after the reform was implemented. The treatment group is thus men born 48 months after the 30th June 1937, while the control group are men in the 48 months before the reform. The parametric estimation is run using linear (first set of results in Panel B) and quadratic (second set of results in Panel B) polynomials of the running variable, i.e. age at the time of the reform measured in months.

The first row of Panel B in table 11 shows that the first stage estimates are in line with those found in the non-parametric analysis. The effect of the reform was to increase the probability of having been conscripted by 19 percentage points for men in the treatment group. The second stage estimates are presented in the following line and show that the effect of conscription on right-wing ideology is positive and non statistically significant. The magnitude of these effects varies in the two specifications but bounces around the average of 21

[9] The width of the bandwidth changes as I include more control variables in the estimation.

percentage points estimated in the non-parametric analysis. In particular, when using linear function of the month of birth the effect is an increase of between 2 and 4.8 percentage points in the probability of being right-wing, while it jumps to 41 percentage points when including quadratic polynomials.

Overall, the results for right-wing ideology show a positive but not statistically significant effect of conscription. The magnitude of the coefficients is estimated to be of around 20 percentage points. However, the estimation of the coefficient is not precise and shows very large confidence intervals. This issue is recurrent throughout the analysis, however, the estimates on rightwing ideology presented in table 11are those presenting a higher variation in the magnitude of the coefficients and in the standard errors. I will partly address this issue in the next section when presenting the robustness checks of this analysis. There I will be able to rule out that the effects are as high as those presented in panel B, but I will face similar issues in terms of interpretation of the magnitude and sign of the effect. Given the strong and stable first-stage results and the robustness checks conducted in the following section, I conclude that this is plausibly indicating the absence of a consistent effect of conscription on right-wing ideology. Thus, the policy recommendation that can be drawn from these results is not as clear as if I could estimate a precise coefficient of a zero effect. However, given the costs associated with this policy, basing its re-introduction on an imprecisely estimated positive effect would come with a high risk of the policy to be ineffective. This same reasoning should be kept in mind also when interpreting the results on voting participation. In fact, although the range around the coefficients is narrower, the estimated effects could also be negative suggesting that conscription might reduce the probability of going to vote.

Table 12 shows the results of the analysis on the effect of conscription on voting participation in the general elections of 2009 and 2013. The outcome variable is a binary variable equal to 1 if the individual reports to have voted either in 2009 or in 2013. Column 1 and 2 present the results when no additional controls are included in the regression, columns 3 and 4 include controls for year of the survey and columns 5 and 6 include controls for year of survey and age of the individual at the time of the survey. Panel A presents the results estimated with non-parametric methods, while Panel B presents the parametric specification using, first, linear and, second, quadratic interacted polynomials of the running variable.

Table 12: The effect of conscription on voting participation in West Germany: main results estimated using non parametric and parametric methods

	(1) FS	(2) RF SS	(3) FS	(4) RF SS	(5) FS	(6) RF SS
Panel A: Non-Parametric						
Treated	0.334 ***	0.040	0.292 ***	0.044	0.282 ***	0.036
	(0.070)	(0.045)	(0.062)	(0.039)	(0.063)	(0.035)
Conscription		0.106		0.135		0.135
		(0.124)		(0.125)		(0.115)
Year+Age	N	N	Y	Y	Y	Y
Controls	N	N	N	N	Y	Y
Bandwidth		47.7		59.5		58.1
N		5385		5385		4514
N Eff. obs.		605		742		666
Panel B: Parametric with interactions						
Treated	0.283 ***	0.041	0.281 ***	0.034	0.265 ***	0.052
	(0.065)	(0.040)	(0.062)	(0.039)	(0.063)	(0.038)
Conscription		0.139		0.160		0.160
		(0.133)		(0.135)		(0.126)
Year+Age	N	N	Y	Y	Y	Y
Controls	N	N	N	N	Y	Y
Window		48 months		48 months		48 months
Pol. Order		1		1		1
F-Stat		21.1		24.1		22.4
N		609		609		559
Treated	0.411 ***	0.047	0.408 ***	0.064	0.382 ***	0.070
	(0.092)	(0.063)	(0.089)	(0.062)	(0.089)	(0.062)
Conscription		0.074		0.092		0.046
		(0.134)		(0.136)		(0.130)
Year+Age	N	N	Y	Y	Y	Y
Controls	N	N	N	N	Y	Y
Window		48 months		48 months		48 months
Pol. Order		2		2		2
F-Stat		23.6		25.8		23.0
N		610		610		560

Note: Each coefficient is estimated using a separate regression. Voted is a binary variable equal to 1 if the individual reports to have voted in the 2009 and 2013 elections. Treated is a binary equal to 1 if the individual was born after the introduction of mandatory conscription, i.e. in or after July 1937. Conscription is the predicted probability of having done military service estimated from the FS regression. Panel A presents the coefficients of the estimation with non-parametric methods around the optimal bandwidth (Calonico et al. 2014). Panebl B presents the coefficients of the parametric estimation with interactions of the running variable before and after the cutoff date. All regressions include controls for: year of the survey, age and region of residence at the time of the survey, marital status, labor force participation, years of schooling, place of residence during childhood and parental education. Regressions include interacted second order polynomials of the running variable (month of birth). Standard errors are clustered at the month of birth.

Panel A in table 12 shows the results obtained when using non-parametric techniques for the optimal bandwidth around the cutoff. These show that the effect of conscription on voting participation is not statistically significant, but positive and robust to the inclusion of controls for year of survey and age of the respondent. The results of the first stage show a stronger effect of the reform for the sample used in this analysis[10]. The effect of the reform is an increase in the probability of having been conscripted of around 30 percentage points for men in the treatment group compared to men in the control group.

The optimal bandwidth when no additional controls are included is of 47.7 months, but it increases to 58 when adding the full set of controls. This happens because information on some of the covariates is missing and the drop in sample size makes it so that, when using non-parametric techniques, the bandwidth needed to estimate the coefficients becomes larger. The second stage estimates remain stable across the different specifications and show a statistically insignificant increase in the probability of having voted for men born after the introduction of conscription. The sign of the coefficient is positive and of around 12 percentage points (with coefficients ranging from 10.6 to 13.5 percentage points) across the different specifications for the treated, but again this effect is not statistically different from zero.

Panel B shows that the results obtained when using linear and quadratic parametric specifications are consistent with the ones of the non-parametric analysis. The effect of conscription is positive and between 7 and 13 percentage points on average when no controls are included in the estimation, but not statistically different from zero. It is worth noting that including the full set of controls in the regressions does not change the estimated coefficients by significant amounts, pointing to the fact that no other factors affect either the probability of being conscripted nor the outcomes of interest.

As for the estimates presented in table 11, the results presented in table 12 show very large confidence intervals. Differently from the former, however, the results for voting participation estimated with non-parametric and linear specifications show that the effect is surely positive or very close to zero. When using quadratic polynomials of the running variable the coefficients decrease in magnitude, but still show large standard errors possibly hiding a negative effect, but still with a magnitude closer to zero. Overall, the estimated effect of conscription on voting participation ranges from a negative effect of 8 percentage points to an increase of 28 percentage points. These estimates show less variation compared to those on right-wing ideology (ranging from -0.30 to 0.69 percentage points), but again do not allow me to identify

[10]Notice that the difference in the first-stage estimates depends on the fact that voting participation questions are asked in less waves and that the sample size is smaller than the one used in the analysis for right-wing ideology.

whether the effect of the policy is zero.

To conclude, the analysis presented so far does not corroborate the hypothesis that conscription has long-term effects on voting participation or right-wing ideology. Although the results on voting participation show a positive effect, this is not statistically different from zero. Hence, policies aimed at re-introducing conscription with the motivation of increasing civic engagement of citizens are not backed up by these empirical findings. At the same time, these findings do not show a change in political ideology. Although all of the estimates suffer from the limitation of being imprecisely estimated, they rule out that conscription has a significant and clear-cut effect on political attitudes, which is what would be needed to motivate costly policies aimed at bringing conscription back.

2.7 Robustness checks

I carry out three robustness checks to test the validity of the main results.

First, table 29 (Appendix 2), shows results for different measures of civic engagement and political ideology. Columns 1 and 2 show the results when using the original variable of political ideology as an outcome, i.e. a discrete variable on a scale from one (extreme left) to ten (extreme right). The results are consistent with those of the main estimation and find that conscription increases the probability of having more right-wing ideology, but the coefficients are not statistically different from zero. It should also be noted that the results of this specification help shedding light on the direction of the effect. In fact, although still imprecisely estimated, the standard errors for the coefficient of the second stage are smaller than the coefficient itself pointing to either an estimate very close to zero or to a positive effect of around 1.4 units in the ideology scale. Columns 3 and 4 show the results for right-wing ideology when the top and bottom 10 percent of the political attitudes scale are trimmed. I conduct this second test to check whether the results are biased by the presence of outliers.

The results show similar magnitude (between 22 and 32 percentage points on average) and sign of the coefficients as those presented in table 11. Conscription has a positive effect on having more right-wing ideology, but the effect is not statistically significant. The remaining of the results presented in table 29 investigate whether different types of civic engagement, other than voting participation, might be positively affected by conscription. Columns 5 and 6 show that men born in the 48 months after the introduction of conscription are not more likely to participate in local politics. This effect is not statistically significant and, if something, it has a negative sign. Columns 7 and 8 do corroborate the absence of statistically significant effects when looking at whether individuals participate in volunteering activities.

I then look at support for the two main political parties in Germany. Table 30 (Appendix 2), shows the results for political party supported. These show that conscription does no change the probability of supporting either SPD or the CDU/CSU coalition. These results show that there is no effect of conscription on other measures of political engagement. The robustness checks conducted so far show that the effect of conscription is not only statistically insignificant when looking at voting participation and ideology, but also when looking at different outcomes and when increasing the sample size to gain power in the estimation.

Second, I address possible issues in the estimation of the main results. The coefficients in table 11, for example, bounce around the average effect of 21 percentage points depending on the functional form used in the parametric estimation, suggesting an imprecise estimation of the second stage regression once higher-order polynomials are used in the regression. One of the possible limitations behind these results is that the sample size for the optimal bandwidth is too small and I do not have enough power to estimate the coefficients precisely[11]. To overcome this limitation, I conduct two robustness checks.

First, the results presented in table 31 (Appendix 2), keep the panel-component of the SOEP sample, i.e. repeated observations for the sample of men for which information on conscription and month of birth is available. Following Lee and Lemieux (2010), the use of panel data as pooled cross-section is correct provided that within-individual correlations are accounted for with individual-level clustered standard errors. Hence, this is the approach followed to estimate these results.

Columns 1-4 in table 31 present the estimation for right-wing ideology and columns 5-8 the results for voting participation. All the regressions control for year of the survey, and the specifications in columns 3,4,7 and 8 include controls for age of the individual at the time of the survey. The findings remain in line with those of the main analysis. However, the magnitude of

[11]It should be noted, however, that this is possibly driven by the absence of a clear effect of this policy, rather than a problem with the data or the estimation as the estimation of the first stage is consistent throughout the different models and show a strong internal validity of the instrument used.

the coefficients for right-wing ideology is much smaller and comparable in the non-parametric (Panel A) and parametric (Panel B) specification. These results still show a positive and non statistically significant effect on right-wing ideology, but the magnitude of the effect is between 4 and 14 percentage points for men born after the introduction of conscription. These results address the variation in the magnitude of the second-stage regressions coefficients presented in panel B of table 11. However, the confidence intervals around these coefficients remains very large and spanning from a negative effect of -26 to a positive effect of 50 percentage points.

The coefficients for voting participation remain of the same sign and magnitude as those presented in the main analysis and, although remaining imprecisely estimated, continue to show narrower confidence intervals (with an effect estimated to be positive or very close to zero). Column 6 in table 31 shows that, when controlling for year of survey and age, conscription increases voting participation of 18 percentage points on average and this result is statistically significant at the 10 percent. However, this result is not robust when including the whole set of controls nor when using parametric methods in the estimation, suggesting that the statistical significance of one coefficient is not enough to provide clear evidence of the effectiveness of conscription.

Second, I reduce the bandwidth around the cutoff to 36 months. The regressions are estimated using linear interacted polynomials of the month of birth. Table 32 (Appendix 2) shows that the results hold when estimating the effect of conscription on ideology and voting participation on men born within the 36 months after the introduction of conscription compared to men born in the 36 months before. Moreover, the coefficients estimated using this bandwidth are more precisely estimated and allow me to say something more on the upper and lower bounds of the effects, especially for right-wing ideology. Columns 2 and 4 in table 32 show that the effect of conscription on right-wing ideology and voting participation is either positive or very close to zero even if imprecisely estimated. The more controls are added in the estimation the wider the confidence intervals around the coefficients, but looking at the results of the simpler models allows me to be more confident on the sign of the estimated effects.

Lastly, I check that the policy was binding and there are no other jumps in the probability of been conscripted away from the cutoff. Following from the graphical evidence presented in figure 15 (Appendix 2) the results confirm the absence of anticipation or delays in the implementation of the reform. This is important as, for example, if a delay in the implementation of the reform is correlated with individual characteristics, e.g. less right-wing individuals were more likely to postpone their conscription, the estimates might be biased. Table 33 shows the results of the estimations when the cutoff date is postponed to 12 months after the actual introduction of conscription. I define a placebo treatment equal to one if the individual was born after July 1938 and equal to zero otherwise. The results of the first stage estimation show that there are no jumps in the probability of having been conscripted if the cutoff is set to a different date.

To conclude, I do not find evidence of delays in the implementation of the reform when looking at jumps in the probability of being conscripted away from the cutoff. Moreover, restricting the bandwidth around the cutoff does not affect the results presented in the main analysis. I now conduct a final check of the results that provides evidence of the external validity of this analysis.

2.7.1 Spain

This section completes the analysis on the effect of conscription on political attitudes by implementing the same methodology used so far in a different context and using a different dataset. I do this evaluating the effect of the abolition of conscription in Spain. This allows me to show the external validity of the results presented so far, as well as to show shorter-term effects of conscription and to explore the effect of conscription on the formation of Spanish national identity.

I start by outlining the details of the reform used in the analysis, which mirrors the one used in the analysis for Germany. After the end of Franco's regime in 1975, Spain incorporated several changes to conscription both in terms of the duration of the service and in the introduction of alternative forms of service for conscientious objectors. The topic of abolishing conscription was particularly debated in the country, as in all Western Europe, after the end of the Cold War. In 1991, the Netherlands were the first country to abolish conscription, followed by Belgium, France and in the same year, Spain. In 1996, the newly elected prime minister Jose Maria Aznar (Partido Popular) announced that the military forces of the country will move to a fully voluntary basis armed force. Thus, although the reform was fully enacted and passed as law in 2002, men turning eighteen in 1997 were not obliged to be drafted. Although, there was the possibility of delaying the military service for reasons of study or to have an exemption, the data used in the analysis allow to conduct placebo analysis to control for anticipation or delays in the implementation of the reform.

The analysis on Spain uses data collected by the Center for Sociological Research (CIS-Centro Investigaciones Sociologicas) in the survey on National Defense and the Armed Forces (Defensa Nacional y de las Fuerzas Armadas). The survey was conducted in eleven years in the period of time between 1997 and 2015. The full dataset is composed of repeated cross-sectional data representative at the national level and collecting information for approximately 2500 men and women in each wave. I use half of this sample as conscription was mandatory only for men. The objective of the survey was to assess the attitudes and opinions of the Spanish population towards the Armed Forces and the national military service. Of interest for this research is the fact that the survey asks all respondents whether they completed the military service.

The sample with information on military service is of 11,333 men who were at least eighteen at the time of the survey. The variable measuring military service is equal to one if the individual reports to have done or be doing the military service. Notice that while in the 1997-2000 waves the question asked was “Did you do the mandatory military service?”, from 2005 the question asks “What is your relationship with the Armed Forces?” and the military service variable is equal to one if the individual reports “I did the mandatory military service”.

The outcome variables of interest are asked yearly in the Spanish survey and are very similar to the German counterpart. This allows to construct comparable measures of right-wing ideology and voting participation to the ones used so far. In fact, the survey asks each individual “When talking about politics, usually terms like right and left-wing are used. You can find a scale from zero to ten below, where would you place your political views”, the individual can then choose a value from 1 “Left” to 10 “Right”. In the analysis the variable “Rightwing” is transformed to be equal to 1 if the individual has right-wing views above the median and zero otherwise. The question on voting participation asks “Did you vote in the past general election?”, the variable used in the analysis is equal to one if the individual reported to have gone to vote and zero otherwise, missing values are attributed to men who report to have been too young to vote at the moment of election.

Additionally, the questionnaire asked in 2009, 2011, 2013 and 2015 whether the individual "Which of these sentences better expresses your feeling? (1) I feel only Spanish, (2) I feel more Spanish than$x$, (3) I feel as Spanish as x, I feel more x than Spanish, I feel only as x." where x identifies the regional origin of the surveyed, e.g. Catalan if the individual reports to be form Catalonia. I follow Clots-Figueras and Masella (2013) and define a variable for "Spanish Identity" to be equal to one if the individual chooses options (1) to (3) from the ones listed above, and zero otherwise.

In the empirical analysis, I use this reform to estimate the probability of having been conscripted for men turning eighteen right before and right after the reform. In particular, I define as my treatment group men born in or after 1979 and turning eighteen in 1997, when the suspension of conscription was enacted, and no longer eligible for conscription. When estimating the first-stage equation (using equation 9), $Treated_i$ in the Spanish analysis is a binary variable equal to one if the individual was born in or after 1979 and zero if born before that year. The sign of α_1 is expected to be negative in this second case, indicating a decrease in the probability of having been conscripted as a result of the reform.

In line with this prediction, figure 5shows the discontinuity in the probability of doing military service for men born before and after the cutoff date. I provide further evidence of the internal validity of the identification strategy in figures 16 (McCrary test) and 17 (change in cutoff year) in Appendix 2. The main difference with the main analysis is that the CIS data provide information of individual date of birth only measured in years. Thus, the analysis is performed only using parametric methods with interactions of the linear function of year of birth before and after the reform. This is because in absence of information at the month of birth, using non-parametric estimation is likely to produce biased estimates as it needs to use a wider bandwidth around the cutoff to include more point estimates in the regression (Lee and Lemieux, 2010).

Figure 5: Discontinuity in the probability of having been conscripted by month of birth for men born 4 years before and after the abolition of conscription in Spain

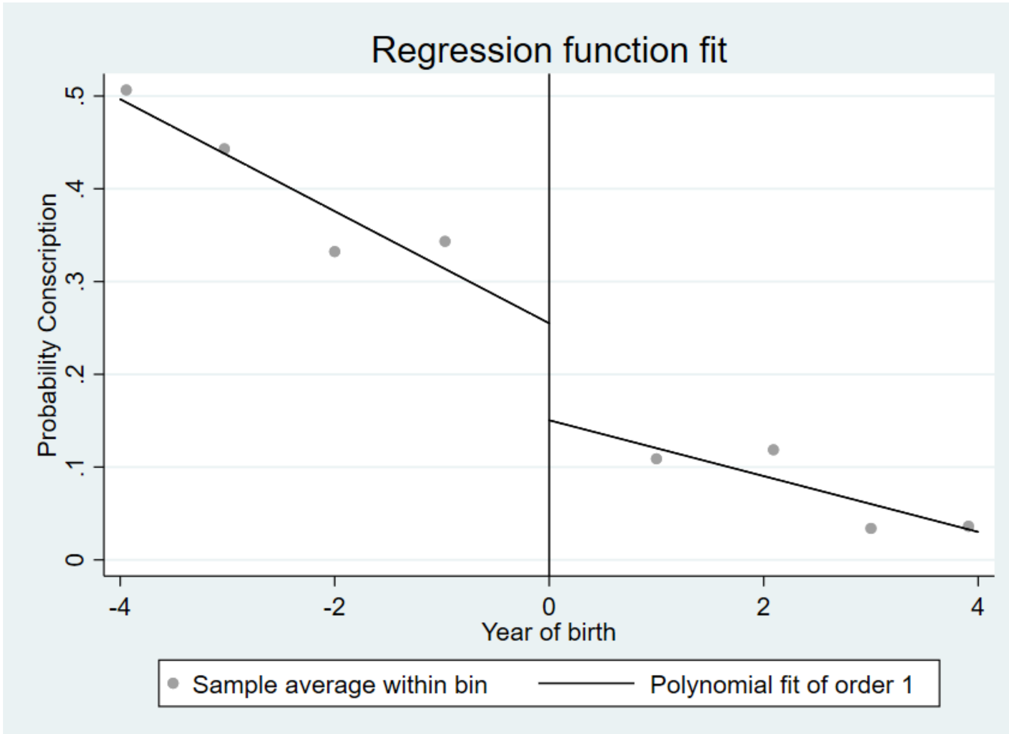

Table 13 presents the summary statistics for the Spanish data.

The first two columns present information on the whole sample of men for which military service information is available. The last two columns contain information for the sample used in the analysis and born in the four years before and in the four years after the abolition of military service (i.e. born between 1975 and 1982). These statistics show that the sample in the analysis is younger than the rest of the population, and that 44 percent of those in the sample of analysis were born after the change of reform. They are less likely to be right-wing, less likely to identify with Spain and less likely to have voted. Additionally, they are more educated than average for the whole population.

Table 13: Summary statistics of the individual characteristic in Spain: full sample of men and sample used in RD estimation.

	All		RD sample	
	(1)	(2)	(3)	(4)
	mean	sd	mean	sd
Conscription	0.58	0.49	0.28	0.45
Born Post 1979	0.19	0.39	0.44	0.50
Age	43.94	18.07	26.43	6.20
Ideology (1left-10right)	4.74	1.92	4.56	1.86
Rightwing 1/0	0.57	0.50	0.54	0.50
National identity (1-5)	3.22	1.07	3.06	1.06
Spanish Identity	0.82	0.39	0.77	0.42
Voted in last GE	0.81	0.39	0.72	0.45
No education	0.07	0.26	0.00	0.05
Primary edu	0.30	0.46	0.14	0.34
Secondary edu	0.42	0.49	0.63	0.48
University degree	0.21	0.40	0.23	0.42
Unemployed	0.13	0.34	0.17	0.38
N	11333	11333	2211	2211

Note: Summary statistics of individual characteristics. Data from the CIS survey on National Defence and Armed forces. Columns 1 and 2 show the summary statistics for the whole sample. Columns 3 and 4 show the summary statistics for the sample used in the analysis.

As in the main analysis, I provide graphical evidence of the discontinuity in the probability of having more right-wing ideology and of having voted in general elections (figure 18 in Appendix 2). The graphical evidence is in line with the causal effects identified in the regression analysis. Table 14 shows the results of this analysis and estimates equations 9-11 using Spanish data. Panel A in table 14 shows the results obtained when no controls are included in the analysis, while Panel B includes controls for year of the survey. All estimations are run using parametric methods and linear polynomial functions of the running variable interacted with the treatment dummy to allow for a change in slope of the function.

Table 14: The effect of mandatory conscription on voting participation, right-wing ideology and Spanish identity: linear parametric estimation with interactions.

	Voted in GE		Righwing		Spanish Identity	
	(1)	(2)	(3)	(4)	(5)	(6)
	FS	RF	FS	RF	FS	RF
		SS		SS		SS
Panel A						
Treated	-0.222 **	0.033	-0.164 ***	0.032	-0.164 ***	0.024
	(0.064)	(0.022)	(0.046)	(0.035)	(0.040)	(0.125)
Conscription		**-0.128 ****		**-0.182**		**-0.263**
		(0.056)		(0.146)		(1.048)
Controls	No	No	No	No	No	No
Panel B						
Treated	-0.167 ***	0.007	-0.167 ***	0.034	-0.167 ***	0.014
	(0.047)	(0.027)	(0.047)	(0.036)	(0.047)	(0.127)
Conscription		**-0.028**		**-0.192**		**-0.159**
		(0.097)		(0.144)		(1.204)
Controls	Year	Year	Year	Year	Year	Year
Window		4 Years		4 Years		4 Years
Pol. Order		1		1		1
N		1684		1695		720

Note: Each coefficient is estimated using a separate regression. Treated is a binary variable equal to 1 if the individual was born in or after 1979 and turned 18 in the year of abolition of mandatory conscription in Spain. Conscription is the predicted probability of having done military service estimated from the FS regression. Voted in GE is a binary variable equal to 1 if the individual states to have voted in the last general election. Rightwing is a binary variable equal to 1 if the individual reports to have a political ideology more rightwing than the median. Spanish Identity is a binary variable equal to 1 if the individual reports to feel more identified with Spain than with his region of origin. Panel A shows the results of the estimation when no controls are included, while panel B includes controls for year of the survey. All regressions are run using parametric methods and interacted linear function of the running variable, i.e. year of birth. Standard errors are clustered at the running variable level

Columns 1 and 2 in table 14 show the results for voting participation. The first stage and reduced form equation results are reported in the first row. The first-stage results show that men in the treatment group are on average 16 percentage points less likely[12] to have been conscripted relative to men born in the four years before the change in reform (column 1). The results of the second stage regression are presented in the second row of column 2. They show that the abolition of conscription decreases by 12 percentage points the probability of having voted in the last general election. The coefficient on voting participation has similar sign and magnitude compare to the German analysis. In Panel A the coefficient for voting

[12]Notice that the sign of the coefficients in the first-stage regressions is now negative as I am evaluating the probability of having been conscripted after the abolition of conscription, as opposed to the German case where I exploit the variation in this probability caused by the introduction of military service.

participation is significant at the 5 percent level. However when controlling for year of the survey this significance disappears and the magnitude of the coefficients drops to 2 percentage points, but remains negative. Columns 3 and 4 in table 14 show the results for right-wing ideology. These coefficients have the same sign and magnitude of those presented in the main analysis. In fact, the abolition of conscription decreases by 18 percentage points the probability of having right-wing ideology above the median.

Lastly, columns 5 and 6 show novel evidence of the effect of conscription on Spanish identity. These results indicate that conscription has a positive effect (of between 15.9 and 26.3 percentage points) on the probability of feeling identified with Spain, but again the effect is not statistically different from zero. The negative sign of the effect of the abolition of conscription, however, is in line with the hypothesis that conscription might reinforce national identity.

2.8 Conclusions

This paper analyzes the effect of conscription on political ideology, voting participation and national identity.

Following from a renewed policy interest in re-introducing conscription, the aim of this paper is to provide rigorous empirical evidence of the causal effect of conscription in changing political attitudes and national identity. This analysis aims at informing policy makers that propose such reform as a way to promote civic duty and national cohesion among the youth. The economic costs of conscription are non-negligible and a large literature identifies both public and private costs associated with peace- and war-time conscription. In light of these costs, this paper aims at evaluating the effectiveness of conscription in achieving its proposed goal.

The analysis is conducted using a regression discontinuity design. Using individual level data from West Germany (and Spain), it exploit the introduction (and abolition) of conscription as an exogenous shock to the probability of being conscripted for men turning eighteen right before and right after this change in reform. I use this exogenous shock to evaluate the causal treatment effect of conscription on a series of outcomes measuring political ideology and voting participation. In the main analysis, I use the introduction of conscription in West

Germany in 1955 to define the treatment group as men born in or after the 1 July 1937 and turning eighteen at the time of the reform, and the control group as men born before this cutoff date. Restricting the analysis to cohorts born in the 48 month before and after this cutoff, I obtain comparable treatment and control groups and evaluate the long-term effect of conscription on political ideology and voting participation.

The key findings are that conscription does not affect political ideology nor voting participation. The effect on right-wing ideology is positive, as well the one on the probability of having voted in parliamentary elections. However, these effects are not statistically different from zero and the coefficients are imprecisely estimated. These findings are robust to a series of robustness checks, including performing the analysis on the effect of the abolition of conscription in Spain on similar outcomes. Exploiting the abolition of conscription in 1997, I compare men turning eighteen in that year (and no longer eligible to be conscripted) to men turning eighteen right before the cutoff year to evaluate the treatment effect of conscription on ideology, voting and national identity. I find similar results to those presented in the main analysis for right-wing ideology and voting participation. I also find a positive and statistically insignificant effect of conscription on Spanish identity.

Before concluding, I would like to draw the attention of the reader to the main limitation of this empirical analysis. This is the estimation of the second stage coefficients measuring the effect of conscription on the outcomes of interest. In most estimations the coefficients show large confidence intervals leading to an imprecise estimation of the effects and in some specification to an ambiguity regarding the sign and the magnitude of the effects. The coefficients showing the highest variation are those presented in the most demanding specifications presented in the main analysis and, especially, the estimates on rightwing ideology. I address part of this issue in section 2.7 when presenting the robustness checks. In particular, restricting the bandwidth closer to the cutoff date and using only linear polynomials in the estimation allows me to restrict the lower and upper bounds of the effect to be very close to zero and positive, rather than ranging from a negative to a positive effect. Still, the effects estimated in this analysis do not allow me to say that the policy has no effect on the outcomes considered, rather that this effect cannot be precisely estimated. Oppositely, the estimation of the first stage and of the validity of the instrument is very strong and consistently estimated throughout the analysis leading me to conclude that this issue is not driven by data issues or by lack of statistical power. The consequence of the lack of precision in the estimation of the second-stage coefficients affects the type of policy recommendation we can put forward from this analysis. In fact, I am only able to say that the policy does not have a statistically significant effect rather than this effect is equal to zero.

To conclude, I do not find statistical significant evidence of a causal effect of conscription on the formation of political preferences nor on voting participation. Although some of the mechanisms suggested by the literature and by policy-makers are corroborated by the sign of the effects, none of these is statistically different from zero nor precisely estimated to have positive effects. Hence, given the costs involved with conscription, both in terms of government spending and in terms of individual labor outcomes|, this paper does not find evidence proving the effectiveness of conscription in shaping individual civic participation. Thus, based on these findings, I would not be able to recommend the re-introduction of conscription as a cost-effective policy to achieve this goal.

3 The effect of migration on household consumption: evidence from rural Ethiopia.

3.1 Introduction

Migration is at the center of the political and development agenda for the coming decade (Sustainable Development Goals 2030, see here). The latest UN estimates show that there are around 272 million international migrants in the world, with around 40 percent of them migrating in the global South (UNPD 2019, see here). Moreover, these numbers do not include internal migrants, more difficult to count, and estimated to be three times as much when using a conservative estimate (Skeldon, 2018). Most migration, especially in the form of labor mobility of the poor, takes place within and between developing countries (Mendola, 2012). This is also true in African countries, simultaneously serving as source and hosts to large shares of migrants (Lucas, 2006).

Individuals migrate to improve their economic condition, moving to a location where they expect to receive higher incomes (Harris and Todaro, 1970) and to countries where they can have better economic opportunities (Borjas, 1989). However, migrant's income maximization motives are not enough to explain the decision to migrate and do not account for the effects of migration on migrant-sending households. In fact, as migrants continue to send remittances back home for extended periods of time after they move, a large literature investigates the direct and indirect effects that this income channel has on origin household welfare (Mendola, 2012).

Seminal work in this field (Stark and Bloom, 1985) has led to developing The New Economics of Labor Migration, which focuses on the role played by origin households in the decision of one of their members to migrate and highlights the role of migration as a way for the migrant-sending household to better manage risk, diversify income and alleviate liquidity constraints. These positive effects on origin household welfare are often linked to the inflow of remittances sent back home by migrants (Stark and Levhari, 1982). In rural contexts, receiving remittances might reduce poverty as it allows the household to insure itself against the risky nature of agricultural production (Rosenzweig, 1988; Yang and Choi, 2007) by both increasing household liquidity (Katz and Stark, 1986; Yang and Martinez, 2006) and by increasing household's productive investments (Yang, 2008).

Empirical research shows evidence in line with these predictions and calculates the welfare gains from migration to be considerably larger than other development interventions (Clemens and Ogden, 2014), with direct positive effects on origin household's income and consumption (Gibson and McKenzie, 2014, Bryan et al., 2014), human capital investment and savings (Yang, 2008; Clemens and Tiongson, 2017). Moreover, migration can have effects on the household welfare that work beyond the income channel, e.g. transmission of social norms and of health practices (McKenzie and Rapoport, 2007;Sasin and McKenzie, 2007; Clemens et al., 2014). A large part of this empirical literature focuses on the effect of migration on migrant's welfare, on economic outcomes in host countries and on the effects of international migration (Clemens et al.), with several studies looking at the consequences of migratory flows from Mexico to the U.S. (Mendola, 2012). Less is known on the effects of migration, and in particular internal migration, on origin households in developing countries (Klugman, 2009). This paper aims at contributing to this literature by providing evidence of the effect of migration on household consumption in rural Ethiopia.

The existing empirical evidence on Ethiopia confirms that migration has been effective in increasing migrant's welfare and in reducing poverty (de Brauw et al., 2017; Blunch and Laderchi, 2015). This is of particular relevance as Ethiopia is one of the poorest countries in the world, with most of its poorer population living in rural areas. However, and notwithstanding these positive effects, the migration rate in Ethiopia is considerably lower than in the rest of the region (De Brauw and Mueller, 2012). To explain this, de Brauw et al. (2017) conclude that leaving one's household is particularly costly in this context and that this prevents larger shares of households to have access to migration (Kosec et al., 2017; De Brauw and Mueller, 2012). Reforms aimed at lowering such costs might allow more households to invest in migration as a income diversification strategy (Clemens, 2010). Moreover, to the best of this author's knowledge the literature on migration in Ethiopia investigates the effect of migration on migrant's welfare (de Brauw et al., 2017), its relationship with land availability and youth unemployment (De Brauw and Mueller, 2012; Mueller et al., 2018), but little is known on the direct effect of migration on household welfare. This paper aims at contributing to this gap in the literature.

The analysis uses panel data for 1200 households in the Amhara, Oromyia, Tigray and SNNP regions of Ethiopia between 2014 and 2018. By using variation over time and a rich set of covariates, I investigate the effect of having at least one migrant in the household on total, food and non-food consumption. The definition of migration used in the analysis follows from the challenges in measuring this phenomenon. In fact, households might have very different migration experiences, each producing different effects on consumption: e.g. households might have more than one migrant, have migrants returning to the household between waves, have only international or both internal and international migrants etc.. To summarize this information in a more generalized measure, I define the treatment to be whether the household has at least one migrant away.

As the data show enough variation over time in this migration status and the attrition rate in the panel is very low, I conduct the analysis separately for households treated in both waves or in just one period and compare them to households that never had a migrant. This allows me, first, to reduce the heterogeneity in the treatment group, and, second, to show evidence on the effects of migration on consumption for three types of households: households with at least a migrant in both waves, with a migrant only in 2014 and with a migrant only in 2018. This strategy partly overcomes the initial loss of information of defining migration using a more generalized measure, as it allows me to partly capture the effect of different migration experiences.

I find that having at least one migrant away affects household food and non-food consumption. I find that household per capita consumption decreases by around 20 percentage points for households treated in both waves and for those treated only at follow-up when compared to households that never had a migrant. I also find that overall household-level consumption increases for these two groups. An increase in household size, number of adults and number of men emerges as a possible explanation of these results. This finding is in line with previous empirical evidence from Ethiopia showing that it is usually more productive and young head of household's children who migrate. This leads to a change in the household composition and a change in the distribution of labor supply within the origin household (Mueller et al., 2018).

I combine this explanation with additional qualitative evidence on the changes in migrant's destination and in the flow of remittances sent back home. In particular, I present qualitative evidence showing a decrease in international migration, in the number of migrants away from the household and in the probability that the household receives remittances and goods. This evidence might offer an additional explanation for the negative effect on per capita consumption. In particular, all these mechanisms might reduce the amount of remittances sent back home, especially by international migrants (sending five to ten times larger remittances amounts than internal migrants). This possible mechanism is of particular relevance in Ethiopia, where barriers to international migration are generally very high (Mueller et al., 2018). Moreover, internal migration has likely become riskier in the period of analysis, with migrants migrating closer to the village and for shorter periods. This has mostly been due to an increase in inter-ethnic conflicts that plausibly not only made migration, but also sending remittances back home riskier (Tsegay and Litchfield, 2019).

The paper contributes in several ways.

First, it contributes to the literature on the relationship between migration and migrant-sending households welfare. Seminal theoretical work on migration focused on the idea that individuals move in search of a better life (Clemens and Ogden, 2014). Following from the Harris-Todaro model (Harris and Todaro, 1970), part of the literature has focused on the individual decision to migrate to maximize the migrant's expected income. Another strand of the literature started understanding migration as a more complex phenomenon motivated not only by wage differentials, but by several market failures at origin (Stark and Bloom, 1985). Empirical evidence shows that such investment might be highly remunerative for the origin household and can cause increases in income and consumption (Gibson and McKenzie, 2014), higher investment in human capital (Yang, 2008), higher food and non-food expenditure (Bryan et al., 2014), as well as increases in spending and savings (Clemens and Tiongson, 2017).

However, Gibson et al. (2011) show that migration causes a short-run decrease in welfare when the loss of labor income caused by the migrant's absence does not compensate the gains from migration, e.g. increase in remittances. Negative effects of migration might also depend on a change in the distribution of decision making within the households, on shocks to knowledge transfers and on negative effects on mental health caused by household separation. These factors lead the authors to conclude that "the overall impact of migration on the welfare of remaining family members is theoretically uncertain" (Gibson et al., 2011). Following from this mixed evidence, additional research needs to be carried out to evaluate the effect of migration on household outcomes in different contexts.

Second, this paper contributes to the empirical literature evaluating the effect of migration on household welfare using panel data. This literature has started growing faster in the last two decades as more data have become available (Clemens et al., 2014). Traditionally, empirical studies on migration were based on cross-sectional and correlation analysis. This is mostly due to the difficulty of studying this topic using experimental methods and from the fact that migrants are not a random sample of the population (Taylor and Martin, 2001). More recent literature has started overcoming these challenges thanks to the use of experimental and quasi-experimental methods (McKenzie and Yang, 2010; Gibson et al., 2011), and to the growing availability of panel data (McKenzie et al., 2010; Beegle et al., 2011). Although still far from the ideal experimental setting, panel data and propensity score matching are commonly used in this literature to attenuate the concerns of unobserved heterogeneity and selection into treatment (Beegle et al., 2011; Gibson and McKenzie, 2014; de Brauw et al., 2017).

Third, this paper contributes to the literature on the effects of migration in Sub-Saharan Africa. This literature has traditionally been limited by the lack of adequate data on migration (Lucas, 2006). Empirical evidence from South-East Asia (Yang, 2008; Gibson and McKenzie, 2014; Clemens and Tiongson, 2017) suggests that migration has positive effects on the welfare of migrant-sending households. As more than half of the extreme poor live in rural areas of Sub-Saharan Africa (World Bank, 2019), understanding how migration can alleviate poverty by increasing household consumption in the region is key in the development agenda.

Fourth, the paper investigate the effect of different types of migration on household consumption. It shows that migration has different impacts depending on its duration, contributing to the literature on temporary migration (Clemens and Tiongson, 2017). In line with some of the evidence presented in this paper, Gibson et al. (2011) find that migration might have a negative short-run effect on household welfare due to the absence of the migrant and the loss of his input to household production, especially in rural contexts.

Moreover, this paper contributes to the growing literature on internal migration in developing countries. Internal migration might differ remarkably from international re-location, as the former tends to be shorter and seasonal. Moreover, remittances sent from abroad are considerably larger than those sent from internal migrants. Empirical studies focusing on internal migration provide evidence for developed (Molloy et al., 2011) and for South-Asian countries (Bryan et al., 2014; Bryan and Morten, 2019; Morten, 2019), while less is known on internal migration in Sub-Saharan Africa. However, most of the migration in Africa happens internally and in the form of rural-urban relocation (see here). In particular, Ethiopia has a large potential for internal migration as a large share of its population still lives in rural areas and non-agricultural employment opportunities are rising as the country grows (de Brauw et al., 2017). This paper contributes by showing evidence of both internal and international migration in this context.

Lastly, it provides evidence on the possible mechanisms explaining the results. By showing that migration is accompanied by a change in household composition, this paper provides additional evidence of the trade-offs faced by Ethiopian households investing in migration. This is in line with previous literature on Ethiopia showing that household might decide not to migrate to keep their right to use land (De Brauw and Mueller, 2012) and, when they do, this results in a change to its labor supply (Mueller et al., 2018). These costs plausibly drive the low migration rate in the country. Moreover, this paper provides qualitative evidence on changes in migration patterns over the period of analysis in Ethiopia. The country experienced several shocks that directly, e.g. a ban on migration to the Middle East for domestic workers, or indirectly, e.g. increasing violence due to conflicts and droughts., affected migration. This evidence highlights an increase in the risks faced by migrants, which might translate negatively on origin household's welfare. Moreover, this evidence aims at complementing the analysis based on a rough measure of migration and highlights the need to study migration as a more complex phenomenon (Mendola, 2012).

3.2 Related literature

First, this paper contributes to the literature on the effects of migration on origin household's welfare. Migration has long been studied as a way to access better economic opportunities. The Harris-Todaro model led the way in this field by demonstrating that individuals decide to engage in rural-urban migration based on their own expected income maximization problem

(Harris and Todaro, 1970). This seminal work in the micro-economic study of migration overcame the limitations of the neoclassical Lewis's two-sectors model (Lewis, 1954), formalized by Ranis and Fei (1961). According to Ranis and Fei (1961), migration is driven only by wage differentials between rural and urban areas (and potentially between countries) and, as such, it should cause these differentials to tend to zero in equilibrium. Several limitations to this model, i.e. persistent urban unemployment in low income countries and the persistence of wage differentials across rural areas and between sectors (see Taylor and Martin, 2001 for a review), make the Harris-Todaro model still the benchmark to understand migration decisions from a micro perspective.

The presence of continuing interactions between the migrant and its origin household motivated the New Economics of Migration (NELM) literature (formalized in Stark and Bloom, 1985 and Stark, 1991). In this model, households jointly decide to invest in migration not only to maximize income, but, and mostly, to minimize risk, diversify earnings and loosen financial constraints. This is particularly relevant in rural contexts, where migration might insure the household against negative agricultural shocks (Rosenzweig, 1988). Thus, the NELM motivates migration as a household-level decision driven by rural areas market failures and asymmetries (Taylor and Martin, 2001). As these are likely to be more severe (and the gains from migration are likely to be larger) most of the NELM literature is focused on developing countries. Following from this model, a growing literature has looked at the relationship between migration and development. Migration might reduce poverty and positively impact household welfare via a direct income effect, i.e. remittances, but can also affect household labor supply, decision-making dynamics, the transmission of health practices and of social norms (Sasin and McKenzie, 2007).

More recently, and thanks to increase in data availability, research has started investigating empirically the presence and magnitude of these effects (Mendola, 2012), with a particular focus on the role of remittances (Clemens and Ogden, 2014). Empirical studies have proved the existence of high returns from migration. Yang (2008) finds that a positive shock to the value of remittances increases investment in human capital and entrepreneurship in origin households, as well as causing a 25 percent increase in income. Bryan et al. (2014) find that migration increases household food and non-food expenditure by 30 to 35 percent and improves household's calories intake in rural Bangladesh. Gibson and McKenzie (2014) provide experimental evidence on the effects of seasonal work migration on household consumption, income, savings and assets, showing an increase in income of around 35 percent for households with at least one labor migrant away. Similar positive effects on households expenditure on human capital investment and savings are documented in Clemens and Tiongson (2017). This paper relates to this literature as it provides novel evidence on the relationship between migration and household consumption in rural Ethiopia.

It should also be noted that the experimental and quasi-experimental evidence on the effects of migration on household welfare is still scarce and has started growing only in the last decade. If "migrants are not a random sample of the population" (Harris and Todaro, 1970), the main issue in any empirical analysis of migration is selection into treatment. The increasing availability of panel data has helped addressing issues of unobserved heterogeneity between the treated and control groups and has been crucial in providing evidence on the effects of migration when experimental settings are not available to the researcher (McKenzie and Yang, 2010; Beegle et al., 2011; Gibson and McKenzie, 2014). Thus, this paper contributes further to the literature that aims at estimating the treatment effect of migration using panel data (Sasin and McKenzie, 2007; McKenzie et al., 2010).

Second, this paper relates to a growing literature studying migration in low-income countries, where a large part of individual's relocation happens within national borders (Bryan et al., 2014). A large amount of evidence exists to explain international and high-skilled migration. This literature aims at explaining the brain drain phenomenon, and how immigration affects wages in receiving countries (Bhagwati and Hamada, 1974; Gibson and McKenzie, 2012). When looking at the effect of migration on origin household, the existing empirical findings are largely focused on permanent international migrations (Bryan et al., 2014; Clemens et al., 2014). Less is known on the effects of internal and seasonal migration on household welfare, although these might differ considerably from the long-run gains of having a member migrating permanently and to another country.

The amount of remittances sent back home is considerably lower. Second, the duration of migration might be shorter and seasonal, not giving the household enough time to benefit from the migrant's higher wages and network. Empirical evidence shows that internal migration allows migrants to benefit from higher wages and increases aggregate productivity (Bryan and Morten, 2019), it allows households to rely less on informal risk-sharing (Morten, 2019) and this reduces the probability of informal-risk sharing within the village network (Meghir et al., 2019). At the same time, internal migration is usually seasonal and can be costly for the household. Moreover, migrants might underestimate the income in urban areas and under-invest in migration (Baseler, 2018).

Following from this discussion, this paper contributes to the literature on the short-term effects of migration and on the effects of temporary migration. When looking at the short-run effects of migration, Gibson et al. (2011) discuss that migration might have mixed effects on origin household welfare. It might cause a reduction in non-remittance agricultural income because of the absence of the breadwinner, it might affect the decision making dynamics within the household, labor supply decisions, and it might cause mental health problems due to the separation from a family member. In line with this prediction, Gibson et al. (2011) find that in the short-run income per capita fell by 22 to 25 percent in Tongan households as a result one of their members migrating. Although these households receive more remittances, these are not enough to compensate for the loss in labor earnings.

These negative effects are not corroborated in Clemens and Tiongson (2017). The authors evaluate the effect of having a member of the household migrating temporarily overseas, highlighting the importance of studying temporary migration separately from permanent household separation. They find that temporary overseas migration has very large positive effects on household welfare. It causes household expenditure to increase by approximately one third (estimate in line with Beegle et al., 2011; Bryan et al., 2014), savings to go up and expenditure in health and education to go up by hundreds of percent. Moreover, Yang (2008) finds that shocks to remittances values do not affect household consumption, but increase the probability of carrying out productive investments. Following from these results, Yang (2011) concludes that the use of remittances might change over time. While at the beginning households might use remittances to increase their consumption level, later shocks to the value of remittances might result in an increase in entrepreneurial activities. Given these mixed evidence, this paper contributes by showing evidence of the effect of different migration experiences, i.e. of households having at least a migrant away in both waves or just in one.

Third, this paper contributes to a growing literature on migration in Sub-Saharan Africa. The existing literature on the effects of migration on household welfare focuses mostly on migration from South and Central America to the U.S. (McKenzie and Rapoport, 2007), from South-East Asia (Yang, 2008), but less is known on Sub-Saharan countries (Lucas, 2006). Some studies have looked at the effect of migration on migrant's income and welfare. Beegle et al. (2011) find that migration increased household consumption by 36 percentage points in Tanzania. Similarly, de Brauw et al. (2017) use panel data to estimate the effect of internal migration on migrant's consumption and income in Ethiopia, finding that income increases, non-food consumption increases by 165 percent for migrants relative to non-migrants and that individuals who migrated between 2004 and 2009 have better diets. However, and to the best of this author's knowledge, no evidence exists on the effects of migration on origin household's welfare, rather than migrant's, in Ethiopia.

3.3 Context

Ethiopia, with a per capita annual income of around $780, is one of the poorest countries in the world and in the Sub-Saharan region (World Bank, 2019). It is also the second most populous country in Africa after Nigeria and it has been experiencing a decade of steady economic growth. For this growth to be sustainable, reforms aimed at reducing poverty further need to be implemented (World Bank, 2019). Migration constitutes one of the possible ways to improve the economic conditions of rural households in Ethiopia.

Thanks to its location in the Horn of Africa, Ethiopia has a long tradition as both a country of origin and of destination (Marchand et al., 2017), with many migrants transiting through to access the route to the Middle East from Djibouti. In the period between 2014 and 2018, several migratory phenomena have interested the country. Conflicts in neighboring countries (South Sudan, Eritrea) and ethnic-related internal conflicts (escalating to violence in 2018) forced masses of refugees to enter Ethiopia and made internal migration riskier. In 2015 and 2017, the country experienced severe droughts which have been found to increase work related migration for men and decrease marriage related migration for women (Gray and Mueller, 2012).

Moreover, work-related emigration, particularly in the form of domestic work, has been affected by a ban imposed by the Ethiopian government on migrants traveling to the Middle East (following from the increase in reports of violence on the workplace against women in Saudi Arabia). This is also evident in our data, where I find that 30 percent of returning migrants who were in the Middle East returned because they were deported. In general, the Ethiopian government has traditionally been adverse to international migration (Mueller et al., 2018) contributing to the low out-migration rates in the country and to the decline in international migration documented in our data. All of these factors contribute to the high costs of migration for financially-constrained households in the country, making it more and more difficult for them to invest in migration as a way to improve their living conditions.

The context of Ethiopia is of particular relevance for two reasons. First, the majority of poor people live in rural areas. For this reason, de Brauw et al. (2017) highlight that given the magnitude of their results (an increase of 145 percent in non-food consumption), policies lowering the costs of migration could be an effective strategy in facilitating poverty reduction. Second, and in line with the NELM puzzle of "why are not more people migrating?", migration is accompanied by high costs for the household. In this regard, it is worth discussing in more detail the role played by land regulations.

The Ethiopian economy is largely dependent on agriculture. Thus, rural households depend on the right of using land for their livelihood. Land regulations are particularly relevant in this context and directly affect migration decisions (De Brauw and Mueller, 2012). This is both because of strict regulations on land transferability and of the increasing scarcity of land for younger generations (Bezu and Holden, 2014). Land is publicly owned in Ethiopia and redistribution of land was common before the implementation of land certificates allocating to each household the right to use a given area of agricultural land (Deininger et al., 2008). These land rights have traditionally been enforced weakly, causing households to avoid sending out migrants in fear of expropriation (De Brauw and Mueller, 2012). Although progress has been made in reinforcing land rights (Deininger et al., 2008), strict restrictions in renting out and exchanging land still affect migration patterns in Ethiopia. In particular, out-migration rates in Ethiopia are considerably lower than in other Sub-Saharan countries and are usually higher for men and women between 15 and 24 years of age (Mueller et al., 2018). This results partly from the Federal Land Use Law which precludes the household head to migrate in order not to lose her land (Bezu and Holden, 2014).

Thus, as the household head cannot migrate and the possibility of renting out and selling land are very limited, it is often the head's children who migrate out of the household (Mueller et al., 2018). Moreover, the amount of land available to each household depends on agreements often based on household size. Thus, if one of the household members migrates, it is common for the household to substitute him with a member of the extended family or hired workforce.

This pattern is also accentuated by an increasing scarcity of land. As more and more land has been allocated to existing households, the country is experiencing growing rates of landlessness among the younger generations, who are then forced to migrate in search of better livelihood (Bezu and Holden, 2014). Part of the findings of this paper, namely the negative effect on consumption and the change in household composition, are in line with some of the mechanisms discussed so far. In particular, the regulations on land and the out-migration of more productive household members are plausible explanations of why per capita consumption decreases for households having a member migrating between waves and of why the composition of the household changes over time.

3.4 Data

The data used in this analysis were collected by the Ethiopia country team of the Migrating out of Poverty Research Programme Consortium. This project aimed at analyzing the effect of migration on migrant-sending household's outcomes. Hence, the selection of households at baseline was carried out in a way such that 65 percent of the households in each of the sixteen sampled kebeles had at least one of the household's member away as a migrant, i.e. living outside of the village for three months or more in the past ten years. The dataset is a panel of around 1200 rural households of the four regions of Amhara, Tigray, Oromyia and Southern Nations Nationalities and People (SNNP)[13].

The survey was conducted for 1207 rural households in September 2014 and it includes both information on households economic situation and demographic information on each member of the household (both for all individuals in the roster and only for migrants). The follow-up

[13]Although not statistically representative at the regional level, these data are qualitatively representative of the migration flows at the aggregate level.

survey was conducted in the months of September and October 2018. The attrition rate of the panel is surprisingly low (less than 1 percent) for the context of Sub-Saharan Africa. In fact, the second wave of data is available for 1202 of the households interviewed at baseline[14]. The team collected household-specific tracking data (using Computer-Assisted Personal Interviewing) in the initial phase of fieldwork, and asked the questionnaire respondent's consent to be interviewed at follow-up. This ensured that the respondents engaged with the research project and made them aware that the team would return for a follow-up interview round[15].

The key information for this analysis is on the migration status of each member in the household. For each household, I have information on whether any member is away as a migrant. A migrant is define to be a person who has lived outside of the kebele (village) for three months or more, in the past 10 years[16]. Thus, for each household I define a binary treatment variable to indicate whether at least one of the members is a migrant. The treatment variable is defined at the household level. As I will discuss in the next section, I will define this treatment variable for three treatment types depending on whether they are treated in both waves or only in one of the periods. The control group will be defined as those households who had no migrants both at baseline and at follow-up.

The dataset comprises three sources of information: at the individual, household and community level. Individual level information for non-migrants and migrants includes questions on age, gender, marital status, employment status, higher level of education achieved, number of children, mother tongue and religion. Information on income, consumption and remittances is, instead, available at the household level. Demographic characteristics are calculated at the household level for the following variables: the ratio of dependent household members (calcu-

[14]Notice that this is peculiar to the context of Ethiopia and is plausibly possible because of the rigid land regulations in Ethiopia. Land is publicly owned in Ethiopia and, although, the land certification reform established rights to use the land for households, the migration of the household head leads to the loss of such right.

[15]Each household received a consent agreement in 2014 and was asked to sign it. The team collected evidence showing that the respondents kept such written agreements and presented them to the team when surveyed back in 2018.

[16]This definition builds on the literature aiming at measuring migration in a consistent and comparable way. There are several dimensions that should be taken into consideration. First, the period of migration should be long enough for the household to see economically meaningful changes. Second, the distance from origin household should be taken into account. Third, it is usually more useful to refer to a relatively recent fixed-period of time during which the migration happen (as opposed to one's lifetime). Further discussion on the definition of migration and its measurement in survey data can be found in Bilsborrow et al. (1984)and Bilsborrow et al. (1997).

lated separately for members below the age of 15 and above the age of 65); the ratio of men and the ratio of adults in the household who indicate to have received no schooling. Information on household infrastructure and hectares of land owned is available for each household. In particular, I will use information on whether the household has access to electricity, to piped water in the dwelling or in the village and to private toilet facilities, as well as the hectares of agricultural and homestead land owned.

Additional information is available for community level characteristics at baseline. For each of the sixteen kebeles, the dataset includes information on population and number of households, number of shocks (droughts, epidemics, flooding) and investments in infrastructures (training centers and schools, health posts, roads, water supply) affecting the community in the past twenty years, as well as distances to main destinations of interest (region's capital, main road, bus stop, micro-credit facility).

Although the dataset includes a section on the use and value of remittances receives, this will not be used in the analysis as this information is only available for households with migrants. However, in the summary statistics section I provide evidence on the main use of remittances and on the value of goods and remittances received by households with migrants in both waves and in only one of the two periods. I construct a binary variable equal to one if at least one member of the household sends remittances ($ReceivesRemittances$) or goods ($ReceivesGoods$) back home. Additionally, I construct two variables indicating the monetary value of remittances and goods received by the household. I do this by summing the value of remittances and, separately, of goods received by the household from all the members away as migrants and adjusting such values for inflation. The questionnaire asks the head of household to indicate the value of remittances (and goods) received in the past twelve months, but I report these values at the monthly level for ease of comparison with the consumption figures. I also present qualitative evidence on the use of remittances. I report the main use of remittances as indicated by the head of household for each of the following categories: everyday consumption, land and agricultural productivity, home improvement, savings, health and education costs.

Lastly, I provide descriptive evidence of a change in the destination of migrants over time. Exploiting information on whether household members migrate internally or internationally. In particular, I classify household to have only internal migrants, only international migrants or both.

I now turn to the description of the outcome variables used in the analysis. The household data contains information on weekly food consumption, monthly non-food consumption and annual income received by the household. I construct a measure of annual income[17], but, although useful when discussing possible mechanisms, I do not use this variable as the main outcome of interest. Given the context of analysis, characterized by high levels of poverty and a strong dependence on agricultural production, together with the availability of rich information at the weekly and monthly level, the analysis focuses on consumption measures to measure household welfare. I discuss these measures in more detail in the following section.

3.4.1 Consumption measures

The main outcome variables are derived from a detailed section on household expenditure in consumed goods. Information is available for a list of food and non-food items from which I construct two separate measures of monthly food and monthly non-food consumption.

In the non-food consumption section, each household is asked to report the amount spent in the last month for thirteen non-food items, and the amount spent in the last year for another fourteen non-food items consumed less regularly (clothes, furniture, ceremonial expenses, education, healthcare). For each household, I construct a non-food consumption value to be the natural logarithm of the sum spent for all these items in a period of one month.

Constructing food consumption values proved less straightforward as for each of the twenty-five items listed in the household questionnaire, each household could report to have purchased, produced or received as a gift said item. Thus, while quantity consumed in the past week is available for all items in the list, information on the amount spent for each item is only available for purchased ones. I construct a median price per unit at the kebele level for each of the items. Then, I use this median kebele-level price as an indication of amount spent per unit for items that were produced at home or received as a gift. I then construct the food consumption variable by taking the natural logarithm of the sum spent by each household in

[17] I do this using information on cash income received from eleven different sources (non-agricultural and agricultural waged work, trade and business, renting out assets, government benefits, payments from NGOs, remittances, gifts from relatives and other members of the community). The variable $AnnualIncome$ presented in table 21 is the natural logarithm of the sum of the values reported by the household for each category. Values for the 2018 survey are adjusted for inflation and expressed in 2014 Birr.

a week and multiply this value by four to obtain a monthly-level figure.

Total monthly household consumption is defined as the natural logarithm of the sum of food and non-food consumption expressed in monthly values. All per capita consumption figures are constructed by dividing household consumption by household size. All consumption values are in real terms and expressed in 2014-Birr. I adjust prices using regional consumer price indices, constructed separately for food and non-food items and made available by the Central Statistical Agency of Ethiopia (for full tables of the Ethiopian CPI, see here).

3.5 Methodology

I am interested in estimating the effect of migration on household consumption.

There are several challenges in estimating the treatment effect of migration. These challenges are likely to be unaccounted for if looking at the correlation between our measure of migration, i.e. having at least one migrant in the household, and household monthly consumption using a simple OLS regression. In fact, this regression is likely to produce biased estimates due to unobserved heterogeneity and selection issues (McKenzie et al., 2010). For example, it might be that poorer households with lower levels of consumption are more likely to migrate, but it is also possible that having at least one migrant in the household causes consumption to increase because of remittances. However, by just looking at the correlation between migration and consumption I could not distinguish between these two mechanisms.

I can overcome some empirical challenges thanks to the structure of the data. By using a difference-in-difference specification I will be able to control for time-invariant unobservable differences between the treatment and control group. This will allow me to isolate the effects of omitted variables affecting consumption provided that their effect is constant over time. Given the panel structure of the data, I can also include household fixed effects to account for household-specific time invariant unobservable characteristics. I present the results of this second specification in the robustness checks. Using a difference-in-difference estimator is a standard practice in the literature looking at the effect of migration on consumption (Gibson and McKenzie, 2014; de Brauw et al., 2017; Beegle et al., 2011).

Similarly to (Gibson and McKenzie, 2014), I define migration has having at least one migrant in the household. I follow the literature (Beegle et al., 2011; de Brauw et al., 2017) in evaluating the effect of migration on consumption by estimating the following equation:

$$Consumption_{h,t} = \beta_0 + \beta_1 Treated_h \times Wave_t + \beta_2 Treated_h + \beta_3 wave_t + \epsilon_{h,t} \tag{12}$$

where $Consumption_{h,t}$ is the natural logarithm of monthly consumption for household h. $Wave_t$is a binary variable indicating the wave of the survey.

$Treated_h$ is a binary treatment variable equal to one if the household has at least one migrant in the household. One of the contribution of this paper is in the definition of such treatment variable. In order to exploit the full information in the dataset and to provide evidence on how the effect of migration differs by migration experience, I define three treatment types:

- households with at least one migrant in both wave ($AlwaysTreated_h$),
- households with no migrant at baseline and with at least one migrant at follow-up ($OnlyWave_2$),
- and households with at least one migrant at baseline and none at follow-up ($OnlyWave_1$).

I am able to study these effects thanks to consistent variation in the treatment group between the two waves. In fact, there is a 43 percent of the households who were treated in wave one are no longer treated in wave two (N=473), a 12 percent of households with at least one member migrating between waves, and 27 percent who had at least one migrant in both waves (N=327). $AlwaysTreated_h$, $OnlyWave2_h$ and $OnlyWave_1$ are binary variables equal to one if the household has at least one migrant in both waves, only in wave two, or only in wave one, respectively, and equal to zero if the household never had migrants. Notice that the definition of these three treatment types is highly dependent on the structure of the data.

Being households the sampling unit in the panel, I define the three treatment types according to whether the household has at least one member away in each period or in both waves. Although this is arguably a simplified measure of migration, it is the only viable definition consistent with the data structure. Few things should be kept in mind when interpreting the results. First, $AlwaysTreated_h$ is a binary variable equal to one if the household has at least a member away in both waves. This does not necessarily mean that the same individuals are away in both waves[18], nor that the migrants were away for the full period between wave one and wave two. Second, $OnlyWave_1$ is equal to one if the household had at least one migrant in wave one, but none at follow-up. This does not necessarily mean that the migrant has returned to the household[19]. Third, $OnlyWave_2$ is equal to one if the household has at least one migrant at follow-up and had none at baseline. This might be due to both migration of one household member between waves and to the migration of a new member of the household[20].

I estimate equation 12 separately for each of these three treatment groups, as follows:

$$Consumption_{h,t} = \gamma_0 + \gamma_1 AlwaysTreated_h \times Wave_t + \gamma_2 AlwaysTreated_h + \gamma_3 wave_t + \eta_{h,t} \tag{13}$$

$$Consumption_{h,t} = \delta_0 + \delta_1 OnlyWave2_h \times Wave_t + \delta_2 OnlyWave2_h + \delta_3 wave_t + \zeta_{h,t} \tag{14}$$

$$Consumption_{h,t} = \lambda_0 + \lambda_1 OnlyWave1_h \times Wave_t + \lambda_2 OnlyWave1_h + \lambda_3 wave_t + \theta_{h,t} \tag{15}$$

The coefficients of interest are given by the interaction terms. γ_1 measures the change in household consumption between the two waves for households having a migrant in both waves compared to households that had no migrants in either period. δ_1, measures the change in consumption for households having at least one member of the household migrating over time relative to the control group of households that never had a migrant. λ_1 measures the change in consumption for households having at least one migrant in wave one but none in wave two compared to households that had no migrants in either wave. Estimating the effect of migration using these three separate treatment categories allows me to explore the heterogeneity in consumption patterns for three groups of households having a more or less

[18]Restricting this group to only households with the same migrant away in both waves would drop the sample size to 29 observations.

[19]In our dataset, I have only 240 households who had at least one returning migrant in wave one and had at least one migrant in wave one. The discrepancy in the two sample sizes might be due to two reasons. First, the question used to define the migration status of one member "Did X move away from the household for a period of three months or more in the past 10 years?" was not updated between the two waves. Hence if X was a migrant within the 10 years before wave one, this definition might not apply to wave two. Second, the questionnaire defines returning migrants those who are living in the household. It is possible that the member has returned to live in the village, but not to the household resulting in the treatment variable being equal to zero.

[20]In fact, the question defining each individual as migrant is asked to all household members, including new members of the household. The number of individuals in the baseline data is 6,405, while at follow-up is 7,365.

consolidated history of migration.

δ_1 (and λ_1) capture the effect of having a member of the household migrating (or returning) between two waves on household consumption. While the interpretation of those coefficients is more straightforward and follows previous literature looking at the effect of a change in migration status on household consumption, it is worth discussing more in detail the interpretation of γ_1. The decision to look at the change in consumption for households treated in both waves is motivated by the summary statistics presented in the next section. In general, and as noted in the paper so far, we observe a deterioration of several aspects of the migration experience between the two waves.

First, we see that international migration decreases over time and that the value of remittances also goes down. Second, we observe that the household has less migrants away between the two periods and that they send lower remittances. In absence of an exogenous shock or quasi-experiment to study the effect of this decrease in migration and in the receipt of remittances, I opt to look at the dynamic effect of having at least one member away in both waves on household consumption. Using first-differences in the main analysis (and household fixed effects in the robustness checks) and making the treatment groups less heterogeneous allows me to control for part of the unobserved heterogeneity that might affect both the decision to migrate and consumption. However, in lack of additional data to control for pre-trends in consumption patterns for the treatment and control groups, my estimates of the effect of migration on consumption is still likely to be biased.

In order to attenuate this issue, I follow the literature that uses propensity score matching to match households in the treatment group to comparable control units (Gibson and McKenzie, 2014; de Brauw et al., 2017), in this case households with at least one migrant in the household to households with none. I follow this approach in an attempt to control for any potential remaining bias in the estimates. Using a rich set of observable baseline characteristics, I estimate for each household the probability of being treated. I then use this propensity score to match households in the treatment group to *comparable* households in the control group. If the matching is successful, units in the treatment and control group with similar propensity scores differ only because of the treatment. Thus, by taking the difference in means of the outcome variable I am able to estimate the effect of migration. This approach combined with a diff-in-diff specification allows me to take care both of the selection bias and of time-invariant unobservable characteristics, and has been often used when experimental data is not available (McKenzie et al., 2010).

Notice that, although using first-differences and panel data constitutes an important advantage in estimating the effect of migration, these estimates should be interpreted with care. This is because of unobserved time-varying heterogeneity that might affect both the decision to migrate (or return) and consumption. This been said, restricting the sample to only comparable households in the treatment and control group, by using propensity scores, increases the probability that the treated and control households exhibit similar trends in unobservable characteristics that might bias the results (Gibson and McKenzie, 2014).

Splitting the sample in the three treatment types proves again important when performing this matching. In fact, if I were to match households with at least one migrant to households without any migrants at baseline and take the difference in their average consumption I might be measuring both a change in consumption and a change in treatment status between baseline and follow-up. This is due to the fact that households might switch their status between waves. The possible switches can be illustrated in the following way:

$$Treated_{h,1} \begin{cases} Treated_{h,2} & (AlwaysTreated_h) \\ Control_{h,2} & (OnlyWave1_h) \end{cases} \tag{16}$$

$$Control_{h,1} \begin{cases} Treated_{h,2} & (OnlyWave2_h) \\ Control_{h,2} & (NeverTreated_h) \end{cases} \tag{17}$$

The decision to return or to migrate between the two waves might be endogenous to household consumption, which would bias the estimation of the difference-in-difference model. Moreover, pooling in the treatment group at baseline both households that are treated in both waves and returning migrants households (and similarly in the control group households with members migrating between waves and households that never had a migrant) would make our treatment and control group less comparable.

As the results show, there is a large heterogeneity between households with a consolidated history of migration and households that are treated in only one of the periods. This might result in the violation of the key assumption for propensity score matching, i.e. the conditional independence assumption, stating that any systematic difference between treatment and control with the same baseline observable characteristics is attributable to the treatment.

Thus, by focusing on each of the three treatment groups separately I make sure that the matching is done on unit as comparable as possible within both the treatment and the control group. To re-estimate equation 13, I match household that have at least one migrant in both waves ($AlwaysTreated_h$) with households that never had a migrant. Separately, I estimate equation 14 matching that had a migrant in wave two but none in wave one ($OnlyWave2$) to households that never had a migrant, and to estimate equation 15 I match households that had a migrant in wave one and none in wave two ($OnlyWave1$) to households that never had a migrant.

I estimate the propensity scores separately to predict each of the three treatments using baseline characteristics. I use the following variables to estimate the propensity scores: demographic characteristics of the household (household size, children below 15 years ration, elderly above 65 ratio, men ratio, illiterate ratio), household infrastructure (number of rooms, square meters, having piped water, having access to electricity), land (land index for land ownership[21], area of homestead land owned, area of agricultural land owned), community level characteristics (population, number of households, number of households with migrants, number of shocks experienced in last years, number of investments carried out in the past years, distance to regional capital, distance to road, distance to bus, distance to bank, cost of reaching capital, region dummies). I estimate the propensity scores using a logit model in all specifications.

Lastly, I estimate equations 13, 14 and 15 keeping only observations with propensity scores that lie on the common support (Crump et al., 2009). Figures 19, 20 and 21 (in Appendix 3) show, respectively, the result of the matching of $AlwaysTreated$, $OnlyWave1$ and $OnlyWave2$ households with households in the control group. In particular, figures 19-21 show the distribution of the propensity scores before the matching and after the matching of treated and control

[21] This is calculated using principal component analysis. For each household the index will have higher value if the household owns more types of land among: homestead land in village, homestead land in urban area, agricultural land, commercial land.

units. The right-hand side figures show that the matching is successful in reducing the difference in observable characteristics between the treated and the control groups. However, this matching is more successful for the third treatment type, $OnlyWave2$. This is plausibly due to the fact that both the treated and the control group do not have any migrants at baseline and might be more comparable in terms of both observable and unobservable characteristics [22].

I include village level fixed effects in all estimations to control for common shocks that might affect households at the local level. This partly controls for village-specific shocks that might affect consumption for both the treated and control groups. I do not include any additional covariates in the estimation as I only have information on time-varying household characteristics that might be also affected by migration (Gibson and McKenzie, 2014).

Section 3.8 shows that the results of the analysis are robust to two additional specifications. First, I estimate equations 13, 14 and 15 including inverse propensity weights directly in the regressions. Second, I use household fixed effects to account for household-level unobserved time-invariant heterogeneity.

3.6 Summary Statistics

In this section, I document the differences in characteristics between waves, between the treatment and control group and between the three treatment groups defined in the previous section. This allows me to unpack qualitative differences useful in interpreting the results.

Table 15 shows the summary statistics for the whole sample, separately for the two waves of the study. It shows a clear reduction in migration between 2014 and 2018, seeing households with at least one migrant passing from being 67 percent of the sample in wave one to less than 40 percent in wave two. This drop in households with at least one migrant is accompanied by a drop in households with only international migrants dropping to 19 percent in 2018, and a drop in the proportion of households receiving remittances or goods. This suggests a clear decline in the main channel through which migration can improve origin household

[22] This is the reason why most of the studies looking at the effect of migration on consumption (Beegle et al., 2011; de Brauw et al., 2017) focus on households that have at least one member migrating over time and compare them to households that never had a migrant. Although acknowledging this to be the more convincing set of results, I conduct the estimation for the other two categories in an attempt to show the heterogeneity of the effects for different household types.

consumption (Yang, 2008).

Table 15 also shows that household size increases between the two waves, and that this increase is not driven by an increase in the share of household members aged less than fifteen or more than sixty-five. I also see an overall improvement in household infrastructures in terms of access to electricity, to piped water in the household or village and of the average size of dwellings. As expected, not much change has been experienced in terms of hectares of land owned by the household. This follows plausibly from the strict regulations on land allocation in the Ethiopian context that limit purchase, mortgage or exchange of the land the household has the right to use (De Brauw and Mueller, 2012).

Table 15: Summary statistics of migration experience, household and community characteristics by survey year 2014 and 2018.

	(1) 2014	(2) 2018	(3) Diff
Migration			
Has at least one migrant	0.67	0.39	-0.271 ***
-internal migrant	0.57	0.72	0.155 ***
-international migrant	0.29	0.19	-0.100 ***
-both internal and international	0.14	0.09	0.055***
Nr migrants in HH	1.91	1.61	-0.294 ***
Receives remittances	0.60	0.39	-0.217 ***
Receives Goods	0.28	0.15	-0.133 ***
Value remittances (monthly)	387.99	573.86	185.88 **
Value Goods (monthly)	173.40	139.55	-33.84
Household characteristics			
Household size	5.31	6.12	0.816 ***
% members<15	0.37	0.29	-0.084 ***
% members>65	0.08	0.09	0.007
% Employed	0.22	0.21	-0.013 *
% Illiterate	0.41	0.28	-0.125 ***
% Men	0.49	0.47	-0.019 **
Has electricity	0.25	0.40	0.152 ***
Nr Rooms	1.75	2.27	0.522 ***
Piped water	0.58	0.82	0.239 ***
Toilet in HH	0.87	0.86	-0.012
Hct homestead land	0.06	0.11	0.050 ***
Hct Agricultural land	0.86	0.81	-0.042
Community Characteristics			
Population	5,351.19	5,349.65	-1.531
Distance to capital	42.59	42.55	-0.041
Distance to road	4.30	4.29	-0.015
Nr shocks last 20 yrs.	3.37	3.38	0.002
Tigray	0.25	0.25	-0.002
Amhara	0.25	0.25	0.000
Oromyia	0.25	0.25	0.002
SNNP	0.25	0.25	-0.000
Outcome variables			
Tot consumption	2002	1991	-11.322
Food consumption	1366	1340	-25.612
Non-food consumption	637	649	12.606
Tot consumption p.c.	409	351	-58.586 ***
Food consumption p.c.	281	241	-39.975 ***
Non-food consumption p.c.	128	110	-18.919 ***

Note: Summary statistics for the pooled sample by wave. Column 1 shows the average for each variable in 2014, column 3 the average in 2018 and column 3 the difference in means between the two.

I now look more in detail at the differences over time between the treated and control groups. As outlined in the previous section, this paper presents evidence on the effects of migration for three different groups. Table 16 shows the change between waves in sample characteristics for each of the groups used in the analysis.

Table 16: Summary statistics of household characteristics and household consumption for the three treatment types and for the control group by year of survey.

	AlwaysTreated		OnlyWave1		OnlyWave2		Control	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	2014	2018	2014	2018	2014	2018	2014	2018
Household characteristics								
Household size	5.32	6.98	5.37	5.57	5.25	6.76	5.20	5.68
% members<15	0.35	0.23	0.36	0.32	0.39	0.22	0.42	0.35
% members>65	0.08	0.07	0.08	0.10	0.07	0.08	0.08	0.09
% Employed	0.23	0.17	0.20	0.23	0.24	0.18	0.24	0.23
% No School	0.38	0.24	0.40	0.30	0.45	0.26	0.43	0.31
% Men	0.50	0.47	0.47	0.47	0.49	0.46	0.49	0.47
Has electricity	0.28	0.39	0.25	0.41	0.22	0.40	0.21	0.38
Nr Rooms	1.77	2.46	1.82	2.28	1.71	2.16	1.59	2.05
Piped water	0.58	0.78	0.62	0.84	0.51	0.79	0.56	0.84
Toilet in HH	0.89	0.87	0.89	0.87	0.91	0.88	0.81	0.82
Hct homestead land	0.06	0.12	0.07	0.11	0.06	0.11	0.06	0.12
Hct Agricultural land	0.93	0.84	0.94	0.81	0.69	0.88	0.69	0.73
Outcome variables			'					
Tot consumption	2093	2227	2005	1924	2052	1991	1860	1810
Food consumption	1429	1449	1353	1313	1406	1333	1288	1254
Non-food consumption	664	777	652	609	646	658	572	553
Tot consumption p.c.	421	336	404	378	432	311	394	342
Food consumption p.c.	291	220	274	262	294	212	274	244
Non-food consumption p.c.	129	116	130	115	139	98	120	98

Note: Summary statistics are for each treatment type and for the control group by wave. Columns 1-2 contain the summary statistics for the analysis on households that had at least one migrant in both waves. Columns 3-4 contain the summary statistics for the sample used in the analysis on households that had at least one migrant at baseline, but not at follow-up. Columns 5-6 contain the summary statistics for the sample used in the analysis on households that had no migrants at baseline and at least one at follow-up. Columns 7-8 present the summary statistics for the control group, i.e.households with no migrants in either wave.Information is provided for each sample by wave of the survey.

Table 16 is structured as follows. Columns 1-2 present the summary statistics for households with at least one migrant in both waves ($AlwaysTreated$). Columns 3-4 present the summary statistics for households with at least one migrant in wave one, but none in wave two ($OnlyWave1$). Columns 5-6 present the summary statistics for households with at least one migrant in wave two, but none in wave one ($OnlyWave2$). For each group, I present the average value of household characteristics in wave one (2014) and in wave two (2018). The summary statistics presented in columns 7-8 contain information on the control group used throughout the analysis, i.e. households that did not have any migrants in either wave of the survey. Note that the control group is the same throughout the analysis, but the regressions are run separately for each of the three treatment types.

First, I present the average value of household characteristics in the top part of table 16. These summary statistics motivate the need to account for selection into treatment in our empirical analysis. Thus, I use this set of characteristics to estimate the propensity scores in the empirical analysis and ensure the comparability of the treatment and control groups. Moreover, some of these figures are useful to interpret the results. I notice that household size increases between the two waves, but this increase is larger for households with at least one migrant in both waves (column 2) and households with at least one migrant only in wave two (column 6).

Then, I present the summary statistics of the outcome variables used in the analysis in the bottom part of table 16. All the consumption values are monthly and expressed in 2014 Birr. Columns 1-2 show that while household consumption increased (and particularly, non-food consumption) over time for households with at least one migrant in both waves, per capita consumption decreased over time. Although this dip in per capita consumption is also registered for households in the control group, the decrease is larger for households in the $AlwaysTreated$ category. Several channels might explain this decrease in consumption: from a drop in the value of remittances received, and in the number of migrants away, to a decrease in the duration of migration (see table 17).

Columns 3-4 of table 16 show that all consumption measures decrease over time for households with at least one migrant only in wave one. This dip in both household and per capita consumption is smaller compared to the decrease in per capita figures reported in columns 1-2. This might suggest two things. First, no longer having one migrant away might decrease household consumption, plausibly because of a drop in remittances. Second, the return of productive members in the household might offset the negative effect due to the loss of remittances income.

Lastly, columns 5-6 of table 16 show the change in average consumption for households with at least one migrant only in wave two. These figures show a decrease in all the measures of consumption, with the exception of non-food consumption. This might suggest two different explanations. One, the migration of one of the members might have been too recent (the median duration of migration is of 24 months for this group) for household consumption to increase. Two, that the household might need to hire additional labor to compensate for the loss in workforce due to the migration of younger household members.

It should be noted that consumption is higher on average for the three treatment types compared to the control group. This is the case throughout the three treatment types suggesting a plausible selection into migration of richer households. This motivates further the need to evaluate the effect of migration with empirical methods aimed at reducing selection bias.

The summary statistics presented so far aimed at showing the difference in characteristics between the treated and control group over time. However, additional information is needed on the three treatment groups used in the analysis. In fact, observable and unobservable differences between households treated in both waves or only in one of the periods largely motivate the empirical specification used in the analysis. I investigate these differences further exploiting information available from the migrant-only questionnaire. While these variables cannot be used in the empirical analysis, qualitative evidence on different migration experiences might help interpreting the results.

Table 17 shows information on the migration experience, on the demographic characteristics of the migrants, on the value and use of remittances and on the destination of migrants in the household. Columns 1-2 show the characteristics for households that have at least one migrant in both waves, by wave of the survey. Column 3 shows the summary statistics for households treated only in wave one, column 4 for households treated only in wave two. Notice that by definition information for the two latter groups is only available for the wave of the survey when the household has at least one migrant away.

Table 17: Summary statistics of characteristics of migration experience for the three treatment types.

	(1) Always migrant	(2)	(3) Only wave 1	(4) Only wave 2
	Wave 1	Wave 2		
Destination				
Internal migrants	0.56	0.74	0.57	0.67
International migrants	0.30	0.18	0.28	0.21
Both	0.14	0.08	0.15	0.12
Household characteristics				
Nr Migrants	1.85	1.63	1.93	1.56
Avg Age	25.40	23.76	25.17	24.46
Avg Months Away	35.05	25.83	35.52	29.49
% Men	0.77	0.79	0.77	0.80
% Head's child	0.98	0.97	0.98	0.96
Nr Migrants remit	1.38	1.25	1.37	1.06
Receives remittances	0.62	0.37	0.59	0.41
Receives Goods	0.30	0.14	0.26	0.16
Value Remittances (monthly)	447.17	521.86	346.11	677.87
Value Goods (monthly)	236.31	152.34	124.80	114.51
Main Use Remittances				
Everyday Consumption	0.56	0.64	0.53	0.43
Land and Agriculture	0.09	0.12	0.16	0.24
Home improvement	0.14	0.09	0.11	0.09
Savings	0.04	0.03	0.04	0.03
Health and Educ	0.07	0.06	0.04	0.12
N	327	327	473	147

*Note:*Average values of sample characteristics for the three treatment types. Columns 1 and 2 present information for households with at least one migrant in both waves. Column 3 for households with at least one migrant only in 2014 and column 4 only in 2018.

Table 17 confirms a decline in migration and, particularly, in the probability of having international migrants across the three treatment groups. This is relevant to interpret the results as international migrants tend to send more remittances and goods back home relative to internal migrants. Thus, if the value of remittances received by the household decreases over time, the reduction in international migration might have a negative effect on consumption. Evidence of the difference in remittances value by destination of the migrant can be found in table 18. I report monthly average remittances and value of goods [23] received by the

[23]These measures are constructed as follows. Each household indicating to have at least one migrant sending

household. The value of remittances received by households with only international migrants is between five and ten times the amount received by household with only internal migrants. This is also true for the value of goods received[24]. Similarly, households with both international and internal migrants receive higher remittances and goods than households with only internal migrants. Thus, the drop in international migration might negatively affect household consumption for households treated in both waves and for those treated only in wave two. The results of the next section confirm this intuition.

Table 17 also shows a decrease in the number of migrants per household between the two waves, both for households in the $AlwaysMigrant$ category, and for households with at least a migrant only in wave two compared to households with at least one migrant in wave one. When looking at remittances, I notice a decrease in the number of migrants that send remittances to the household and in the probability of receiving both goods and remittances between the two waves. The drop in the probability of receiving goods and remittances, together with the decrease in the number of migrants over time and with the drop in international migration might help interpret the results presented in the next section and, in general, a negative effect of migration on consumption. Although this discussion relies on qualitative evidence, it helps investigating potential mechanisms not captured by our simplified binary measure of migration.

Lastly, table 17 provides some information on the use of remittances for the three groups. The larger share of remittances is used for everyday consumption by all the treated groups, with households in the $OnlyWave2$ spending relatively less for this category. Spending in everyday consumption is followed by use of remittances to improve agricultural and land productivity[25], or to improve one's home[26]. I notice that households with migrants only in wave two differ both in terms of the share of remittances devoted to improving land and

remittances or goods is asked to report their value. I sum the value sent home by all migrants currently away and adjust it for inflation using 2014 as base year. The original value is reported at the annual level, but I report monthly figures for comparison with the measures of consumption used in the analysis.

[24]Notice that for households with at least one migrant in both waves, the data shows a substitution effect between remittances and goods received. While remittances value goes up between waves for households with only international migrants, the value of goods goes down. This might suggest that as households have a more consolidate history of migration, they start receiving more income and less in kind transfers.

[25]This category includes the purchase of land, the purchase of agricultural equipment, seeds, irrigation systems, water, paying wages to agricultural employees.

[26]This category includes household goods such as furniture and household utensils, electronic goods and construction and development of homestead.

agriculture and the one invested in human capital[27]. The share of remittances invested in the latter is also higher for those who always had a migrant relative to those with migrants only in wave one. This information on the use of remittances suggests that poor households do not only use remittances to increase their consumption, but might invest it in human capital and agricultural productivity (Mendola, 2012).

Table 18: Summary statistics of remittances value for the three treatment types and by destination of migrants

	(1)	(2)	(3)	(4)
	Always migrant		Only wave 1	Only wave 2
	Wave 1	Wave 2		
Value Remittances (monthly)				
Only Internal Migrants	109.06	146.93	91.57	220.13
Only International Migrants	792.54	1173.58	504.30	1125.23
Both Internal and International	580.96	567.49	606.98	874.79
Value Goods (monthly)				
Only Internal Migrants	59.81	58.68	83.20	53.29
Only International Migrants	1025.24	622.66	388.78	225.39
Both Internal and International	222.86	123.55	143.79	84.11
N	327	327	473	147

*Note:*Average values of monthly remittances by destination of migrants in the households. Columns 1 and 2 present information for households with at least one migrant in both waves. Column 3 for households with at least one migrant only in 2014 and column 4 only in 2018. All values are monthly and expressed in 2014 Ethiopian Birr.

Aside from providing information on the migration experience of these three groups, the information in table 17 is meant to motivate further the empirical specification used in the analysis. In particular, by showing that the three treatment types differ in their observable characteristics, I want to motivate the choice of splitting the sample in these three categories when performing the analysis. In fact, looking at the summary statistics presented so far we can infer that households treated in both waves and those treated only in one will differ also in terms of unobservable characteristics. The results presented in the next section confirm this assumption by showing heterogeneous effects of migration on the consumption of these three treatment groups.

[27]This category includes remittances spend in health and medical expenditure and in education.

3.7 Results

The results of the main analysis are presented in table 19 and 20. The results presented in this section show that households with at least one migrant in both waves and those treated only in wave two experience a decrease in per capita consumption and a change in household composition. This is not the case for households treated only in wave one. Moreover, household overall consumption increases for households with at least one migrant in both waves, plausibly in line with the increase in household size experienced by this group.

Table 19: The effect of migration on per capita monthly household consumption by treatment type: unmatched and matched sample.

				PSM		
	Total	Food	Non-Food	Total	Food	Non-Food
	(1)	(2)	(3)	(4)	(5)	(6)
A: Always migrant						
Treated x Wave	-0.092 *	-0.123 **	-0.079	-0.066	-0.121 *	-0.007
	(0.055)	(0.061)	(0.080)	(0.060)	(0.067)	(0.089)
N	582	582	582	500	500	500
B: Only wave 2						
Treated x Wave	-0.187 **	-0.225 ***	-0.120	-0.184 **	-0.238 ***	-0.070
	(0.075)	(0.079)	(0.107)	(0.082)	(0.089)	(0.117)
N	402	402	402	318	318	318
C: Only wave 1						
Treated x Wave	0.053	0.046	-0.016	0.047	0.020	0.018
	(0.052)	(0.058)	(0.076)	(0.057)	(0.065)	(0.084)
N	728	728	728	609	609	609

Note: All coefficients are estimated using a different regression. Columns 1,2 and 3 are estimated on the unmatched sample. Columns 4,5 and 6 are estimated using propensity score matching and observations on the common support. Always migrant is a binary variable equal to one if the household has at least one migrant in both waves, only wave 2 is a binary variable if the household has at least one migrant only in wave 2, and only wave 1 is a binary variable equal to one if the household has at least one migrant only in wave 1. The control group are households with no migrants in either wave. Outcome variables are real monthly per capita consumption expressed in natural logs.All regressions include village level fixed effects. All standard errors are clustered at the household level.

Firstly, I look at table 19 showing the results on the effect of migration on per capita household consumption, conventionally considered as a better measure of welfare in contexts of poverty (De Brauw and Harigaya, 2007). Columns 1 to 3 present the results when using the unmatched sample. Columns 4 to 6 show the results when using propensity score matching. The table shows the estimation of equation 13-15 and presents the coefficients separately for each of the three treatment types. I report the coefficients for the interaction terms of interest γ_1, δ_1, λ_1 , full tables of the main results are presented in appendix 3 (see table 34). Two main findings emerge from table 19. First, the results obtained implementing propensity scores matching differ from those obtained running the regressions on the unmatched sample. Although the differences are not large, the magnitude of the coefficients is smaller and the statistical significance of the coefficients changes in the two specifications for households with at least one migrant in both waves. This confirms a bias in the estimates due to selection into migration for households in this treatment group. Second, the effect of migration on consumption varies depending on the treatment type.

The results are heterogeneous across the three treatment types. Panel A of table 19 shows an overall negative effect on consumption for households with at least one migrant in both waves of the survey. However, none of these effects is statistically different from zero, with the exception of a weak negative effect of around 12 percent on food-consumption. I interpret these coefficients as showing that there is no statistically significant change in overall household consumption for households with at least one migrant in both waves relative to households who had no migrants in either period. The negative sign of the coefficients, however, indicates a plausible worsening in household welfare. This is in line with a deterioration of the income channel through which households might benefit in terms of consumption, i.e. lower probability of receiving remittances, lower number of migrants away, and lower value of remittances received, as discussed qualitatively in the previous section.

Panel B of table 19 shows that migration has overall negative effects on per capita consumption for households who have at least one member migrating between the two waves: total consumption drops by 18.4 percent, and this is driven by a larger dip in food-consumption of around 24 percent. Notice that the magnitude of these effects is in line with that presented in Gibson et al. (2011). This might suggest that in the short-run the loss of agricultural labor is not offset by the receipt of remittances. It should also be kept in mind that internal migration increased between waves and that the amount of remittances sent back home by internal migrants is much lower than in the case of international migration. Moreover, migration became riskier over the four years of analysis and this might translate in a higher difficulty for the migrant to send money and goods back home.

Lastly, panel C of table 19 shows that the effect of migration for households treated only in wave one is not statistically significant and smaller in magnitude. The interpretation of these coefficient is less straightforward as the households in this category are less homogeneous. In fact, it might be that migrants returned to the household of origin bringing back their savings and positively affecting consumption, but increasing household size might have the opposite effect. Moreover, as I am unable to measure whether the migrant returned to the household or just to the village, these coefficients might indicate that little has changed for the household.

To conclude, these findings suggest that migration reduces per capita consumption in the short-run in line with previous findings for Tongan working migrants in (Gibson et al., 2011). Moreover, although the coefficients are not statistically different from zero, the sign of the estimates differs for households in the $OnlyWave1$ group compared to households in the other two categories. These results might be driven by a reduction in the access to migration and in receiving remittances between waves, as documented in the previous section. These negative changes would not affect households that were treated in wave one and are no longer treated in wave two and might explain the positive sign of the coefficients in panel C of table 19.

Secondly, I look at whether these results hold when using as outcome variable overall household consumption measures. Table 20 presents the results. I report the coefficients for the interaction terms of interest γ_1, δ_1, λ_1 , while full tables of these results are presented in appendix 3 (see table 35). Using this second set of outcome variables, I find evidence that migration increases overall household consumption.

Panel A of table 20 shows that this increase is statistically significant only for households that were treated in both waves. In fact, households with at least one migrant in both waves see their total consumption increase by 16 percent relative to households that never had a migrant, and non-food consumption increase by 22 percent. The estimated effect on food consumption is of 10.6 percent, but not statistically different from zero (however, the p-value for this coefficient is of 0.12).

Panel B and Panel C of table 20 show that the effect on total household consumption for those treated in only one of the two waves is not statistically significant. On one hand, panel B shows that this effect is positive but not statistically significant for households having at least one of their member migrating between waves. Although not statistically different from zero, the sign of the coefficients turns positive and oppositely from the results on per capita consumption suggests a positive effect of migration on household consumption. On the other, panel C shows that the sign of the effect for households treated only in wave one is negative for food and non-food consumption and the magnitude of the effect is very small, in line with the coefficients for this category in terms of per capita consumption.

The findings presented so far suggest two things. First, that the effect of migration on consumption differs across the three treatment types. Second, that something might be affecting per capita figures differently than overall household-level consumption for households in the $AlwaysMigrant$ and $OnlyWave2$ categories.

Table 20: The effect of migration on monthly household consumption by treatment type: unmatched and matched sample.

				PSM		
	Total (1)	Food (2)	Non-Food (3)	Total (4)	Food (5)	Non-Food (6)
A: Always migrant						
Treated x Wave	0.112 *	0.081	0.128	0.162 **	0.106	0.226 **
	(0.059)	(0.063)	(0.087)	(0.065)	(0.070)	(0.096)
N	582	582	582	500	500	500
B: Only wave 2						
Treated x Wave	0.006	-0.032	0.073	0.063	0.010	0.179
	(0.079)	(0.081)	(0.114)	(0.091)	(0.095)	(0.129)
N	402	402	402	318	318	318
C: Only wave 1						
Treated x Wave	-0.025	-0.008	-0.099	0.001	-0.002	-0.028
	(0.056)	(0.059)	(0.083)	(0.061)	(0.066)	(0.090)
N	728	728	728	609	609	609

Note: All coefficients are estimated using a different regression. Columns 1,2 and 3 are estimated on the unmatched sample. Columns 4,5 and 6 are estimated using propensity score matching and observations on the common support. Always migrant is a binary variable equal to one if the household has at least one migrant in both waves, only wave 2 is a binary variable if the household has at least one migrant only in wave 2, and only wave 1 is a binary variable equal to one if the household has at least one migrant only in wave 1. The control group are households with no migrants in either wave. Outcome variables are real monthly per capita consumption expressed in natural logs. All regressions include village fixed effects. All standard errors are clustered at the household level.

The discrepancy in the findings for household consumption and per capita consumption leads to explore possible mechanisms. The most straightforward explanation is that migration affects household size. If more productive younger household members migrate (as shown in the sample summary statistics and in line with the literature Mueller et al., 2018), it can be expected that the household needs to increase its sample size to compensate for the loss in labor force. This might be even truer if exchanging and renting out land is particularly problematic or impossible as in this context (De Brauw and Mueller, 2012).

Given the negative and statistically significant effect on per capita consumption limited to households switching to migration in wave two, it is also possible that this effect is negative only in the short-term. Table 21 presents the results of the estimation of equations 13, 14 and 15 using several outcome variables and propensity score matching. I report the coefficients for the interaction terms of interest γ_1, δ_1, λ_1 , while full tables of these results are presented in appendix 3 (see table 36)

Column 1 in table 21 shows that migration increases household size by 1.3 units for households with at least one migrant in both waves and by 1.2 units for those treated only in wave two relative to households in the control group. This suggests that over time households treated in both waves might need to compensate for the worsening in the migration experience documented in section 3.6. As they receive less remittances and have less migrants away, households might need to compensate with extra agricultural labor by either hiring more workforce or by asking help to the extended family. Similarly, households that have at least one of their members migrating between waves, might need to insure themselves against the loss in labor productivity if the income received via remittances is not high enough to compensate for migration.

Moreover, this increase in household size is largely driven by an increase in the number of adults in the household (rather than of dependent members below the age of 15 or above 65). In fact, I find an increase in the number of adults in the household for households treated in both waves, experiencing an increase of around 0.9 units relative to households that never had migrants, and for households treated only in wave two, experiencing an increase in adults of 1 unit. This is in line with the qualitative evidence showing that the average age of migrants is 25 years and with previous findings showing that the majority of migrants in Ethiopia are head of household's children in the age between 25 and 34 (Mueller et al., 2018).

Column 3 of table 21 shows the change in the gender composition of the household. I find that the number of men increases for the first two categories (Panel A and B) in table relative to households that never had a migrant, while there is no effect for households treated only in wave 1 (Panel C). Given the importance of gender-roles in agricultural practices in Ethiopia, it is plausible that this increase in number of men is due to a higher need of productive workforce in the farm. This is in line with the qualitative evidence presented in table 17 showing that 80 percent of the migrants are men in the sample of analysis, suggesting that the household might need to substitute this loss in workforce for the agricultural tasks carried out only by men (e.g. plowing).

Panel C of table 21 shows that households with at least one migrant at baseline and none at follow-up did not experience a change in household composition. This is in line with the findings on household and per capita consumption for $OnlyWave1$ households, showing no change in consumption over time. Several reasons might be driving this result. First, as discussed in section 3.5, households in this category do not necessarily have migrants coming back to the household. Thus, if household members return to the village, but are still out of the household, per capita consumption might not change (especially if they substitute remittances with agricultural work or by exchanging goods with their household of origin). Second, if the migrants do return to the household, they might offset the negative effect on per capita consumption, due to an increase in household size, with an increase in labor productivity. In both of these cases, household per capita consumption would not change.

Lastly, column 4 in table 20 shows the results of the effect of migration on household annual income. This is a different variable from our consumption measures and it is calculated to be the natural logarithm of the sum of any source of income that the household reports to have received in the previous 12 months (this includes the category "Money received by members abroad"). I find that households treated in wave one and no longer in wave two experience a significant drop in income of 48 percent. This is plausibly due to a mechanical decrease in remittances and goods received from members outside the household included in the measure of annual income for households both in the treatment and control group.

Table 21: The effect of migration on household composition and annual total income by treatment type: matched sample.

	(1) HH size	(2) Nr adults	(3) Nr Men	(4) Annual Income
A: Always migrant				
Treated x Wave	1.294 ***	0.873 ***	0.547 ***	-0.200
	(0.245)	(0.184)	(0.179)	(0.142)
N	500	500	500	500
B: Only wave 2				
Treated x Wave	1.210 ***	1.019 ***	0.498 **	-0.004
	(0.321)	(0.229)	(0.226)	(0.177)
N	318	318	318	318
C: Only wave 1				
Treated x Wave	-0.124	-0.211	-0.024	-0.480 ***
	(0.230)	(0.180)	(0.164)	(0.137)
N	609	609	609	609

Note: All coefficients are estimated using a different regression. Columns 1,2 and 3 are estimated using propensity score matching and observations on the common support and a diff-in-diff estimator. Always migrant is a binary variable equal to one if the household has at least one migrant in both waves, only wave 2 is a binary variable if the household has at least one migrant only in wave 2, and only wave 1 is a binary variable equal to one if the household has at least one migrant only in wave 1. The control group are households with no migrants in either wave.HH Size is a variable indicating the number of members in the household. Nr adults is indicating the number of household members aged between 15 and 65. Number of men is the number of men in the household. Annual income is the natural log of real household yearly income. All regressions include village level fixed effects. All standard errors are clustered at the household level.

To conclude, table 21 shows that the change in household consumption is accompanied by a change in the composition of the household. This is true for households treated in both waves and treated only in wave two, but not for households treated only in wave one. The discrepancy in per capita and overall household consumption is driven by an increase in household size and in the number of adults. Moreover, the results show that gender plays a role in this change of household composition.

Overall, I find mixed evidence on the effect of migration on consumption. This effect seems to be negative, at least in the short run. The specific context of Ethiopia, where land regulations impose high costs on households who invest in migration, seems to matter in explaining the mechanisms behind the results. In fact, migrant-sending households experiencing a decrease in per capita consumption, also register an increase in household size over the same period. This result might depend on the need of substituting migrants with productive workforce back home.

3.8 Robustness checks

This section presents the results of the robustness checks conducted to validate the main analysis. First, I include household fixed effects in the estimation to net out household-specific time-invariant characteristics that might affect consumption (Gibson and McKenzie, 2014). I use the following specifications:

$$\Delta Consumption_h = \gamma_h + \gamma_1 AlwaysTreated_h + \zeta_h \tag{18}$$

$$\Delta Consumption_h = \delta_h + \delta_1 OnlyWave2_h + \epsilon_h \tag{19}$$

$$\Delta Consumption_h = \beta_h + \beta_1 OnlyWave1_h + \eta_h \tag{20}$$

Tables 22 and 23 show results very similar to those presented in the previous section in terms of magnitude, sign and statistical significance of the coefficients. First, table 22 shows the main results on per capita household consumption for the three different treatment groups. The results presented in table 22 are estimated using household fixed effects for the unmatched sample (columns 1 to 3) and for the matched sample using the same estimation of the propensity scores as in the main analysis (columns 4 to 6). Then, table 23 presents the results for overall household consumption (columns 1 to 6) and for the mechanisms explored in the analysis, i.e. household size, number of adults, number of men and annual household income (columns 7 to 10).

Table 22: The effect of migration on per capita monthly household consumption: household fixed effects, unmatched and matched sample.

	(1)	(2)	(3)	(4)	(5) PSM	(6)
	Total	Food	Non-Food	Total	Food	Non-Food
A: Always migrant						
Treated	-0.088	-0.125 **	-0.070	-0.062	-0.120 *	-0.003
	(0.055)	(0.060)	(0.080)	(0.060)	(0.066)	(0.088)
N	582	582	582	500	500	500
B: Only wave 2						
Treated	-0.184 **	-0.225 ***	-0.116	-0.181 **	-0.238 ***	-0.064
	(0.074)	(0.078)	(0.106)	(0.082)	(0.089)	(0.117)
N	402	402	402	318	318	318
C: Only wave 1						
Treated	0.057	0.070	-0.010	0.048	0.045	0.020
	(0.052)	(0.057)	(0.076)	(0.057)	(0.062)	(0.084)
N	728	728	728	609	609	609

Note: All coefficients are estimated using a different regression and household fixed effects. Columns 1,2 and 3 are estimated on the unmatched sample. Columns 4,5 and 6 are estimated using propensity score matching and observations on the common support. Always migrant is a binary variable equal to one if the household has at least one migrant in both waves, only wave 2 is a binary variable if the household has at least one migrant only in wave 2, and only wave 1 is a binary variable equal to one if the household has at least one migrant only in wave 1. The control group are households with no migrants in either wave. Outcome variables are real monthly per capita consumption expressed in natural logs. All standard errors are clustered at the household level.

Second, I use inverse probability weights directly in the estimation of equation 13, 14 and 15 as an alternative way to make the treatment and control group more comparable. Differently from propensity score matching, this methodology allows me to use all of the observations in the analysis, without having to pre-screen the sample to keep only observations on the common support.

I estimate the inverse propensity weights to be equal to $1/p$ for the treatment group and $1/(1-p)$ for the control group, where p is the probability of being treated estimated using a logit model. I estimate the probability of being treated as a function of the same set of baseline characteristics used in the main analysis: demographic characteristics of the household (household size, children below 15 years ration, elderly above 65 ratio, men ratio, illiterate ratio), household infrastructure (number of rooms, square meters, having piped water, having access to electricity), land (land index for land ownership, area of homestead land owned, area of agricultural land owned), community level characteristics (population, number of households, number of households with migrants, number of shocks experienced in last years, number of investments carried out in the past years, distance to regional capital, distance to road, distance to bus, distance to bank, cost of reaching capital, region dummies). I estimate the weights separately for each of the three samples used in the main analysis.

The results are presented in table 24. I find a large effect of migration on household consumption for households with at least one migrant in both waves. The estimated coefficients are strongly significant and larger than those estimated using propensity score matching. This is plausibly due both to the larger sample size and to a higher bias in the estimation due to households with weights very close to zero or to one.

Table 23: The effect of migration on monthly household consumption and household composition: household fixed effects, unmatched and matched sample.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
					PSM				PSM	
	Total	Food	Non-Food	Total	Food	Non-Food	HH size	Nr adults	Nr Men	Annual Income
A: Always migrant										
Treated	0.114 *	0.081	0.130	0.163 **	0.106	0.228 **	1.297 ***	0.863 ***	0.540***	-0.183
	(0.059)	(0.062)	(0.087)	(0.064)	(0.069)	(0.095)	(0.243)	(0.183)	(0.178)	(0.143)
N	582	582	582	500	500	500	500	500	500	500
B: Only wave 2										
Treated	0.007	-0.032	0.075	0.065	0.010	0.182	1.210 ***	1.019 ***	0.498**	-0.010
	(0.078)	(0.080)	(0.113)	(0.090)	(0.093)	(0.127)	(0.317)	(0.226)	(0.222)	(0.176)
N	402	402	402	318	318	318	318	318	318	318
C: Only wave 1										
Treated	-0.025	-0.009	-0.097	0.001	-0.002	-0.027	-0.124	-0.211	-0.024	-0.509 ***
	(0.056)	(0.059)	(0.083)	(0.061)	(0.065)	(0.090)	(0.229)	(0.179)	(0.162)	(0.138)
N	728	728	728	609	609	609	609	609	609	609

Note:All coefficients are estimated using a different regression and using household fixed effects. Columns 1,2 and 3 are estimated on the unmatched sample. Columns 4 to 10 are estimated using propensity score matching and observations on the common support. Columns 4 to 6 show the results for overall household consumption, while columns 7 to 10 explore possible mechanisms. Always migrant is a binary variable equal to one if the household has at least one migrant in both waves, only wave 2 is a binary variable if the household has at least one migrant only in wave 2, and only wave 1 is a binary variable equal to one if the household has at least one migrant only in wave 1. The control group are households with no migrants in either wave. Outcome variables are real monthly per capita consumption expressed in natural logs. All standard errors are clustered at the household level.

Table 24: The effect of migration on household consumption, household composition and annual income: inverse probability weighted estimation.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Household level			Per-capita			Mechanisms			
	Total	Food	Non-Food	Total	Food	Non-Food	HH size	Nr adults	Nr men	Annual income
A: Always migrant										
Treated	0.270 ***	0.172 **	0.398 ***	-0.002	-0.099	0.121	1.490 ***	1.378 ***	0.589***	-0.146
	(0.075)	(0.080)	(0.109)	(0.069)	(0.077)	(0.101)	(0.258)	(0.228)	(0.203)	(0.152)
N	500	500	500	500	500	500	500	500	500	500
B: Only wave 2										
Treated	0.107	0.033	0.267 *	-0.170 *	-0.246 **	-0.010	1.310 ***	1.179 ***	0.547**	-0.159
	(0.097)	(0.099)	(0.146)	(0.091)	(0.099)	(0.130)	(0.328)	(0.245)	(0.239)	(0.194)
N	318	318	318	318	318	318	318	318	318	318
C: Only wave 1										
Treated	0.093	0.050	0.133	0.074	0.033	0.113	0.181	0.349	-0.032	-0.245
	(0.073)	(0.077)	(0.103)	(0.065)	(0.072)	(0.092)	(0.241)	(0.231)	(0.186)	(0.153)
N	609	609	609	609	609	609	609	609	609	609

Note: Each coefficient is estimated using a separate regression and using inverse propensity weights in the estimation. Columns 1 to 3 estimate the effect of having at least one migrant in the household on overall household total, food and non-food consumption. Columns 4 to 6 estimate the effects on per-capita consumption. Columns 7 to 10 show the results when investigating possible mechanisms. All the regressions are estimated including inverse propensity weights in the regressions and clustering the SE at the household level. The propensity scores are estimated using a logit model to predict the probability of having at least one migrant in the household. The propensity scores are estimated separately for the three samples in panel A, B and C using a set of baseline characteristics including demographic characteristics, household infrastructure, land ownership and hectares of land owned, community level characteristics.

3.9 Conclusion

Migration has long been studied as a way to improve the welfare of migrants (Harris and Todaro, 1970) and of origin households (Stark and Bloom, 1985).

Its recent inclusion in the 2030 Sustainable Development Goals marks the centrality of this topic for the development agenda of the next decade. This is particularly true in contexts where poverty rates are high and migrant-sending households are more subject to income vulnerability, as it is often the case in rural areas of low income countries, and where migration can help insuring households against negative income shocks (Rosenzweig, 1988; Rosenzweig and Stark, 1989).

Growing empirical evidence has shown that migration can have large welfare effects on origin households income and consumption, as well as on human capital investments and entrepreneurship (Yang, 2008; McKenzie et al., 2010; Gibson and McKenzie, 2014; Clemens and Tiongson, 2017). However, other studies find that in the short-run this effect might be negative if the increase in remittances income does not offset the loss in agricultural labor productivity (Gibson et al., 2011) and a literature on internal migration highlights that seasonal re-location has different effects from international permanent migration (Clemens and Tiongson, 2017).

Given this mixed evidence, additional studies aimed at measuring the effect of migration on origin household welfare are still needed. The scarcity of experimental settings and panel data to study migration (McKenzie et al., 2010), further motivate this study. This paper contributes to this literature by using panel data to conduct an empirical analysis of the effect of migration on household consumption in rural Ethiopia.

The context of Ethiopia is of particular relevance both for its high levels of poverty, highest in its rural areas, and for the large potential gains of migration (de Brauw et al., 2017). In fact, empirical evidence on Ethiopia shows that the migration rate is relatively lower than in other Sub-Saharan countries (Mueller et al., 2018) and that this is partly due to the high costs of migration faced by rural households (Bezu and Holden, 2014; De Brauw and Mueller, 2012; Mueller et al., 2018). Following from these findings, additional evidence is needed to understand the consequences of migration in this particular context.

This paper contributes to the literature by evaluating the effect of migration on consumption for 1200 rural households in the Amhara, Tigray, Oromyia and SNNP regions of Ethiopia. Using household panel data from the Migrating out of Poverty programme, it studies the effect of having at least one migrant in the household on total, food and non-food consumption. In order to reduce issues of selection and unobserved heterogeneity between the treatment and control groups, the analysis uses a difference-in-difference matched estimator to evaluate these effects.

Given the variation in the probability of being treated between the two waves of data, the results are estimated differently for three types of households: those with at least one migrant in both years, those with at least one migrant only in 2014 and those with at least one migrant only in 2018. By splitting the sample in these three groups, the analysis attempts to control for the endogeneity in the decision of switching status between waves as well as providing richer information on different migration experiences.

The results show evidence of a 20 percent decrease in per capita consumption for households that have at least one member migrating between the two waves. Oppositely, total household consumption increases by 16 to 22 percent for households with at least one migrant in both waves compared to households that never had a migrant, while the effect is not statistically significant for the other two groups. To explain these mixed results, the paper shows that changes in consumption are accompanied by changes in the household composition. In particular, it shows that the household size increases over time for those households experiencing a drop in consumption.

Moreover, this increase in household size is not driven by higher fertility, but with an increase in the share of adults and men in the household. This is in line with previous literature on rural Ethiopia showing that as younger and more productive members (mostly, the head's sons) migrate, the household might need additional workforce to compensate for its loss in labor (Mueller et al., 2018). The results hold when using household fixed effects and inverse probability weights directly in the estimation.

Overall, the main conclusion of this paper is that, at least in the short-term, migration can be negative for household consumption and that the magnitude of the effects is in line with evidence from other contexts (Gibson et al., 2011). In line with previous studies on Ethiopia, this negative effect is associated with a change in household composition and the loss in agricultural labor workforce resulting from the migration of young male head of household's children.

To conclude, I would like to draw the attention of the reader to some important limitations of this analysis. The main drawback of the analysis is to not be able to rule out issues of reverse causality and selection into treatment. Thus, the results presented in the analysis cannot be interpreted as causal, notwithstanding the effort made by the author of using panel data and matching techniques to partly address the above mentioned issues.

In absence of an experimental or quasi-experimental setting to study migration, issues of reverse causality might bias the estimates in the analysis. This means that it could be negative changes to household consumption driving the decision to migrate or to return to the household rather than the opposite. This is plausible in either case: household experiencing a negative shock to consumption (e.g. a drought or conflict in neighboring villages) might decide to send a member of the household away as insurance mechanism, but might also see a member of the household return (in the case of rising violence and conflict, having a son away might make the household more vulnerable). Even considering this issue, the results might be helpful for policy makers. In particular, I find that migration is associated with a decrease in per capita consumption for households having a member migrating between waves. This information and the heterogeneity analysis for other groups of households can help identify more vulnerable groups of the population and help target policies aimed at alleviating poverty.

Moreover, the results of this analysis should be interpreted with care as, because of data limitations, I am not able to test for the core underlying assumption of difference-in-differences estimators, i.e. the pre-trend assumption. Although using panel data and propensity score matching techniques reduces the concerns of unobserved heterogeneity and selection, one might plausibly argue that treated households differ from the control group not just because of the treatment. Thus, these estimates are likely to be biased compared to the ideal experimental setting (McKenzie et al., 2010). This been said, the contribution of using panel data when analyzing migration is non trivial (Beegle et al., 2011) and combining this methodology with propensity score matching is a valid alternative in non-experimental settings (Gibson and McKenzie, 2014; de Brauw et al., 2017).

Then, the analysis uses a simplified measure of migration. While the control group is arguably homogeneous enough, the definition of the three treatment categories might be too generic, e.g. households treated in both waves have different members of the household migrating over time and I do not consider the fact that they might be returning home during the two waves, moreover one household might have more than one migrant away and this might affect consumption differently. However, by splitting the sample in the three treatment groups, I reduce partly the heterogeneity among households that experience migration at some point during the period of analysis, while facing the trade-off of sample size reduction.

Lastly, the consumption measures might be biased because of the large portion of food-items that are produced by the household. By attaching to these items a kebele-level price, I might be overestimating their actual value. This might be an issue as it is mostly animal products, usually a better measure of household diet, that fall into this category.

4 Conclusions

This thesis comprises three papers.

First, it investigates the effect of an increase in wealth inequality on political preferences. Building on a long-standing literature in economics, it investigates the effect of wealth inequality on party support and preferences for redistribution. In particular, the first paper looks at the role of wealth inequality in shaping one's preferences for redistribution and on the probability to support the Conservative party in England and Wales. Following from this interest in understanding political preferences, this thesis analyzes their formation. In particular, it evaluates whether specific policies, i.e. conscription, have a short- or long-run effect in shaping civic and political engagement and national identity. The motivation of this second analysis spurs from a renewed interest among policy-makers in promoting policies that increase civic engagement and national and social cohesion, of which conscription is just an example.

This political discourse around national identity has become increasingly linked with the topic of migration. Managing the increase in migratory flows from lower-income countries to Europe (and to the U.S.) has a central role in the agenda of political parties, as it does in the development agenda. While most of the political debate around this topic focuses on the effects of migration at destination, the empirical and theoretical literature in economics has long studied migration as a more complex phenomenon. In particular, migration has important effects on migrant-sending household's welfare, both in terms of increasing income and reducing poverty. Following from the relevance of this topic, the third paper presents an empirical analysis of the effect of migration on origin household consumption in rural Ethiopia.

The first paper analyses the effect of wealth inequality on preferences for redistribution and support for political parties. Motivated by a growing interest on the effects of rising wealth inequality, this paper provides novel evidence on its effect on political preferences. The aim of the paper is to test two mechanisms. First, it looks at whether rising wealth inequality affects the probability of supporting the Conservative party in England and Wales. Second, it looks at whether this change in support is accompanied by a change in preferences for redistribution. With this analysis, I contribute to the literature on the relationship between inequality and preferences for redistribution, as well as providing evidence on the effect of wealth inequality as separate from income and wages (Piketty et al., 2014). Empirical evidence has long been testing the validity of the median voter theory (Meltzer and Richard, 1981) finding that income inequality matters in shaping one's preferences for redistribution (Alesina and Giuliano, 2009). Less is known on whether this is the case for changes in the distribution of wealth (Di Tella et al., 2007; Caprettini et al., 2019). Moreover, and to the best of this author's knowledge, there is no causal empirical evidence testing that wealth inequality changes both preferences and party support. Following the rationale of Barth and Moene (2016), this paper investigates whether as wealth inequality rises richer individuals are more likely to support the Conservative party and less likely to have pro-public sector preferences.

I investigate this relationship exploiting the house price boom in England and Wales between 1995 and 2007. Using local-level house price shocks (as in Disney et al., 2010), I evaluate whether homeowners react differently from renters in terms of political preferences. I find that a 10 percent house price shock makes homeowners 2 percentage points more likely to support the Conservative party. However, the effect on preferences for redistribution is not statistically different from zero, although negative as predicted by the literature. To explain this result I explore two mechanisms. First, I find that the average effect hides important heterogeneity among homeowners. I find that owners of more expensive properties are more likely to support the Conservative party and less likely to support the government's ownership of public services relative to owners of cheaper ones. Second, I find evidence of endowment effects in the housing market, and I find that self-reported house value is associated with a change in preferences in line with the hypotheses tested. The overall conclusion of this paper is that wealth inequality does have an effect in shaping political preferences, but that the average effect on preferences for redistribution is not as predicted in the literature. This might depend, as shown in this analysis, on heterogeneous effects across homeowners or on their perception of own house value.

The second paper investigates the effect of conscription on political ideology, voting participation and national identity. The motivation for this paper is a renewed policy interest in bringing back national service as a way to foster civic duty and social cohesion. Interest resulting in some countries following through with this reform (e.g. France, seehere). The consequences of such reform might be non trivial in terms of costs. The economic literature has long studied the effect that conscription has on labor market outcomes. Among others, Angrist (1990) first found that conscription decreases draftees wages by 15 percent compared to comparable non-draftees, Bedard and Deschênes (2006) find that it increases mortality rates, and Galiani et al. (2011) find that it increases crime rates. Given these non-negligible costs, convincing empirical evidence on the effects of conscription in shaping civic participation and political attitudes is crucial to inform policy-makers. This paper aims at providing such evidence.

Exploiting a quasi-experimental setting and data from Germany and Spain, the paper investigates whether peacetime conscription affects voting participation, political ideology and national identity. The main analysis evaluates the effect of the introduction of conscription in West Germany for men born just before and just after the reform. After ensuring that these two groups are comparable, I find no statistically significant evidence of an effect of conscription on the outcomes considered. Moreover, I find that the reform has no statistically significant effect on other forms of civic engagement (namely, volunteering or participation in local politics), nor on the support for political parties. I perform several robustness checks to test the validity of the results to different specifications. Among these checks, I also test whether the results hold when analyzing the effect of the abolition of conscription in Spain.

The main conclusions of this paper are the following. First, I do not find statistical and economic significant evidence of an effect of conscription in affecting political attitudes. This is key for the policy evaluation of this reform as strong evidence of negative effects of conscription is, instead, available. Second, the effects are comparable in sign and magnitude in the two contexts of analysis plausibly leading one to believe their external validity. Third, the analysis unpacks both short-term and long-term effects and, in line with the impressionable years hypothesis (Giuliano and Spilimbergo, 2014), shows that the effect on political and social values persists over time. Overall, the findings of this analysis do not provide evidence motivating the reintroduction of conscription with the aim recently proposed by policy makers (for media coverage of the political discussion, see here and here).

Lastly, the third paper evaluates the effect of migration on household consumption in rural Ethiopia. Migration is at the center of the political, economic and development agenda (e.g. Sustainable Development Goals 2030, see here). In line with the New Economics of Labor

Migration literature (Stark and Bloom, 1985), growing research is looking at migration as a strategy available to households to improve their livelihood. This is particularly important for poorer households living under financial constraints and income variability (Clemens et al., 2014; Clemens and Ogden, 2014), as in the case of rural households in Sub-Saharan countries. Empirical evidence has shown that migration can have large positive effects on household welfare (Yang, 2008; Gibson and McKenzie, 2014; Clemens and Tiongson, 2017), but that this effect might be negative in the short-run if the absence of the migrant (both in terms of earned wage and role within the household) is not offset by the gains coming from migration, e.g. remittances (Gibson et al., 2011). Given these mixed results, this paper aims at providing additional evidence on this relationship. It does so using panel data from rural Ethiopia, one of the poorest areas in Sub-Saharan Africa and in the world, where identifying strategies that lift more households out of poverty is of particular policy relevance.

This paper uses panel data collected in 2014 and 2018 to investigate whether having at least one migrant in the household affects household consumption. The analysis is carried out using propensity score matching techniques and a difference-in-differences estimator contributing to the literature on the causal effects of migration on household welfare (Gibson and McKenzie, 2014). Exploiting a large variation in the probability of being treated between the two waves and a very low attrition rate, the paper provides several findings. First, I find that the effects varies for the three treatment groups in the analysis: per capita consumption decreases both for households with at least one migrant in both waves and for households treated only at follow-up when compared to households that never had a migrant. This is not the case for households that had at least one migrant at baseline and none at follow-up. Second, I find that overall household consumption increases for the former two groups.

I explore possible mechanisms to explain these results. On the one hand, I find evidence of a change in household composition, explained by an increase in household size and in the number of adults. On the other, I provide qualitative evidence of other channels that might explain the negative effect of migration on per capita consumption, namely, the decrease in international migration (usually characterized by a higher value of remittances), the decrease in the number of migrants per households and in the probability of receiving goods and remittances. The paper concludes the following. First, that it is crucial to conduct the analysis splitting the sample by migration experience to reduce the selection bias. Second, that the high costs of migration in Ethiopia (in terms of barriers to international migration, in terms of conflict in and around the country and in terms of land regulations) might concur to explain the negative effect on per capita consumption. Third, that this effect might be limited to the short-run in line with other literature in the field (Gibson et al., 2011). Overall, I conclude that more research needs to be carried out to fully understand the effect of migration on the origin household. In fact, given the lack of additional data, this analysis addresses only partly issues of selection bias and endogeneity of particular concern when studying migration.

To conclude, I would like to discuss the key findings of this thesis.

First, this thesis shows that wealth, and not only wage or income, inequality affects both the support for political parties and preferences for redistribution. The evidence presented in the first chapter contributes to the current political and economic debate on the effects of rising wealth inequality on political preferences and redistributive policies. Following from the economic evidence on the increasing relevance of wealth accumulation at the top of the distribution (Piketty et al., 2014), this debate hypothesizes that increasing wealth inequality might translate into higher political polarization and lower levels of redistribution (Bonica et al., 2013). The economic literature has long investigated the relationship between inequality and political preferences, but little evidence is focused on wealth. The findings of the paper aim at inform this discussion and policy-makers. The findings presented in the first chapter show that wealth inequality increases the probability of voting for parties less-favorable to redistribution, although on average it does not affect one's beliefs. This is until we consider the heterogeneity in wealth effects for different groups of the population benefiting more or less from rising inequality or the relationship between one's perceptions of own wealth and political beliefs.

Thus, I conclude that political preferences respond to changes in the wealth distribution, that these changes are driven by those who experience larger wealth positive shocks and that one's valuation of own wealth matters in explaining political beliefs. In contexts where wealth is at the center of the political debate, this information might inform political parties of the electoral response to redistributive policies. It also provides novel evidence of the relationship between party support and redistributive beliefs. Understanding what changes this relationship, and in which direction, proves relevant in the current research effort aimed at explaining voting behavior, increasing political and ideological polarization, and the changing policy platform of traditional parties.

Following from this discussion on political preferences, this thesis investigates the formation of political ideology, civic values and national identity. The political debate has recently been focused on several policies aimed at strengthening national cohesion and civic participation, as well as national identity. Several policy proposals have been put forward: from stricter immigration rules to more protectionist trade policies. The analysis carried out in the second chapter aims at informing the policy discussion on the benefits of one of these proposed policies, i.e. re-introducing conscription. Policy-makers in Europe advocate for the positive effects of conscription in increasing civic participation and strengthening national identity. The evidence presented in the second chapter aims at testing whether this is the case.

This evidence is of particular relevance because of the extensive empirical evidence showing the negative effects of conscription on several outcomes, among which crime rates and individual health. Given these costs, policy proposals aimed at bringing back conscription should be grounded on solid evidence of its intended benefits. The findings of the paper suggest that these effects are not present both in the short- and in the long-run, as none of the estimated coefficients is statistically different from zero. Moreover, the results do not seem to be peculiar to a specific context, as they hold when I conduct the analysis using data from different countries, nor to the different measures of political preferences (ideology, political party supported) or civic engagement (voting participation, volunteering, participation in local politics). Thus, the main conclusion of the paper is that policy-makers might want to reconsider introducing conscription as a way to achieve this proposed goal.

Lastly, the analysis carried out in the third chapter analyzes the relationship between migration and migrant-sending household's welfare in a low-income country. Migratory flows and immigration policies are at the center of the political debate in both developed and developing countries, and often accompany political discourses on national identity and social cohesion. Much of this debate focuses on the effects of immigration on receiving countries and communities, although migration is a more complex phenomenon with important consequences on origin countries. The economic literature has long studied migration as a way to improve the migrant and her origin household's living standards. The evidence presented in the third chapter aims at understanding whether this is the case in the particular context of rural Ethiopia.

This context is of interest both for its high poverty rates and for its high barriers to migration. Moreover, its position in the Horn of Africa has put Ethiopia at the center of the migration routes to the Middle East and to Europe. The evidence presented in this paper shows that migration is a complex phenomenon and can have mixed effects on the origin household. In fact, as much of the rural population in Ethiopia depends on agriculture, as land regulations deter individuals from migrating, and as much of the migration flows happen seasonally and withing national borders, the evidence provided in this paper shows an overall negative effect of migration on household consumption, at least in the short-term.

Drawing from previous literature, the main conclusion of this chapter is that this negative effect depends on the high costs of internal and international migration in this context. Evaluating whether this is the case and whether lowering the costs of migration would affect these results is relevant for policy-makers both in low- and high-income countries. If the costs imposed on migrant-sending households are too high, the households might not benefit from migration in terms of welfare and of poverty reduction.

References

Alberto Abadie and Javier Gardeazabal. The economic costs of conflict: A case study of the basque country. *American economic review*, 93(1):113–132, 2003.

Daron Acemoglu, Georgy Egorov, and Konstantin Sonin. A political theory of populism. *The Quarterly Journal of Economics*, 128(2):771–805, 2013.

George A Akerlof and Rachel E Kranton. Identity and the economics of organizations. *Journal of Economic perspectives*, 19(1):9–32, 2005.

James W Albrecht, Per-Anders Edin, Marianne Sundström, and Susan B Vroman. Career interruptions and subsequent earnings: A reexamination using swedish data. *Journal of human Resources*, pages 294–311, 1999.

Alberto Alesina and Nicola Fuchs-Schündeln. Good-bye lenin (or not?): The effect of communism on people's preferences. *The American Economic Review*, 97(4):1507–1528, 2007.

Alberto Alesina, Rafael Di Tella, and Robert MacCulloch. Inequality and happiness: are europeans and americans different? *Journal of Public Economics*, 88(9):2009–2042, 2004.

Alberto F Alesina and Paola Giuliano. Preferences for redistribution. Technical report, National Bureau of Economic Research, 2009.

Facundo Alvaredo, Lucas Chancel, Thomas Piketty, Emmanuel Saez, and Gabriel Zucman. Global inequality dynamics: New findings from wid. world. *American Economic Review*, 107(5):404–09, 2017a.

Facundo Alvaredo, Thomas Piketty, Tony Atkinson, Emmanuel Saez, et al. World wealth and income database. `http://wid.world/data/`, 2017b. Accessed: 2017-06-13.

Joshua Angrist and Alan B Krueger. Why do world war ii veterans earn more than nonveterans? *Journal of labor economics*, 12(1):74–97, 1994.

Joshua D Angrist. Lifetime earnings and the vietnam era draft lottery: evidence from social security administrative records. *The American Economic Review*, pages 313–336, 1990.

Joshua D Angrist, Stacey H Chen, and Jae Song. Long-term consequences of vietnam-era conscription: New estimates using social security data. *American Economic Review*, 101 (3):334–38, 2011.

Vivekinan Ashok, Ilyana Kuziemko, and Ebonya Washington. Support for redistribution in an age of rising inequality: New stylized facts and some tentative explanations. Technical report, National Bureau of Economic Research, 2015.

Anthony B Atkinson, Thomas Piketty, and Emmanuel Saez. Top incomes in the long run of history. *Journal of economic literature*, 49(1):3–71, 2011.

Anthony Barnes Atkinson and Thomas Piketty. *Top incomes over the twentieth century: a contrast between continental european and english-speaking countries*. OUP Oxford, 2007.

Orazio Attanasio, Andrew Leicester, and Matthew Wakefield. Do house prices drive consumption growth? the coincident cycles of house prices and consumption in the uk. *Journal of the European Economic Association*, 9(3):399–435, 2011.

Orazio P Attanasio, Laura Blow, Robert Hamilton, and Andrew Leicester. Booms and busts: Consumption, house prices and expectations. *Economica*, 76(301):20–50, 2009.

David Autor, David Dorn, Gordon Hanson, and Kaveh Majlesi. Importing political polarization? *Massachusetts Institute of Technology Manuscript*, 2016.

David H Autor, Mark G Duggan, and David S Lyle. Battle scars? the puzzling decline in employment and rise in disability receipt among vietnam era veterans. *American Economic Review*, 101(3):339–44, 2011.

Erling Barth and Karl Ove Moene. The equality multiplier: How wage compression and welfare empowerment interact. *Journal of the European Economic Association*, 14(5):1011–1037, 2016.

Travis Baseler. Hidden income and the perceived returns to migration: Experimental evidence from kenya. 2018.

Francesca Bastagli and John Hills. Wealth accumulation in great britain 1995-2005: The role of house prices and the life cycle. 2012.

Thomas K Bauer, Stefan Bender, Alfredo R Paloyo, and Christoph M Schmidt. Evaluating the labor-market effects of compulsory military service. *European Economic Review*, 56(4): 814–829, 2012.

Thomas K Bauer, Stefan Bender, Alfredo R Paloyo, and Christoph M Schmidt. Do guns displace books? the impact of compulsory military service on educational attainment. *Economics Letters*, 124(3):513–515, 2014.

Kelly Bedard and Olivier Deschênes. The long-term impact of military service on health: Evidence from world war ii and korean war veterans. *American Economic Review*, 96(1): 176–194, 2006.

Kathleen Beegle, Joachim De Weerdt, and Stefan Dercon. Migration and economic mobility in tanzania: Evidence from a tracking survey. *Review of Economics and Statistics*, 93(3): 1010–1033, 2011.

Roland Bénabou and Jean Tirole. Incentives and prosocial behavior. *The American economic review*, 96(5):1652–1678, 2006.

Sosina Bezu and Stein Holden. Are rural youth in ethiopia abandoning agriculture? *World Development*, 64:259–272, 2014.

Jagdish Bhagwati and Koichi Hamada. The brain drain, international integration of markets for professionals and unemployment: a theoretical analysis. *Journal of Development Economics*, 1(1):19–42, 1974.

Richard E Bilsborrow, Amarjit Singh Oberai, and Guy Standing. Migration surveys in low income countries: guidelines for survey and questionnaire design. 1984.

Richard E Bilsborrow, Graeme Hugo, Amarjit S Oberai, et al. *International migration statistics: Guidelines for improving data collection systems*. International Labour Organization, 1997.

Niels-Hugo Blunch and Caterina Ruggeri Laderchi. The winner takes it all: Internal migration, education and wages in ethiopia. *Migration Studies*, 3(3):417–437, 2015.

Adam Bonica, Nolan McCarty, Keith T Poole, and Howard Rosenthal. Why hasn't democracy slowed rising inequality? *Journal of Economic Perspectives*, 27(3):103–24, 2013.

George J Borjas. Economic theory and international migration. *International migration review*, 23(3):457–485, 1989.

Gharad Bryan and Melanie Morten. The aggregate productivity effects of internal migration: Evidence from indonesia. *Journal of Political Economy*, 127(5):000–000, 2019.

Gharad Bryan, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak. Underinvestment in a profitable technology: The case of seasonal migration in bangladesh. *Econometrica*, 82(5): 1671–1748, 2014.

Jeffrey V Butler, Paola Giuliano, and Luigi Guiso. The right amount of trust. *Journal of the European Economic Association*, 14(5):1155–1180, 2016.

Sebastian Calonico, Matias D Cattaneo, and Rocio Titiunik. Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326, 2014a.

Sebastian Calonico, Matias D Cattaneo, Rocio Titiunik, et al. Robust data-driven inference in the regression-discontinuity design. *Stata Journal*, 14(4):909–946, 2014b.

Donald T Campbell and Thelma H McCormack. Military experience and attitudes toward authority. *American Journal of Sociology*, 62(5):482–490, 1957.

John Y Campbell and Joao F Cocco. How do house prices affect consumption? evidence from micro data. *Journal of monetary Economics*, 54(3):591–621, 2007.

Bruno Caprettini, Lorenzo Casaburi, and Miriam Venturini. The electoral impact of wealth redistribution. evidence from the italian land reform. 2019.

David Card and Ana Rute Cardoso. Can compulsory military service raise civilian wages? evidence from the peacetime draft in portugal. *American Economic Journal: Applied Economics*, 4(4):57–93, 2012.

David Card and Thomas Lemieux. Going to college to avoid the draft: The unintended legacy of the vietnam war. *American Economic Review*, 91(2):97–102, 2001.

Satya R Chakravarty. *Inequality, Polarization and Conflict: An Analytical Study*, volume 12. Springer, 2015.

Daniel Chandler and Richard Disney. Measuring house prices: a comparison of different indices. Technical report, Institute for fiscal studies, 2014.

Piero Cipollone and Alfonso Rosolia. Social interactions in high school: Lessons from an earthquake. *American Economic Review*, 97(3):948–965, 2007.

Andrew Clark, Emanuela D'Angelo, et al. *Upward social mobility, well-being and political preferences: Evidence from the BHPS*. Università politecnica delle Marche, Dipartimento di economia, 2010.

Michael A Clemens. A labor mobility agenda for development. *Center for Global Development Working Paper*, (201), 2010.

Michael A Clemens and Timothy Ogden. Migration as a strategy for household finance: A research agenda on remittances, payments, and development. *Center for Global Development Working Paper*, (354), 2014.

Michael A Clemens and Erwin R Tiongson. Split decisions: household finance when a policy discontinuity allocates overseas work. *Review of Economics and Statistics*, 99(3):531–543, 2017.

Michael A Clemens, Çağlar Özden, and Hillel Rapoport. Migration and development research is moving far beyond remittances. *World Development*, 64:121–124, 2014.

Irma Clots-Figueras and Paolo Masella. Education, language and identity. *The Economic Journal*, 123(570):F332–F357, 2013.

Richard K Crump, V Joseph Hotz, Guido W Imbens, and Oscar A Mitnik. Dealing with limited overlap in estimation of average treatment effects. *Biometrika*, 96(1):187–199, 2009.

Conchita D'Ambrosio and Edward N Wolff. Is wealth becoming more polarized in the united states? 2001.

Alan De Brauw and Tomoko Harigaya. Seasonal migration and improving living standards in vietnam. *American Journal of Agricultural Economics*, 89(2):430–447, 2007.

Alan De Brauw and Valerie Mueller. Do limitations in land rights transferability influence mobility rates in ethiopia? *Journal of African Economies*, 21(4):548–579, 2012.

Alan de Brauw, Valerie Mueller, and Tassew Woldehanna. Does internal migration improve overall well-being in ethiopia? *Journal of African Economies*, 27(3):347–365, 2017.

Klaus Deininger, Daniel Ayalew Ali, Stein Holden, and Jaap Zevenbergen. Rural land certification in ethiopia: Process, initial impact, and implications for other african countries. *World Development*, 36(10):1786–1812, 2008.

Giorgio Di Pietro. Military conscription and university enrolment: evidence from italy. *Journal of Population Economics*, 26(2):619–644, 2013.

Rafael Di Tella, Sebastian Galiant, and Ernesto Schargrodsky. The formation of beliefs: evidence from the allocation of land titles to squatters. *The Quarterly Journal of Economics*, 122(1):209–241, 2007.

Denise DiPasquale and Edward L Glaeser. Incentives and social capital: Are homeowners better citizens? *Journal of urban Economics*, 45(2):354–384, 1999.

Richard Disney and John Gathergood. House prices, wealth effects and labour supply. *Economica*, 85(339):449–478, 2018.

Richard Disney and Guannan Luo. The right to buy public housing in britain: a welfare analysis. Technical report, IFS Working Papers, 2016.

Richard Disney, John Gathergood, and Andrew Henley. House price shocks, negative equity, and household consumption in the united kingdom. *Journal of the European Economic Association*, 8(6):1179–1207, 2010.

Joan-Maria Esteban and Debraj Ray. On the measurement of polarization. *Econometrica: Journal of the Econometric Society*, pages 819–851, 1994.

Thiemo Fetzer. Did austerity cause brexit? *American Economic Review (forthcoming)*, 2019.

Eleonora Fichera and John Gathergood. Do wealth shocks affect health? new evidence from the housing boom. *Health economics*, 25(S2):57–69, 2016.

Sebastian Galiani, Martín A Rossi, and Ernesto Schargrodsky. Conscription and crime: Evidence from the argentine draft lottery. *American Economic Journal: Applied Economics*, 3 (2):119–36, 2011.

Andrew Gelman and Guido Imbens. Why high-order polynomials should not be used in regression discontinuity designs. *Journal of Business & Economic Statistics*, 37(3):447–456, 2019.

Stephen Gibbons, Stephen Machin, and Olmo Silva. Valuing school quality using boundary discontinuities. *Journal of Urban Economics*, 75:15–28, 2013.

John Gibson and David McKenzie. The economic consequences of âbrain drainâof the best and brightest: Microeconomic evidence from five countries. *The Economic Journal*, 122 (560):339–375, 2012.

John Gibson and David McKenzie. The development impact of a best practice seasonal worker policy. *Review of Economics and Statistics*, 96(2):229–243, 2014.

John Gibson, David McKenzie, and Steven Stillman. The impacts of international migration on remaining household members: omnibus results from a migration lottery program. *Review of Economics and Statistics*, 93(4):1297–1318, 2011.

Vladimir Gimpelson and Daniel Treisman. Misperceiving inequality. Working Paper 21174, National Bureau of Economic Research, May 2015. URL `http://www.nber.org/papers/w21174`.

Paola Giuliano and Antonio Spilimbergo. Growing up in a recession. *The Review of Economic Studies*, 81(2):787–817, 2014.

Clark Gray and Valerie Mueller. Drought and population mobility in rural ethiopia. *World development*, 40(1):134–145, 2012.

Julien Grenet, Robert A Hart, and J Elizabeth Roberts. Above and beyond the call. long-term real earnings effects of british male military conscription in the post-war years. *Labour Economics*, 18(2):194–204, 2011.

John R Harris and Michael P Todaro. Migration, unemployment and development: a two-sector analysis. *The American economic review*, pages 126–142, 1970.

John Hills. *Wealth in the UK: distribution, accumulation, and policy*. Oxford University Press, 2013.

Randi Hjalmarsson and Matthew Lindquist. The causal effect of military conscription on crime and the labor market. 2016.

Andrew Hood and Robert Joyce. Inheritances and inequality across and within generations. Technical report, Jan 2017. URL `https://www.ifs.org.uk/uploads/publications/bns/bn192.pdf`.

Guido Imbens and Wilbert Van Der Klaauw. Evaluating the cost of conscription in the netherlands. *Journal of Business & Economic Statistics*, 13(2):207–215, 1995.

David W Johnston, Michael A Shields, and Peter Siminski. Long-term health effects of vietnam-era military service: A quasi-experiment using australian conscription lotteries. *Journal of health economics*, 45:12–26, 2016.

Eliakim Katz and Oded Stark. Labor migration and risk aversion in less developed countries. *Journal of labor Economics*, 4(1):134–149, 1986.

Jeni Klugman. Human development report 2009. overcoming barriers: Human mobility and development. *Overcoming Barriers: Human Mobility and Development (October 5, 2009). UNDP-HDRO Human Development Reports*, 2009.

Stephen Knack and Philip Keefer. Does social capital have an economic payoff? a cross-country investigation. *The Quarterly journal of economics*, 112(4):1251–1288, 1997.

Katrina Kosec, Hosaena Ghebru, Brian Holtemeyer, Valerie Mueller, and Emily Schmidt. The effect of land access on youth employment and migration decisions: Evidence from rural ethiopia. *American Journal of Agricultural Economics*, 100(3):931–954, 2017.

Rachel E Kranton. Identity economics 2016: Where do social distinctions and norms come from? *American Economic Review*, 106(5):405–09, 2016.

Jon A Krosnick and Duane F Alwin. Aging and susceptibility to attitude change. *Journal of personality and social psychology*, 57(3):416, 1989.

David S Lee and Thomas Lemieux. Regression discontinuity designs in economics. *Journal of economic literature*, 48(2):281–355, 2010.

W Arthur Lewis. Economic development with unlimited supplies of labour. *The manchester school*, 22(2):139–191, 1954.

Robert EB Lucas. Migration and economic development in africa: A review of evidence. *Journal of African Economies*, 15(suppl_2):337–395, 2006.

Katrin Marchand, Julia Reinold, and Raphael Dias e Silva. Study on migration routes in the east and horn of africa. 2017.

Justin McCrary. Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of econometrics*, 142(2):698–714, 2008.

David McKenzie and Hillel Rapoport. Network effects and the dynamics of migration and inequality: theory and evidence from mexico. *Journal of development Economics*, 84(1): 1–24, 2007.

David McKenzie and Dean Yang. Experimental approaches in migration studies. Technical report, 2010.

David McKenzie, Steven Stillman, and John Gibson. How important is selection? experimental vs. non-experimental measures of the income gains from migration. *Journal of the European Economic Association*, 8(4):913–945, 2010.

Costas Meghir, Ahmed Mushfiq Mobarak, Corina D Mommaerts, and Melanie Morten. Migration and informal insurance. Technical report, National Bureau of Economic Research, 2019.

Allan H Meltzer and Scott F Richard. A rational theory of the size of government. *Journal of political Economy*, 89(5):914–927, 1981.

Mariapia Mendola. Rural out-migration and economic development at origin: A review of the evidence. *Journal of International Development*, 24(1):102–122, 2012.

Raven Molloy, Christopher L Smith, and Abigail Wozniak. Internal migration in the united states. *Journal of Economic perspectives*, 25(3):173–96, 2011.

Melanie Morten. Temporary migration and endogenous risk sharing in village india. *Journal of Political Economy*, 127(1):1–46, 2019.

Valerie Mueller, Cheryl Doss, and Agnes Quisumbing. Youth migration and labour constraints in african agrarian households. *The Journal of Development Studies*, 54(5):875–894, 2018.

Tommaso Nannicini, Andrea Stella, Guido Tabellini, and Ugo Troiano. Social capital and political accountability. *American Economic Journal: Economic Policy*, 5(2):222–50, 2013.

Thomas Piketty. Social mobility and redistributive politics. *The Quarterly journal of economics*, 110(3):551–584, 1995.

Thomas Piketty. Brahmin left vs merchant right: Rising inequality and the changing structure of political conflict. *World Inequality Database, WID. World Working Paper Series*, (2018/7), 2018.

Thomas Piketty, Arthur Goldhammer, and LJ Ganser. Capital in the twenty-first century. 2014.

Panu Poutvaara and Andreas Wagener. Conscription: Economic costs and political allure. *The Economics of Peace and Security Journal*, 2(1), 2007.

Robert D Putnam. Bowling alone: Americaâs declining social capital. In *Culture and politics*, pages 223–234. Springer, 2000.

Gustav Ranis and John CH Fei. A theory of economic development. *The American economic review*, pages 533–565, 1961.

Dani Rodrik. Populism and the economics of globalization. *Journal of International Business Policy*, pages 1–22, 2018.

Klaus Roghmann and Wolfgang Sodeur. The impact of military service on authoritarian attitudes: Evidence from west germany. *American Journal of Sociology*, 78(2):418–433, 1972.

Mark R Rosenzweig. Risk, implicit contracts and the family in rural areas of low-income countries. *The Economic Journal*, 98(393):1148–1170, 1988.

Mark R Rosenzweig and Oded Stark. Consumption smoothing, migration, and marriage: Evidence from rural india. *Journal of political Economy*, 97(4):905–926, 1989.

Filipa Sá. Immigration and house prices in the uk. *The Economic Journal*, 125(587):1393–1424, 2015.

Marcin J Sasin and David McKenzie. Migration, remittances, poverty, and human capital: conceptual and empirical challenges. 2007.

Lyle Scruggs, Detlef Jahn, and Kati Kuitto. Comparative welfare entitlements dataset 2. version 2014-03. *University of Connecticut & University of Greifswald*, 2014.

Peter Siminski. Employment effects of army service and veterans' compensation: evidence from the australian vietnam-era conscription lotteries. *Review of Economics and Statistics*, 95(1):87–97, 2013.

Ronald Skeldon. International migration, internal migration, mobility and urbanization: Towards more integrated approaches. 2018.

Oded Stark. The migration of labor. 1991.

Oded Stark and David E Bloom. The new economics of labor migration. *The american Economic review*, 75(2):173–178, 1985.

Oded Stark and David Levhari. On migration and risk in ldcs. *Economic development and cultural change*, 31(1):191–196, 1982.

J Edward Taylor and Philip L Martin. Human capital: Migration and rural population change. *Handbook of agricultural economics*, 1:457–511, 2001.

Asmelash Haile Tsegay and Julie Litchfield. Changing patterns of migration and remittances in ethiopia, 2014-2018. *WP series Migrating our of Poverty*, 2019.

Michael C Wolfson. When inequalities diverge. *The American Economic Review*, 84(2):353–358, 1994.

Dean Yang. International migration, remittances and household investment: Evidence from philippine migrants exchange rate shocks. *The Economic Journal*, 118(528):591–630, 2008.

Dean Yang. Migrant remittances. *Journal of Economic perspectives*, 25(3):129–52, 2011.

Dean Yang and HwaJung Choi. Are remittances insurance? evidence from rainfall shocks in the philippines. *The World Bank Economic Review*, 21(2):219–248, 2007.

Dean Yang and Claudia Martinez. Remittances and poverty in migrantsâ home areas: Evidence from the philippines. *International Migration, Remittances and the Brain Drain Washington DC: World Bank*, pages 81–121, 2006.

5 Appendix

5.1 Appendix 1: supplemental figures and tables to paper 1.

Figure 6: Kernel density of the natural log of self-reported house value for home-owners showing a shift to the right and an increase in the gap between homeowners and renters in England and Wales in the period between 1995 and 2007.

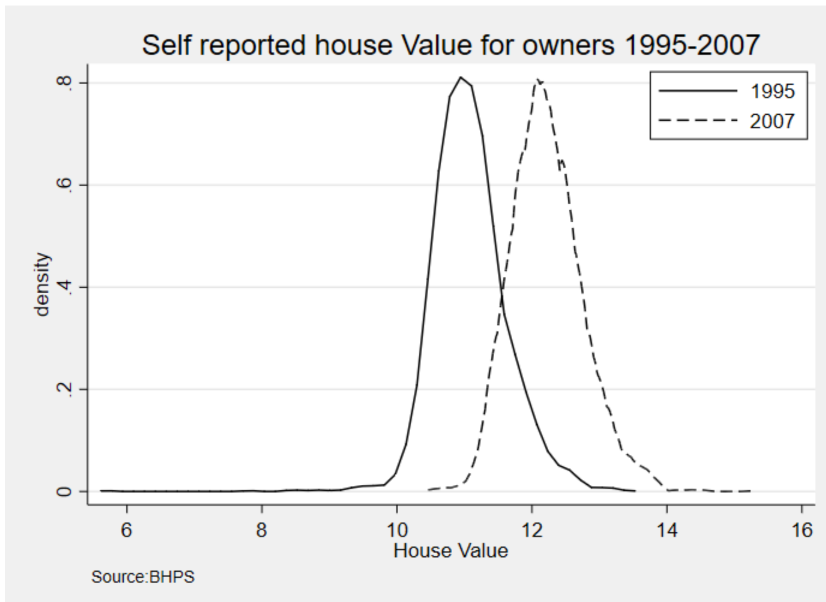

Figure 7: Correlation between average house prices and support for the Conservative Party in England and Wales in the period between 1995 and 2007.

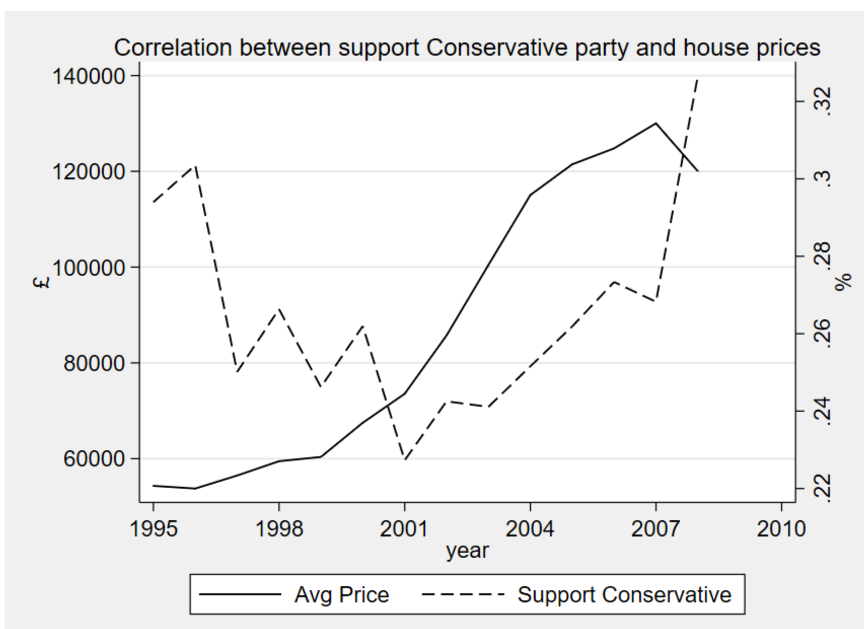

Figure 8: Deviation from linear predicted trend of average house prices in England and Wales for the period between 1995 and 2007.

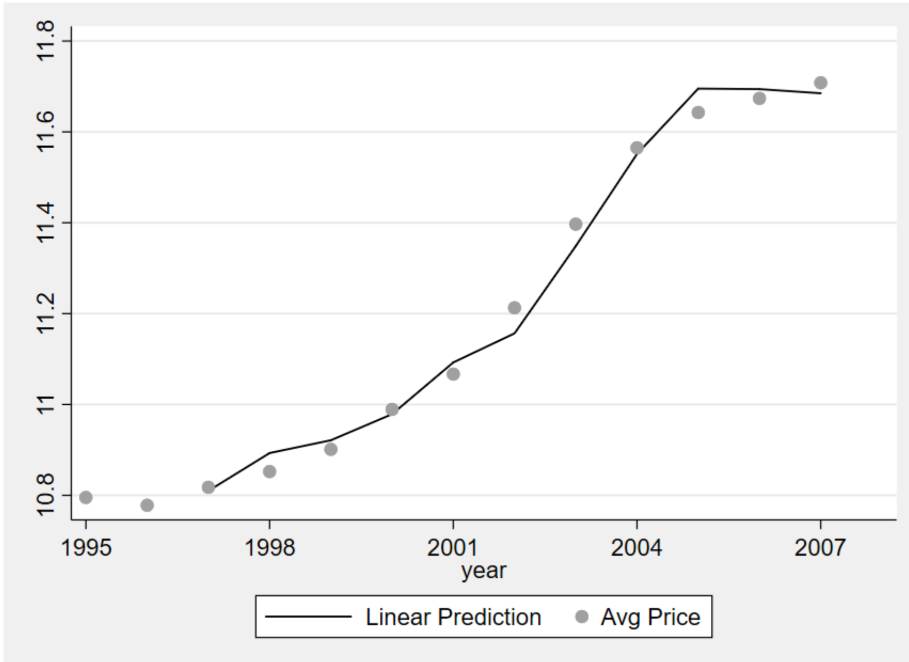

Figure 9: Change in the probability of agreeing with different statements on the role of the government in the economy in England and Wales in the period between 1995-2007.

(a) Pro-Public attitudes

(b) Pro-Private enterprise attitudes

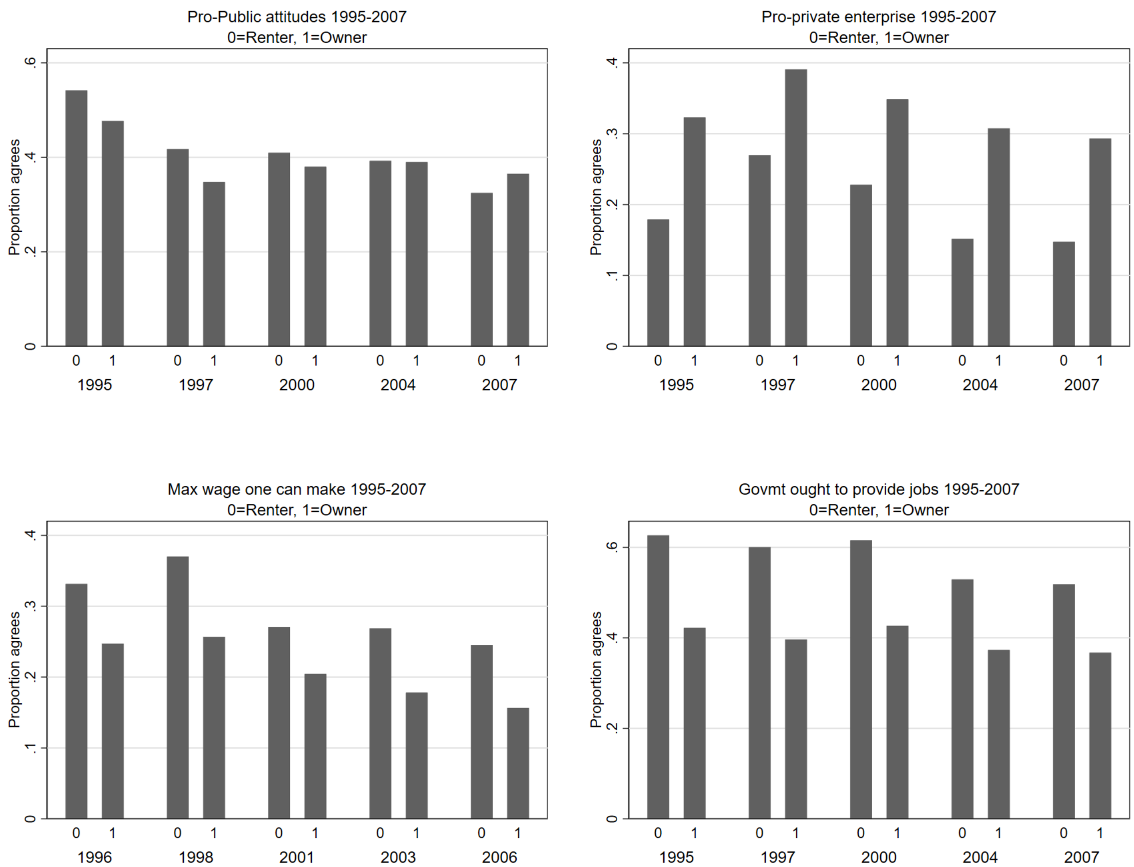

(c) Max wage imposed by government

(d) Government ought to provide jobs

Figure 10: Probability of supporting the Conservative party in England and Wales in the period between 1995 and 2007 : by property type

(a) Detached house (b) Semi-detached

(c) Terraced (d) Flat

Figure 11: Probability of supporting state ownership of public sector in England and Wales between 1995 and 2007: by type of property owned.

Figure 12: Average self-reported house value and average real and nominal house prices in England and Wales between 1995 and 2007.

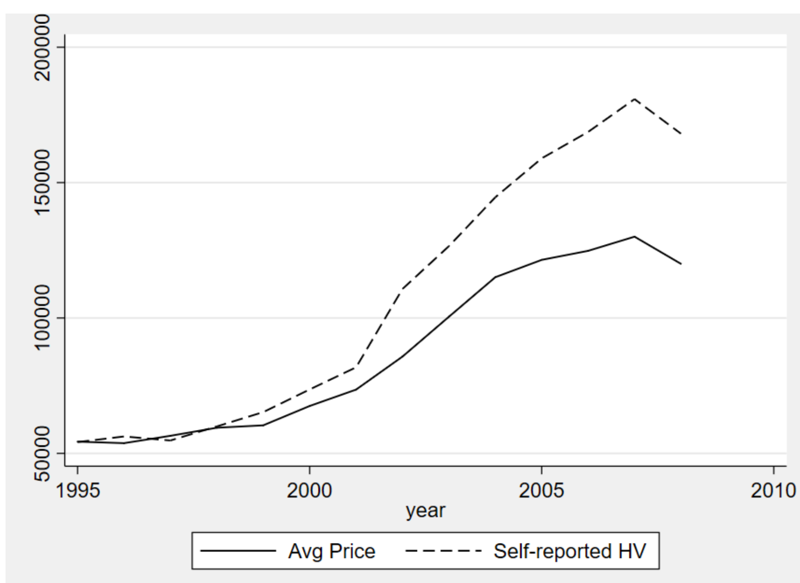

Figure 13: Self-reported attitudes on importance of owning a home in England and Wales in the period between 1995 and 2007: by home ownership status.

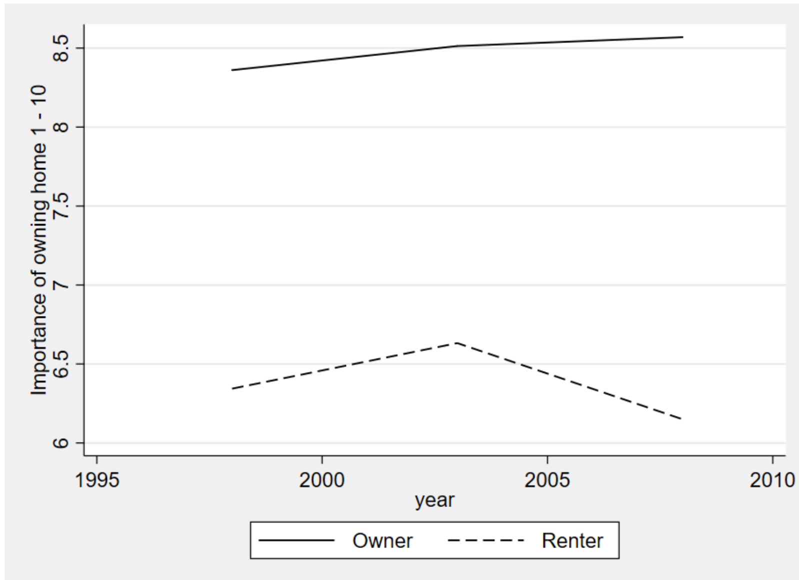

Importance of owning a home is measured as the probability of choosing values between 5 and 10 in when asked “I’m going to read you a list of things that people value. For each one I’d like you to tell me on a scale from 1 to 10 how important each one is to you, where ’1’ equals ’Not important at all’, and ’10’ equals ’Very important’. Owning your own home”. Figure13 shows the statistics for head of households by homeownership status. This variable is available in the BHPS for 3 years : 1998, 2003, 2008.

Table 25: The effect of wealth inequality on the support for the Conservative party and on political attitudes: excluding movers.

	(1) Support Conservative	(2) Closer Conservative	(3) Support Labour	(4) Public
Price shock	-0.095	-0.154	0.147	-0.140
	(0.082)	(0.111)	(0.113)	(0.197)
Price shock x Owner	0.192***	0.236**	-0.103	-0.025
	(0.072)	(0.110)	(0.109)	(0.194)
Owner	0.120***	0.145***	-0.070***	-0.032***
	(0.005)	(0.007)	(0.007)	(0.010)
Couple	0.003	-0.001	0.001	0.006
	(0.006)	(0.008)	(0.007)	(0.012)
Household income	0.046***	0.053***	-0.014***	-0.029***
	(0.005)	(0.006)	(0.004)	(0.007)
Higher Degree	-0.044***	-0.098***	-0.063***	0.012
	(0.012)	(0.017)	(0.015)	(0.026)
Undergraduate	0.003	-0.036***	-0.041***	0.001
	(0.008)	(0.010)	(0.010)	(0.014)
Secondary Edu	0.075***	0.086***	-0.096***	-0.020**
	(0.005)	(0.007)	(0.006)	(0.010)
Male	0.002	-0.001	0.017***	0.075***
	(0.004)	(0.006)	(0.005)	(0.008)
Employed	-0.001	0.000	0.005	0.013
	(0.007)	(0.009)	(0.008)	(0.012)
Household size	-0.001	-0.005*	0.012***	-0.006
	(0.002)	(0.003)	(0.003)	(0.004)
Self Employed	0.093***	0.124***	-0.075***	-0.052***
	(0.010)	(0.012)	(0.012)	(0.017)
Unemployed	0.009	0.008	-0.010	0.019
	(0.011)	(0.017)	(0.018)	(0.029)
Retired	0.030***	0.037***	-0.007	0.006
	(0.011)	(0.011)	(0.012)	(0.018)
Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes
N	40143	28282	40143	15577

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Bootstrapped Standard Errors in parenthesis

Note: All models estimated using OLS. Standard errors are bootstrapped (200 reps).The dependent variables are binary variables equal to 1 if :(1) the individual supports the Conservative Party, (2) the individual feels closer to the Conservative party, (3) the individual supports the Labour party, (4) the individual agrees with state ownership of public services. "Price shock all" is estimated as the residuals of an autoregressive model of order 2 on average house prices at the county level for 53 counties in England and Wales in the period 1995-2007. Owner is a binary variable equal to 1 if the individual is a homeowner and zero if a renter. Sample includes head of households only. Data from the BHPS wave 5-17 (1995-2007) and from the ONS-HPI. The estimation includes a set of controls: binary variables for marital status, gender, education level, retirement status, employment status, log hh equivalised income, county, year of survey and age.

Table 26: The effect of wealth inequality on the support for the Conservative party and on political attitudes: excluding Greater London.

	(1) Support Conservative	(2) Closer Conservative	(3) Support Labour	(4) Public
Price shock	-0.076	-0.068	0.111	-0.315**
	(0.082)	(0.115)	(0.113)	(0.160)
Price shock x Owner	0.158*	0.135	-0.076	0.127
	(0.086)	(0.107)	(0.116)	(0.190)
Owner	0.115***	0.138***	-0.069***	-0.031***
	(0.005)	(0.007)	(0.007)	(0.012)
Couple	0.012**	0.010	-0.007	0.003
	(0.006)	(0.008)	(0.006)	(0.011)
Household income	0.044***	0.050***	-0.011***	-0.029***
	(0.005)	(0.007)	(0.004)	(0.007)
Higher Degree	-0.029**	-0.085***	-0.068***	0.003
	(0.014)	(0.019)	(0.020)	(0.029)
Undergraduate	0.008	-0.031***	-0.051***	-0.007
	(0.008)	(0.011)	(0.010)	(0.017)
Secondary Edu	0.080***	0.091***	-0.100***	-0.021*
	(0.006)	(0.007)	(0.006)	(0.012)
Male	-0.001	-0.005	0.022***	0.077***
	(0.005)	(0.007)	(0.006)	(0.010)
Employed	-0.005	-0.006	0.012	0.016
	(0.007)	(0.008)	(0.008)	(0.013)
Household size	-0.004**	-0.009***	0.014***	-0.005
	(0.002)	(0.003)	(0.003)	(0.004)
Self Employed	0.101***	0.136***	-0.087***	-0.047***
	(0.011)	(0.014)	(0.012)	(0.017)
Unemployed	0.009	0.016	-0.016	0.027
	(0.012)	(0.019)	(0.016)	(0.030)
Retired	0.029***	0.037***	-0.008	-0.002
	(0.009)	(0.013)	(0.011)	(0.023)
Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes
N	37556	26314	37556	14543

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Bootstrapped Standard Errors in parenthesis

Note: All models estimated using OLS. Standard errors in parenthesis are bootstrapped (200 reps).The dependent variables are binary variables equal to 1 if :(1) the individual supports the Conservative Party, (2) the individual feels closer to the Conservative party, (3) the individual supports the Labour party, (4) the individual agrees with state ownership of public services. "Price shock all" is estimated as the residuals of an autoregressive model of order 2 on average house prices at the county level for 52 counties in England and Wales in the period 1995-2007-Excluding the county of Greater London. Owner is a binary variable equal to 1 if the individual is a homeowner and zero if a renter. Sample includes head of households only.The estimation includes a set of controls: binary variables for marital status, gender, education level, retirement status, employment status, log hh equivalised income, county, year of survey and age.

Table 27: The effect of wealth inequality on the support for the Conservative party and on political attitudes: including 2008.

	(1) Support Conservative	(2) Closer Conservative	(3) Support Labour	(4) Public
Price shock	-0.090	-0.158	0.152	-0.158
	(0.074)	(0.098)	(0.094)	(0.153)
Price shock x Owner	0.151**	0.213**	-0.122	0.012
	(0.067)	(0.084)	(0.089)	(0.144)
Owner	0.122***	0.147***	-0.068***	-0.033***
	(0.005)	(0.007)	(0.006)	(0.010)
Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes
Observations	43759	30650	43759	15793

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Bootstrapped Standard Errors in parenthesis

Note: All models estimated using OLS. Standard errors are bootstrapped (200 reps).The dependent variables are binary variables equal to 1 if :(1) the individual supports the Conservative Party, (2) the individual feels closer to the Conservative party, (3) the individual supports the Labour party, (4) the individual agrees with state ownership of public services. "Price shock all" is estimated as the residuals of an autoregressive model of order 2 on average house prices at the county level for 53 counties in England and Wales in the period 1995-2008. Owner is a binary variable equal to 1 if the individual is a homeowner and zero if a renter. Sample includes head of households only. Data from the BHPS wave 5-17 (1995-2007) and from the ONS-HPI. The estimation includes a set of controls: binary variables for marital status, gender, education level, retirement status, employment status, log hh equivalised income, county, year of survey and age.

5.2 Appendix 2: supplemental figures and tables to paper 2.

Figure 14: Plot of the McCrary test of a statistically significant jump in the discontinuity of the running variable function: Germany

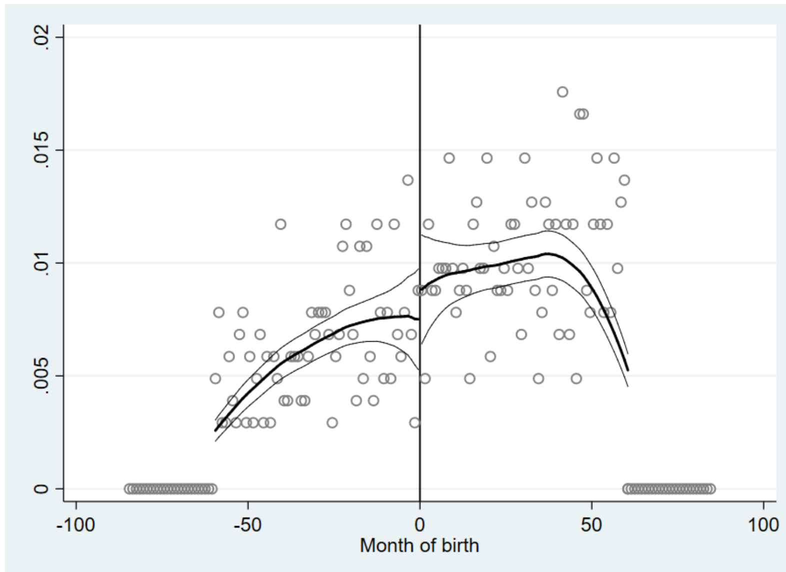

Figure 15: Plot of the discontinuity in the probability of having been conscripted by month of birth: placebo cutoffs +/- 12 months Germany

(a) Cutoff -12 months

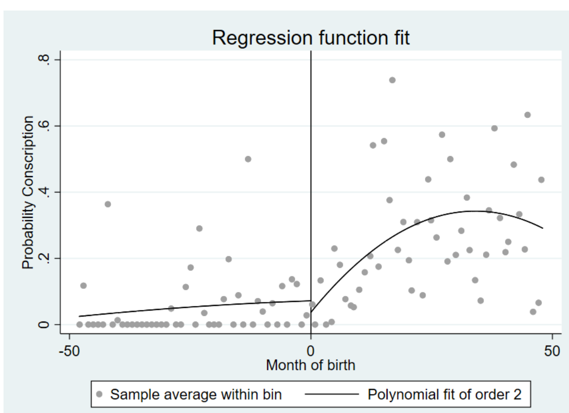

(b) Cutoff +12 months

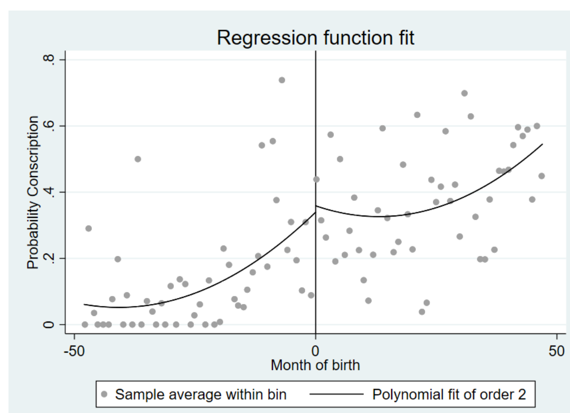

Figure 16: Plot of the McCrary test of a statistically significant jump in the discontinuity of the running variable function: Spain

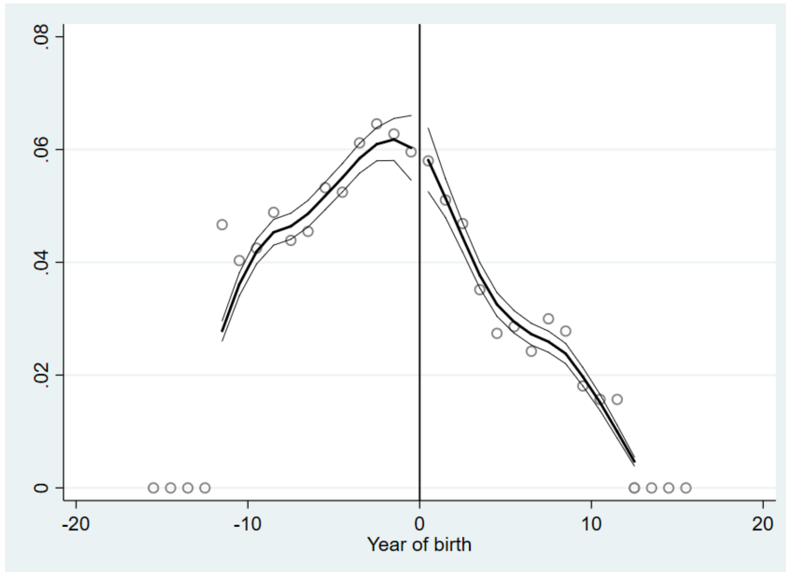

Figure 17: Plot of the discontinuity in the probability of having been conscripted by year of birth: placebo cutoffs +/- 1 year Spain

(a) Cutoff -12 months

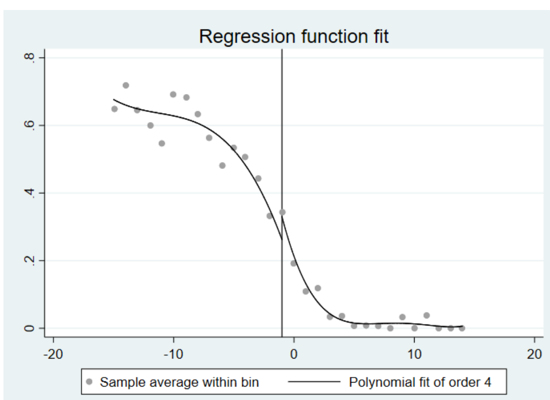

(b) Cutoff +12 months

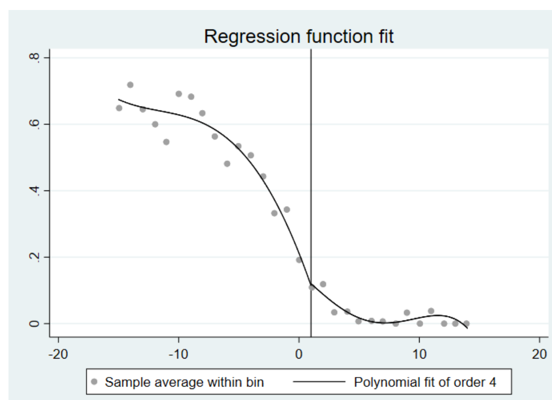

Figure 18: Discontinuity in the outcome variables by year of birth for men born within the 4 years before and the 4 years after the abolition of conscription in Spain.

(a) Rightwing ideology

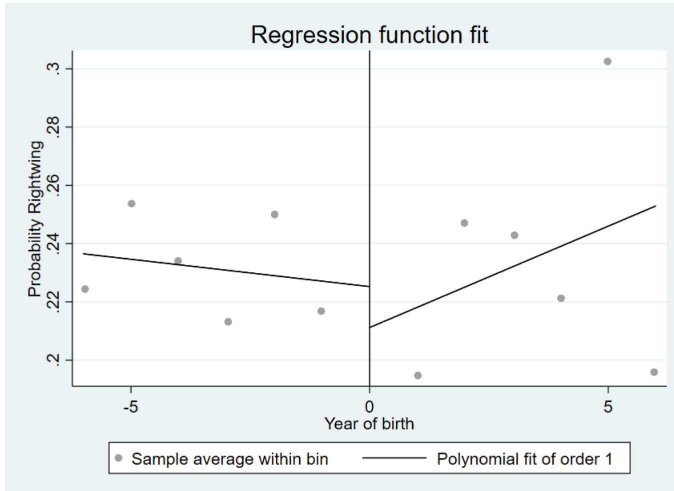

(b) Spanish Identity

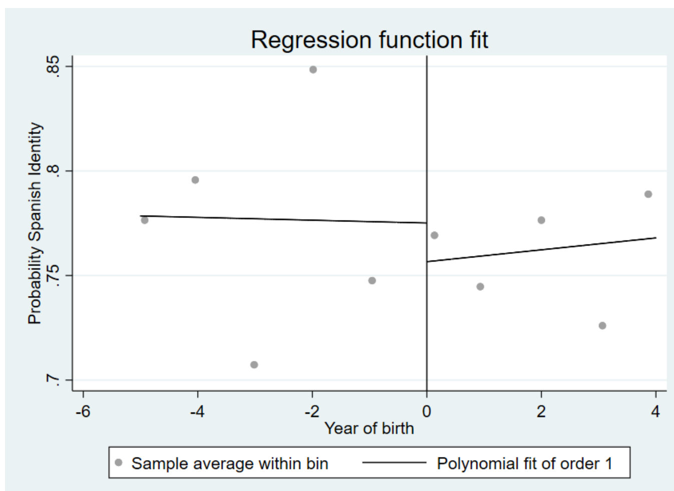

(c) Voted in GE

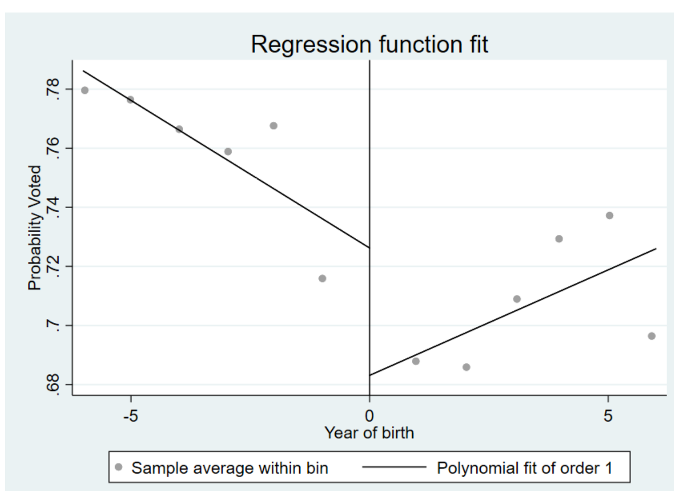

Table 28: Regression analysis of the correlation between conscription voting participation, ideology and civic engagement: analysis military vs. civil service.

	Rightwing	Voted	Rightwing: trim	Local politics	Volunteers	CDU/CSU	SPD
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Conscription	0.155***	-0.023	0.139***	-0.005	0.022	0.142***	0.025
	(0.023)	(0.016)	(0.024)	(0.021)	(0.025)	(0.030)	(0.025)
N	2781	2007	2384	2722	2730	2007	2007
Year + Age	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y
Baseline	Y	Y	Y	Y	Y	Y	Y

Note: Each coefficient is estimated using a separate regression. Rightwing is a binary variable equal to 1 if the individual reports to have a political ideology more rightwing than the median. Voted is a binary variable equal to 1 if the individual voted in the last general election. Conscription is a binary variable equal to 1 if the individual was conscripted. All regressions include controls for: year of the survey, age and region of residence at the time of the survey, marital status, labor force participation, years of schooling, place of residence during childhood and parental education. SE clustered at month of birth

Table 29: The effect of conscription on right-wing ideology, participation in local politics and volunteering: cross-sectional sample, non-parametric and parametric models

	Political attitude 1-10		Rightwing: trim		Local Politics		Volunteers	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	FS	RF	FS	RF	FS	RF	FS	RF
		SS		SS		SS		SS
Panel A: Non-Parametric								
Treated	0.202***	-0.067	0.212 ***	0.001	0.248 ***	-0.034	0.247***	0.096*
	(0.055)	(0.130)	(0.064)	(0.056)	(0.068)	(0.035)	(0.068)	(0.052)
Conscription		0.469		0.223		-0.262		0.408
		(0.846)		(0.293)		(0.260)		(0.365)
Bandwidth		56.5		49.4		35.0		34.6
N		6104		5515		5213		5226
N Eff. obs.		876		706		581		573
Panel B: Parametric with interactions								
Treated	0.257***	-0.163	0.255 ***	0.051	0.272***	0.030	0.232***	0.141
	(0.081)	(0.223)	(0.087)	(0.080)	(0.081)	(0.052)	(0.083)	(0.086)
Conscription		1.321		0.328		-0.264		0.394
		(1.098)		(0.370)		(0.295)		(0.465)
Window		48 months		48 months		48 months		48 months
Pol. Order		2		2		2		2
N		756		689		765		850

Note: Each coefficient is estimated using a separate regression.Political attitude 1-10 is a discrete variable indicating political ideology with values 1 to 10, where 1 is extreme left and 10 is extreme right. Rightwing: trim is a binary variable equal to 1 if the individual reports to have a political ideology more rightwing than the median, and trims the original value of political attitudes of the top and bottom 10percent. Local Politics is a binary variable equal to 1 if the individual ever participates in local politics. Volunteers if a binary variable equal to 1 if the individual reports to ever volunteer. Treated is a binary equal to 1 if the individual was born after the introduction of mandatory conscription. Conscription is the predicted probability of having done military service estimated from the FS regression. All estimations control for:year of survey, age, region, marital status, years of schooling and labor force participation of individual at time of survey, place of residence during childhood and parental education. Panel A presents the results obtained using non-parametric methods following Calonico et al.(2016) for the optimal bandwidth. Panel B presents the results obtained using parametric methods and interacted second order polynomials of the month of birth. Standard errors are clustered at the month of birth.

Table 30: The effect of conscription on support for political parties in Germany: cross-sectional and panel sample, parametric and non-parametric models.

	Cross-section				Panel			
	CDU/CSU		SPD		CDU/CSU		SPD	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	FS	RF	FS	RF	FS	RF	FS	RF
		SS		SS		SS		SS
Panel A: Non-Parametric								
Treated	0.310***	-0.071	0.310 ***	0.055	0.257***	-0.055	0.257***	0.094
	(0.072)	(0.093)	(0.072)	(0.062)	(0.068)	(0.068)	(0.062)	(0.072)
Conscription		-0.229		0.244		-0.321		0.441
		(0.299)		(0.219)		(0.375)		(0.353)
Bandwidth		50.7		50.7		46.6		48.9
N		3831		3831		31575		31575
N Eff. obs.		647		647		5016		5016
Panel B:Parametric with interactions								
Treated	0.251 ***	- 0.099	0.251 ***	0.098	0.243***	-0.078	0.243***	0.177*
	(0.082)	(0.132)	(0.082)	(0.095)	(0.084)	(0.108)	(0.084)	(0.103)
Conscription		-0.293		0.292		-0.397		0.528
		(0.397)		(0.281)		(0.485)		(0.472)
Window		48 months		48 months		48 months		48 months
Pol. Order		2		2		2		2
F-Stat		11.7		11.7		65.6		65.6
N		619		619		5157		5157

Note: Each coefficient is estimated using a separate regression. Treated is a binary equal to 1 if the individual was born after the introduction of mandatory conscription. Conscription is the predicted probability of having done military service estimated from the FS regression. All estimations control for year of survey and age of individual at time of survey. Panel A presents the results obtained using non-parametric methods following Calonico et al.(2016) for the optimal bandwidth. Panel B presents the results obtained using parametric methods and interacted second order polynomials of the running variable (month of birth). All regressions include controls for: year of survey, region, age, marital status, labor force participation and years of schooling of individual at time of survey, place of residence during childhood and parental education. In columns 1-4 the regressions are estimated on the cross-sectional sample, with standard errors clustered at the running variable (month of birth). In columns 5-8 regressions are estimated using the pooled panel data and clustering the standard errors at the individual level.

Table 31: The effect of conscription on right-wing ideology and voting participation in Germany: panel sample.

	Rightwing				Voted			
	(1) FS	(2) RF SS	(3) FS	(4) RF SS	(5) FS	(6) RF SS	(7) FS	(8) RF SS
Panel A: Non-Parametric								
Treated	0.208 ***	-0.008	0.195 ***	-0.029	0.323***	0.058*	0.343***	0.050
	(0.056)	(0.065)	(0.058)	(0.065)	(0.056)	(0.033)	(0.063)	(0.031)
Conscription		0.042		0.149		0.180*		0.147
		(0.308)		(0.339)		(0.108)		(0.097)
Year+Age	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	N	Y	Y	N	N	Y	Y
Bandwidth		53.1		56.6		59.3		47.8
N		11669		11669		7310		6117
N Eff. obs.		1584		1526		979		794
Panel B: Parametric with interactions								
Treated	0.259 ***	0.016	0.236***	-0.016	0.244***	0.055	0.417***	0.027
	(0.091)	(0.096)	(0.089)	(0.101)	(0.081)	(0.050)	(0.089)	(0.044)
Conscription		0.064		0.069		0.124		0.066
		(0.369)		(0.432)		(0.113)		(0.105)
Year+Age	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	N	Y	Y	N	N	Y	Y
Window		48 Months		48 Months		48 Months		48 Months
Pol. Order		2		2		2		2
F-Stat		15.5		11.74		27.57		21.51
N		1424		1306		869		799

Note: Each coefficient is estimated using a separate regression. Rightwing is a binary variable equal to 1 if the individual reports to have a political ideology more rightwing than the median. Voted is a binary variable equal to 1 if the individual voted in the last general election. Treated is a binary equal to 1 if the individual was born after the introduction of mandatory conscription. Conscription is the predicted probability of having done military service estimated from the FS regression. Columns1,2,5,6 include controls for year of the survey and age.Columns 3,4,7 and 8 include controls for: year of the survey, age and region of residence at the time of the survey, marital status, labor force participation, years of schooling, place of residence during childhood and parental education. Panel A presents the results obtained using non-parametric methods following Calonico et al.(2016) for the optimal bandwidth. Panel B presents the results obtained using parametric methods and interacted second order polynomials of the running variable (month of birth). Standard errors are clustered at the individual level to account for within-individual correlation given the panel dimension of the data (Lee and Lemieux, 2010)

Table 32: The effect of conscription on right-wing ideology and voting participation in Germany: parametric methods and 36 months window.

	(1) FS	(2) RF SS	(3) FS	(4) RF SS	(5) FS	(6) RF SS
Panel A: Rightwing						
Treated	0.239 ***	0.055	0.233 ***	0.050	0.223 ***	0.065
	(0.067)	(0.063)	(0.062)	(0.062)	(0.062)	(0.068)
Conscription		0.228		0.215		0.215
		(0.263)		(0.266)		(0.300)
Year + Age	N	N	Y	Y	Y	Y
Controls	N	N	N	N	Y	Y
Window		36 months		36 months		36 months
Pol. Order		1		1		1
F-Stat		12.9		14.1		12.7
N		621		621		571
Panel B: Voted						
Treated	0.351 ***	0.050	0.351 ***	0.051	0.334 ***	0.068
	(0.075)	(0.045)	(0.072)	(0.045)	(0.071)	(0.042)
Conscription		0.113		0.124		0.124
		(0.123)		(0.129)		(0.116)
Year + Age	N	N	Y	Y	Y	Y
Controls	N	N	N	N	Y	Y
Window		36 months		36 months		36 months
Pol. Order		1		1		1
F-Stat		25.8		27.3		25.8
N		453		453		420

Note: Each coefficient is estimated using a separate regression. Rightwing is a binary variable equal to 1 if the individual reports to have a political ideology more rightwing than the median. Voted is a binary variable equal to 1 if the individual voted in the last general election. Treated is a binary equal to 1 if the individual was born after the introduction of mandatory conscription. Conscription is the predicted probability of having done military service estimated from the FS regression. Columns1,2 do not include controls.Columns 3,4include controls for: year of the survey and age. Columns 5, 6 include full set of controls for age, year of survey, region of residence at the time of the survey, marital status, labor force participation, years of schooling, place of residence during childhood and parental education. All estimations use parametric methods and interacted first order polynomials of the running variable (month of birth). Standard errors are clustered at the month of birth level.

Table 33: Placebo regressions with change in cutoff date: +12 months after change introduction of conscription in Germany.

	Cross-section				Panel			
	Rightwing		Voted		Rightwing		Voted	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	FS	RF	FS	RF	FS	RF	FS	RF
		SS		SS		SS		SS
Treated (placebo)	-0.027	-0.121	-0.063	-0.022	-0.068	0.053	-0.024	-0.026
	(0.107)	(0.088)	(0.143)	(0.041)	(0.105)	(0.098)	(0.105)	(0.030)
Conscription		4.41		-0.009		-1.17		0.277
		(16.86)		(0.251)		(3.09)		(0.316)
Year+Age	Y	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Window		48 Months		48 months		48 months		48 months
Pol. Order		2		2		2		2
F-Stat		0.06		1.32		0.55		1.18
N		814		603		1404		868

Note: Each coefficient is estimated using a separate regression. Rightwing is a binary variable equal to 1 if the individual reports to have a political ideology more rightwing than the median. Voted is a binary variable equal to 1 if the individual voted in the last general election. Treated (placebo) is a binary equal to 1 if the individual was born 12 months after the introduction of mandatory conscription, i.e. 12 months after the cutoff date. Conscription is the predicted probability of having done military service estimated from the FS regression. All regressions include controls for: year of the survey, age and region of residence at the time of the survey, marital status, labor force participation, years of schooling, place of residence during childhood and parental education. Regressions include interacted second order polynomials of the running variable (month of birth). In columns 1-4 the standard errors are clustered at the month of birth. In columns 5-8 standard errors are clustered at the individual level to account for within-individual correlation given the panel dimension of the data (Lee and Lemieux, 2010)

5.3 Appendix 3: supplemental figures and tables to paper 3.

Figure 19: Distribution of the propensity scores before and after matching treatment and control group: always treated vs. never treated

(a) Pre-Matching (b) Post-Matching

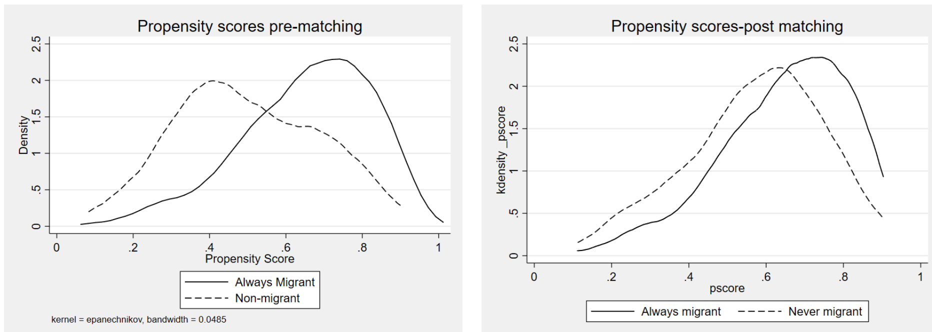

Figure 20: Distribution of the propensity scores before and after matching treatment and control groups: only wave 1 vs. never treated

(a) Pre-Matching (b) Post-Matching

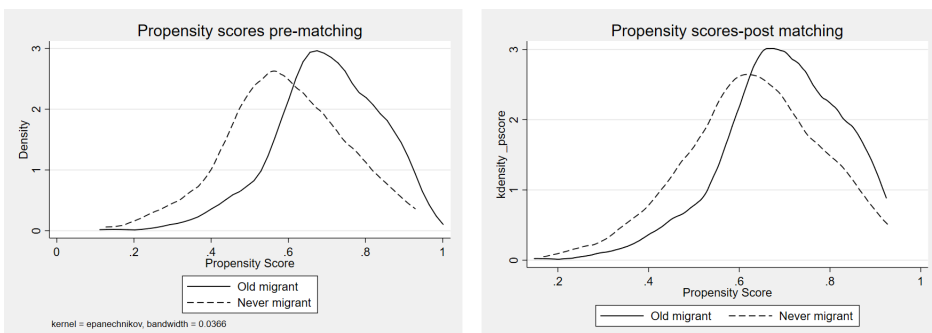

Figure 21: Distribution of the propensity scores before and after matching treatment and control groups: only wave 2 vs. never treated

(a) Pre-Matching (b) Post-Matching

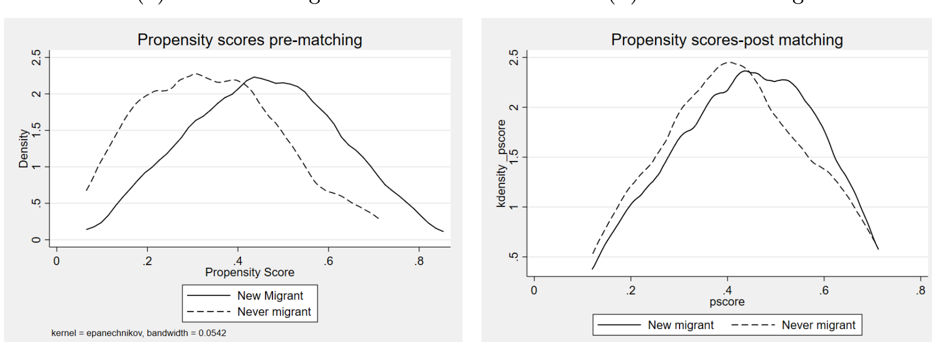

Table 34: The effect of migration on per capita monthly household consumption by treatment type: unmatched and matched sample.

				PSM		
	Total	Food	Non-Food	Total	Food	Non-Food
	(1)	(2)	(3)	(4)	(5)	(6)
A: Always migrant						
Treated	0.078**	0.045	0.171	0.070*	0.048	0.144 **
	(0.038)	(0.041)	(0.053)	(0.040)	(0.045)	(0.057)
Wave	-0.151***	-0.153***	-0.195***	-0.161***	-0.152***	-0.242***
	(0.040)	(0.044)	(0.059)	(0.045)	(0.050)	(0.067)
Treated x Wave	-0.092 *	-0.123 **	-0.079	-0.066	-0.121 *	-0.007
	(0.055)	(0.061)	(0.080)	(0.060)	(0.067)	(0.089)
N	582	582	582	500	500	500
B: Only wave 2						
Treated	0.054	0.044	0.085	0.059	0.052	0.076
	(0.052)	(0.052)	(0.078)	(0.059)	(0.060)	(0.086)
Wave	-0.150***	-0.153***	-0.195***	-0.139**	-0.130**	-0.228***
	(0.040)	(0.044)	(0.060)	(0.046)	(0.051)	(0.069)
Treated x Wave	-0.187 **	-0.225***	-0.120	-0.184 **	-0.238 ***	-0.070
	(0.075)	(0.079)	(0.107)	(0.082)	(0.089)	(0.117)
N	402	402	402	318	318	318
C: Only wave 1						
Treated	0.004	-0.011	0.069	0.011	0.001	0.070
	(0.039)	(0.041)	(0.055)	(0.038)	(0.042)	(0.054)
Wave	-0.151***	-0.153**	-0.194***	-0.151***	-0.141**	-0.235***
	(0.040)	(0.044)	(0.059)	(0.045)	(0.050)	(0.067)
Treated x Wave	0.053	0.046	-0.016	0.047	0.020	0.018
	(0.052)	(0.058)	(0.076)	(0.057)	(0.065)	(0.084)
N	728	728	728	609	609	609

Note:: All coefficients are estimated using a different regression. Columns 1,2 and 3 are estimated on the unmatched sample. Columns 4,5 and 6 are estimated using propensity score matching and observations on the common support. Always migrant is a binary variable equal to one if the household has at least one migrant in both waves, only wave 2 is a binary variable if the household has at least one migrant only in wave 2, and only wave 1 is a binary variable equal to one if the household has at least one migrant only in wave 1. The control group are households with no migrants in either wave. Outcome variables are real monthly per capita consumption expressed in natural logs. All regressions include village level fixed effects. All standard errors are clustered at the household level.

Table 35: The effect of migration on monthly household consumption by treatment type: unmatched and matched sample.

				PSM		
	Total	Food	Non-Food	Total	Food	Non-Food
	(1)	(2)	(3)	(4)	(5)	(6)
A: Always migrant						
Treated	0.113***	0.079*	0.207***	0.097*	0.076	0.171*
	(0.042)	(0.043)	(0.061)	(0.044)	(0.047)	(0.065)
Wave	-0.047	-0.050	-0.090	-0.087*	-0.079	-0.168**
	(0.044)	(0.046)	(0.065)	(0.049)	(0.053)	(0.073)
Treated x Wave	0.112 *	0.081	0.128	0.162 **	0.106	0.226 **
	(0.059)	(0.063)	(0.087)	(0.065)	(0.070)	(0.096)
N	582	582	582	500	500	500
B: Only wave 2						
Treated	0.077	0.067	0.107	0.057	0.050	0.073
	(0.053)	(0.054)	(0.079)	(0.060)	(0.062)	(0.088)
Wave	-0.047	-0.050	-0.091	-0.101*	-0.092*	-0.190
	(0.044)	(0.046)	(0.066)	(0.052)	(0.056)	(0.076)
Treated x Wave	0.006	-0.032	0.073	0.063	0.010	0.179
	(0.079)	(0.081)	(0.114)	(0.091)	(0.095)	(0.129)
N	402	402	402	318	318	318
C: Only wave 1						
Treated	0.053	0.035	0.123**	0.068*	0.058	0.127**
	(0.037)	(0.039)	(0.055)	(0.040)	(0.043)	(0.058)
Wave	-0.046	-0.050	-0.090	-0.082*	-0.074	-0.167
	(0.044)	(0.046)	(0.065)	(0.048)	(0.053)	(0.072)
Treated x Wave	-0.025	-0.008	-0.099	0.001	-0.002	-0.028
	(0.056)	(0.059)	(0.083)	(0.061)	(0.066)	(0.090)
N	728	728	728	609	609	609

Note:All coefficients are estimated using a different regression. Columns 1,2 and 3 are estimated on the unmatched sample. Columns 4,5 and 6 are estimated using propensity score matching and observations on the common support. Always migrant is a binary variable equal to one if the household has at least one migrant in both waves, only wave 2 is a binary variable if the household has at least one migrant only in wave 2, and only wave 1 is a binary variable equal to one if the household has at least one migrant only in wave 1. The control group are households with no migrants in either wave. Outcome variables are real monthly per capita consumption expressed in natural logs. All regressions include village fixed effects All standard errors are clustered at the household level.

Table 36: The effect of migration on household composition and annual total income by treatment type: matched sample.

	HH size	Nr adults	Nr Men	Annual Income
	(1)	(2)	(3)	(4)
A: Always migrant				
Treated	0.055	0.398***	0.032	0.430***
	(0.184)	(0.119)	(0.133)	(0.094)
Wave	0.359*	0.694 ***	0.044	1.634***
	(0.192)	(0.146)	(0.134)	(0.108)
Treated x Wave	1.294 ***	0.873 ***	0.547 ***	-0.200
	(0.245)	(0.184)	(0.179)	(0.142)
N	500	500	500	500
B: Only wave 2				
Treated	0.083	0.129	0.029	0.063
	(0.255)	(0.158)	(0.174)	(0.113)
Wave	0.221	0.574***	-0.010	1.661***
	(0.194)	(0.149)	(0.136)	(0.111)
Treated x Wave	1.210 ***	1.019 ***	0.498 **	-0.004
	(0.321)	(0.229)	(0.226)	(0.177)
N	318	318	318	318
C: Only wave 1				
Treated	0.161	0.389***	0.269	0.430***
	(0.167)	(0.111)	(0.121)	(0.087)
Wave	0.325*	0.665***	0.043	1.63***
	(0.187)	(0.144)	(0.133)	(0.108)
Treated x Wave	-0.124	-0.211	-0.024	-0.480 ***
	(0.230)	(0.180)	(0.164)	(0.137)
N	609	609	609	609

Note:: All coefficients are estimated using a different regression. Columns 1,2 and 3 are estimated using propensity score matching and observations on the common support and a diff-in-diff estimator. Always migrant is a binary variable equal to one if the household has at least one migrant in both waves, only wave 2 is a binary variable if the household has at least one migrant only in wave 2, and only wave 1 is a binary variable equal to one if the household has at least one migrant only in wave 1. The control group are households with no migrants in either wave.HH Size is a variable indicating the number of members in the household. Nr adults is indicating the number of household members aged between 15 and 65. Number of men is the number of men in the household. Annual income is the natural log of real household yearly income. All standard errors are clustered at the household level.