JANI-PETRI LAAMANEN

Essays in Public and Labour Economics

Estimating causal effects using regional variation

■

UNIVERSITY OF TAMPERE

JANI-PETRI LAAMANEN

Essays in Public and Labour Economics

Estimating causal effects using regional variation

UNIVERSITY
OF TAMPERE

Cover design by
Mikko Reinikka

Layout by
Sirpa Randell

441 729
Painotuote

Acknowledgements

I think I had my first serious thoughts of becoming a doctoral student when I was working as an economics teaching assistant at the University of Tampere and as an intern at the Bank of Finland during my master's studies. The cause of those thoughts must have been the combination of the encouraging attitude of the personnel and visiting scholars of both institutions, the intellectual challenges I faced at work, and young student's urge to learn and understand more about the society. So I would first like to thank those who worked at the economics department of the University of Tampere and the research department of the Bank of Finland back then. Without you, I might have chosen a different, and probably a much less satisfying, career.

Professor Jari Vainiomäki was my supervisor already when I was working on my master's thesis. Since then, he has supervised my licentiate thesis and this doctoral dissertation. Jari has supported and guided my work during the years. I would like to thank Jari for his patience and understanding also at those (not infrequent) times when I have stubbornly done things my way. For the recent years, I have been working in Professor Jukka Pirttilä's projects. It has been great to work for Jukka, because he has respected my ideas and has let me work freely on topics that interest me. This thesis, as well as my other works, have benefited greatly from the comments that Jari and Jukka have given.

I would like to thank the two pre-examiners of this thesis, Docent Kari Hämäläinen from the Government Institute for Economic Research (VATT) and Professor Bas van der Klaauw from the VU University Amsterdam. I have been able to improve the thesis considerably because of their comments. Many of their comments have been positive and very encouraging, which is not always the case with comments that you get. My special thanks go to Professor Kaisa Kotakorpi who has co-authored two essays of this thesis. We have known each other already since master's studies and collaborated in the economics student union and later in joint projects. Kaisa has read and given valuable comments to most of my academic writings. Not less important has been Kaisa's advice on how to effectively transform a research idea into a final version of a research paper. I am fortunate to have been able to work with Kaisa and to be able to continue working with her.

Thanks to the financial support, I have been able to study, choose my research topics freely and work full-time on my research. I am very lucky to have had the Yrjö Jahnsson Foundation, the Finnish Doctoral Programme in Economics and the Academy of

Copyright notice

Abstract

The aim of almost any econometric study is to utilise the variation in one or more explanatory variables to identify their effects on, or associations with, some outcome variable. Explanatory variables may vary across cross-sectional units, over time or both in the case of panel data. Any variation across units of observation in both explanatory variables and the outcome variable allows the researcher to estimate simple statistical associations. Since today's econometricians are usually interested in causal inference, simple associations between variables are more and more seldom the key output of an econometric study.

This thesis consists of an introductory chapter and four empirical essays in public and labour economics. In each of the essays, the aim is to estimate causal effects using variation in the explanatory variables of interest across regions. In the Introduction, issues related to the use of regional data in identifying causal effects is discussed. The Introduction provides a review of cases in which it is beneficial, and sometimes imperative, to use regional data to estimate causal effects. Next, the issues related to estimation of causal effects in different cases involving regional data are discussed. There are various econometric approaches that can be taken when studying regional data, and the choice of an appropriate approach is crucial. There are some problems yet to be solved related to methods. Especially the literature on how to estimate causal effects in the presence of externalities, or peer effects, is still relatively young and still needs to evolve before satisfying methods are available to study some research questions. The discussion in the Introduction is more general in nature, because most of the conclusions extend to any case of group-level data.

The empirical essays apply methods reviewed in the Introduction to four research questions in public and labour economics. Two of the essays examine impacts of public policies and the other two essays examine causal effects of factors that are market outcomes rather than policy instruments. In the first essay, the effects of a pre-filled tax return on deduction claims are studied. It is found that receiving a pre-filled income tax return leads to a reduction of over one-fourth in the number of individuals claiming for deductions. In the second essay, the effects of home-ownership on the labour market are examined. The main finding is that while home-owners are less likely to experience unemployment, as documented by many earlier studies, an increase in the rate of home-ownership causes regional unemployment to rise. In the third essay, the effects of public health care spending on subjective well-being are estimated. It is shown that high

expenditures in health care have a positive effect on citizens' life satisfaction. In the last essay, the effects of regional labour market conditions on interregional migration are estimated. The results indicate that hires from unemployment and job separations leading to unemployment have sizeable effects on migration. The effects of hires from and separations to other labour market states, while statistically significant, appear smaller. Changes in the structure of employment were not found to affect migration.

Tiivistelmä

Miltei kaikkien ekonometristen tutkimusten tavoite on hyödyntää yhden tai useamman selittävän muuttujan vaihtelua sen selvittämiseksi, mikä on niiden vaikutus tai yhteys johonkin lopputulemamuuttujaan. Selittävät muuttujat saattavat vaihdella poikkileikkausyksiköiden, ajanhetkien tai paneeliaineiston tapauksessa molempien välillä. Mikä tahansa havaintoyksiköiden välinen vaihtelu sekä selittävissä muuttujissa että lopputulemamuuttujassa mahdollistaa yksinkertaisten tilastollisten yhteyksien estimoinnin. Koska nykyaikana ekonometrikot haluavat yleensä tehdä kausaalivaikutuksia koskevia johtopäätöksiä, pelkät yksinkertaiset yhteydet muuttujien välillä ovat yhä harvemmin ekonometrisen tutkimuksen päätuloksia.

Tämä väitöskirja koostuu johdantoluvusta ja neljästä julkistaloutta ja työn taloustiedettä koskevasta empiirisestä esseestä. Jokaisessa esseessä tavoitteena on estimoida kausaalivaikutuksia hyödyntäen alueiden välistä vaihtelua kiinnostuksen kohteena olevissa selittävissä muuttujissa. Johdannossa keskustellaan alueaineistojen käytöstä kausaalivaikutusten identifioinnissa. Ensin luodaan katsaus tapauksiin, joissa alueaineistojen käyttö on hyödyllistä tai jopa välttämätöntä kausaalivaikutusten estimoimiseksi. Seuraavaksi keskustellaan kausaalivaikutusten estimointiin liittyvistä kysymyksistä tapauksissa, joissa hyödynnetään alueaineistoja. Kun alueaineistoja käytetään tutkimuksessa, on olemassa useita mahdollisia ekonometrisia lähestymistapoja, ja oikean lähestymistavan valinta on olennaista. Menetelmiin liittyy joitakin toistaiseksi ratkaisemattomia ongelmia. Erityisesti kirjallisuus, joka koskee kausaalivaikutusten estimointia ulkoisvaikutusten, tai vertaisvaikutusten, läsnä ollessa, on verraten nuorta ja sen tulee kehittyä ennen kuin joidenkin tutkimuskysymysten tutkimista varten on käytettävissä tyydyttäviä menetelmiä. Johdannossa oleva keskustelu on luonteeltaan edellä kuvattua yleisempää, koska suurin osa johtopäätöksistä voidaan yleistää koskemaan mitä tahansa tapausta, johon liittyy ryhmätason aineistoja.

Empiirisissä esseissä sovelletaan johdannossa esiteltyjä menetelmiä neljään julkistalouden ja työn taloustieteen tutkimuskysymykseen. Kahdessa esseessä tutkitaan politiikkojen vaikutuksia ja kahdessa muussa esseessä tutkitaan sellaisten tekijöiden kausaalivaikutuksia, jotka ovat markkinalopputulemia. Ensimmäisessä esseessä tutkitaan esitäytetyn veroilmoituksen vaikutuksia verovähennysten hakemiseen. Tuloksena on, että esitäytetty veroilmoitus vähentää vähennyksiä hakeneiden määrää neljänneksellä. Toisessa esseessä tutkitaan omistusasumisen vaikutuksia työmarkkinoihin. Päätulos on että vaikka, kuten aiemmissakin tutkimuksissa on

havaittu, omistusasujat kokevat työttömyyttä muita harvemmin, omistusasumisen lisääntyminen johtaa työttömyyden lisääntymiseen aluetasolla. Kolmannessa esseessä estimoidaan julkisten terveydenhuoltomenojen vaikutusta subjektiiviseen hyvinvointiin. Tutkimuksessa osoitetaan, että korkeammat terveydenhuoltomenot lisäävät kansalaisten elämään tyytyväisyyttä. Viimeisessä esseessä estimoidaan alueiden työmarkkinaolosuhteiden vaikutusta alueiden väliseen muuttoliikkeeseen. Tulokset osoittavat, että työttömien työllistymisellä ja työttömäksi joutumisella on merkittäviä vaikutuksia muuttoliikkeeseen. Työllistymisellä muista työmarkkinatiloista ja työsuhteiden katkeamisilla, jotka eivät johda työttömyyteen, on pienempiä, joskin tilastollisesti merkitseviä vaikutuksia. Työllisyyden rakennemuutoksilla ei havaita olevan vaikutusta muuttoliikkeeseen.

Contents

Chapter 1
Introduction

This thesis consists of four empirical essays in public and labour economics. Two of the essays examine impacts of policies and are thus relevant from the perspective of policy-making. The effects of public health care spending on subjective well-being and the effects of a pre-populated tax return on deduction claims are studied. The other two essays examine causal effects of factors that are market outcomes rather than policy instruments. Namely, the effects of home-ownership on the labour market and the effects of regional labour market conditions on interregional migration are estimated. Various policies affect households' housing tenures and the results of the home-ownership essay are informative of the indirect and often neglected consequences of these policies. The essay on interregional migration adds to the understanding of how geographical mobility of the labour force is related to labour market dynamics. In each of the essays, the aim is to estimate causal effects using variation in the explanatory variables of interest across regions. In this introduction, I discuss the issues related to use of regional data in identifying causal effects.

The aim of almost any econometric study is to utilise the variation in one or more explanatory variables of interest to identify their effects on, or associations with, some outcome variable. Explanatory variables may vary across cross-sectional units, over time or both in the case of panel data. Cross-sectional units typically range from geographical entities such as countries or administrative regions to individuals, households or firms. Any variation across units of observation in both explanatory variables and the outcome variable allows the researcher to estimate simple statistical associations. Since today's econometricians are usually interested in causal inference, simple associations between variables are more and more seldom the key output of an econometric study. It is crucial from the point of view of policy that economic studies are able to produce information on causal impacts rather than on mere statistical relationships. In most cases, causal inference requires theoretical reasoning on the subject of the study and other information as to why the explanatory variable varies across observations and, in particular, how this variation might be related to variation in the outcome variable. In other words, it is critical to assess whether variation in the outcome itself has an impact on (or is otherwise related to) the variation in the explanatory variable. If this is the case, simple estimates of the causal effect are contaminated by endogeneity. A large share of the recent econometric literature focuses on the ways to circumvent the endogeneity

problem. A variety of methods have been proposed and the appropriateness of a method depends on the specific research question and setup.

The case of interest in the essays of this thesis is one in which one or more of the key regressors or control variables are regional aggregates. The outcome variable may or may not be similarly aggregated. Thus, the equation to be estimated is either

$$y_{ijt} = \alpha x_{ijt} + \beta x_{jt}^r + \gamma' Z_{ijt} + \epsilon_{ijt}, \tag{1}$$

if the outcome is measured on the disaggregated (individual) level or

$$y_{jt} = \beta x_{jt}^r + \gamma' Z_{jt} + \epsilon_{jt}, \tag{2}$$

if all variables are regional aggregates. In both equations, y is the outcome variable. Variable x is the individual-level variable of interest and x^r is the regional-level variable of interest. Z is a vector of individual-level and regional-level control variables including a constant and ϵ is the error term that is allowed to include a regional component. Note that in some cases of Equation (1), we are not primarily interested in parameter α but only in β. Then variable x is included as an individual-level control variable and is part of Z. In some other cases, we are primarily interested in parameter α and variable x^r is a regional-level control variable and thus part of Z.

Although I focus on the case of data on multiple regions within one country, most of the discussion applies to analysis of cross-country data as well.[1]

1.1   Reasons for the use of regional data in econometric research

Availability of microdata and the computational capacity to analyse it have been rapidly improving in recent decades. This has led most researchers to use micro-level datasets whenever possible. However, regional data is still often used and, in some cases, it should be used to identify the desired parameters. I will next discuss some of the most prevalent reasons to use regional data in econometric research. The main focus is on situations where regional variation in policy is exploited, but I also discuss identification in the case of no identifying policy variation.

---

1   Some parts of the discussion are even more general: They apply to all cases in which data on group averages is used. To keep the focus on the particularities of regional variation and practical examples, I will not generalise the discussion to the case of any group-level data.

1.1.1 Availability of data and the level of variation

The most obvious reason for the use of aggregated region-level data is the lack of suitable microdata. In some cases, microdata is available but at a considerable cost, and if the focus of the research is not on the empirical analysis, aggregated data may be used to illustrate the empirical relationships between variables.

Some key economic concepts are not well defined at the micro-level. This is often the case in macroeconomics. Measures such as GDP are important from the policy perspective, but they cannot be disaggregated to e.g. the level of an individual citizen. In the case of GDP this is especially true with the public sector spending component. Then the options are either to use microdata to obtain results on some substitute measure or to use aggregated data. Often researchers choose to do the latter to be able to produce relevant results and avoid the need to examine the connection between the substitute micro-variables and the aggregate variables of interest. For example, there is an empirical literature on the effect of GDP (or economic growth) on subjective well-being (for two recent much-cited contributions, see Stevenson and Wolfers, 2008 and Di Tella, MacCulloch and Oswald, 2006). Although there is no lack of data on individual or household income and subjective well-being, many studies in this field prefer to examine the effects of GDP because it is a key target of macroeconomic policies. GDP also includes the public spending component which would be difficult to measure at the micro-level since it is very hard to *a priori* identify the citizens who benefit from it.[2] Therefore, estimated impacts of GDP can be interpreted as the impacts that policies affecting GDP might have on well-being of individuals.

There are variables that are defined at the micro-level but do vary between regions or countries only. In modern microeconometric and policy evaluation studies, regional variation in policy is often the variation of interest. Prominent examples are region-specific policies such as local tax rates determined by local governments. All individuals living in the same region face the same local tax policy and therefore there is only variation in these tax rates across regions. Obviously, when there is no within-region variation in an explanatory variable, there is no difference between individual-level information and regional information on this variable.[3]

1.1.2 Identification of causal effects

Many variables have both regional and individual components. For example, income taxes faced by individuals may consist of a proportional local tax and a progressive national

---

2 This is often true for public goods as well because of the externalities created by them. Use of aggregate data in the case of externalities will be discussed later in this introduction.

3 However, the level of variation needs to be taken into account in calculating the standard errors of the coefficient estimates, as I will discuss later.

tax. In some studies, it is assumed that the individual component of an explanatory variable is endogenous, but the regional variation in the variable is exogenous. Based on this assumption, regional variation can be used to identify causal effects. In the tax example, the individual component of the tax rate depends on income and thus cannot as such be used, for instance, to estimate the elasticity of taxable income with respect to tax rates. However, the risk of an endogeneity problem is much lower if one uses only the tax rate variation across regions to identify the elasticity. Although it is not in general correct to say that regional variation cannot be endogenous when running individual-level regressions, in many cases, the endogeneity problem can be credibly alleviated by using regional rather than individual variation in the explanatory variable. An example of such strategy is the study by Blundell *et al.* (2003) who use differential regional changes in housing benefits as a source of exogenous variation in disposable income to estimate models of labour force participation.

A similar but differently grounded case is one in which cross-regional variation in an explanatory variable is considered as a source of exogenous variation in the same variable between units of observation. In this case, regional average of the explanatory variable can be used as an instrumental variable. For example, van Leuvensteijn and Koning (2004) and, following them, Munch *et al.* (2008) use regional rate of home-ownership as an instrument for individual-level home-ownership. Except for a mention of a 'supply effect' in Munch *et al.* (2008), the authors do not discuss the mechanism through which the regional average affects individuals. I presume that they use the regional home-ownership because it may have an effect on individuals' home-ownership instead of getting rid of the presumably endogenous individual variation in the way I discussed earlier.

1.1.3 Externalities

Sometimes the explanatory variables of interest are both defined and measured at the individual level and they also vary between individuals, but there are impacts at a more aggregated level. In the case of regional externalities, individuals' outcomes are affected by the values of regressors of other individuals residing in their region. Regardless of the primary parameters of interest, presence of regional externalities usually means that variables measured at the regional level are needed in the analysis. This is obviously the case in studies that attempt to estimate the externalities. As Imbens and Wooldridge (2009) note, there are fields of study that consider interaction effects (externalities) as the phenomenon of primary interest (e.g. studies on peer effects). In some other studies, regional information is used because ignoring it in the presence of externalities would lead to omitted variable bias in the parameter estimates of primary interest. There is yet another group of studies that estimate both individual effects and externalities.

The branch of econometric literature that focuses on estimation of returns to education has acknowledged external effects and there have also been attempts to estimate them (see e.g. Acemoglu and Angrist, 2000). It has been argued that a higher level of education may lead to higher earnings for not only the ones who educate themselves but also for other individuals in the same local labour market. Some studies on labour market programmes (Crépon *et al.*, 2013; Blundell *et al.*, 2004) have used regional data to control for and identify labour market externalities. In the labour market, externalities are due to so-called displacement or spillover effects i.e. employment of individuals affects employment of other individuals in the same region. Recent surveys of programme evaluation literature note the possibility of external effects in the labour market. Abbring and Heckman (2007) discuss the interaction issue at length in the context of labour markets and Imbens and Wooldridge (2009) relate interactions to programme evaluation in general. Increased attention to externalities means that micro-level datasets are more often augmented with regional information or the whole analysis is conducted at the regional level. I will discuss the estimation issues caused by externalities in section 1.2.4.

1.2 Regional data and estimation of causal effects

In what follows, I discuss the issues related to estimation of causal effects in different cases involving regional data. I do not cover methods of `spatial econometrics' such as models with spatial autocorrelation. Instead, the focus is on typical econometric methods of public and labour economics applied to regional or both to regional and individual-level data. I attempt to cover the different cases in which regional data is used that are discussed in Section 1.

1.2.1 Policy reforms

Policy reforms that are targeted to specific regions are often exploited to identify causal effects.[4] Given that a reform is implemented only in some regions at a time, causal effects of that reform can potentially be estimated by using simple difference-in-differences (DID) estimation. A prerequisite for DID is that the outcome variable changes similarly in time in the reform and non-reform regions *in the absence* of the reform. This means that there is a constant (or possibly zero) difference in the outcome between the two sets of regions that only changes due to the reform that is analysed. Equality (unconditional or conditional on some control variables) of the time paths between the regions can be tested when there is data from more than one pre-reform

---

4   Only aspects of policy reforms relevant to the essays in the thesis are covered here. A thorough and more technical review of the methods and various special cases is Imbens and Wooldridge (2009).

period. However, after the reform, the outcome in the absence of the reform can only be observed in the non-reform regions. This constitutes a problem for identification if the reform region was actually chosen because an exceptional change (compared to the non-reform region) in the outcome was anticipated to take place there. An example of such a policy reform would be a labour market programme implemented in regions where mass lay-offs are expected to take place. The first essay of this thesis examines the effects of a pre-populated tax form on deduction claims by exploiting a reform that was implemented in different regions in different years.

Evaluations of policy reforms are able to guide future policy especially if the reform or its close variant is replicable somewhere else or if the reform can be extended or reversed. But sometimes reforms can be used to estimate the causal effects of their outcomes. In such case, a reform is used to construct instrumental variables for its outcome. As a result, the effect of that outcome on some other outcome can be estimated. In the simplest case of such an instrumental variables strategy, the first stage model is a DID model of the effects of the reform on the explanatory variable. For example, Acemoglu and Angrist (2000) use state schooling law reforms in the US as instruments to state schooling levels. Schooling levels, in turn, are used as an explanatory variable for individual earnings to capture external effects of education. Similar strategy is adopted in the second essay of this thesis, which studies the effects of regional home-ownership on labour market outcomes.

1.2.2 Regional variation in policy in the absence of a reform

In some cases there is no policy reform, but the policy varies across regions (and possibly over time as well). Local tax rates and local public services are examples of such policies. To obtain credible estimates of the effects, the variation needs to be exogenous. To be more specific, variation in the explanatory variable should be independent of outcomes *conditional on* the covariates. It is thus possible to identify the causal effect of the policy if there is enough information on the determinants of the policy variation and their possible relationship with the outcome. Depending on the case, determinants of policy are either controlled for or used as instruments to make the claim of conditional independence credible. While most policies vary for a reason and the variation may thus be related to the outcome, some parts of the variation can often be considered exogenous. Determinants of a policy are usually better known than determinants of some other type of explanatory variable. This is why controlling for the (endogenous) determinants or using exogenous determinants as instruments is usually a more credible strategy in the case of policy variation than in the case of variation in some other explanatory variable of interest. Two examples of studies utilising policy rules to isolate the exogenous part of variation in policy are Gordon (2004) and Dahlberg *et al*. (2008). Both of the studies look at the effects of federal grants on spending behaviour

of administrative units receiving the grants. Kotakorpi and Laamanen (2010) (the third essay of this thesis) study the welfare effects of local public health care spending. In the absence of specific policy rules, they check the robustness of their results by controlling for the (political) determinants of spending variation.

1.2.3 Regional variation in other than policy variables

In the absence of policy reforms and policy variation, determinants of which are known, instrumental variables and credible controlling strategies are sometimes available. However, a large share of econometric research has been conducted with data in which there is neither exogenous variation in the explanatory variables nor proper instruments. Most often, the problem of endogeneity then arises because the explanatory variables are determined as outcomes of the same local economic system (markets) as the outcome variable. Most macroeconomic relationships are of this type. In such a case, there is not enough information on the economic system that determined the variation in the explanatory variable and, thus, that information cannot be credibly controlled for.

These problems have led to development of methods that can, under some conditions, tackle endogeneity. The instrumental variables method, using lagged regressor as the instrument for the present value of that regressor, is one such method. With panel data, one can control for regional fixed effects and time effects. However, already in the early 1980's, Nickell (1981) showed analytically that a dynamic panel data model with fixed effects does not yield consistent estimates if there is autocorrelation in the dependent variable. A widely used solution to these problems is to use a dynamic panel GMM model by Arellano and Bond (1991) and its later variant by Arellano and Bover (1995) and Blundell and Bond (1998). These models are based on the idea of using lagged data on the regressors as instruments for the current data. Given that the constructed instruments are both relevant and valid, the causal relationship between the explanatory variables and the outcome variable can be estimated. The dynamic panel GMM models have been extensively used in empirical macroeconomics, international trade and development economics.

The basic variants of a dynamic panel GMM model work when there is a large number of cross-sectional units and the panel is short. This is often the case with panels of regions. However, the number of regions may not always be very large and consequently the asymptotic properties of the GMM estimators may not hold, resulting in biased parameter estimates and inference problems.[5] Furthermore, dynamic panel GMM methods require some assumptions that may be hard to defend and various choices regarding especially the set of instruments have to be made. On the positive

---

5   In the last essay of this thesis, there are 342 cross-sectional units (pairs of regions) which is a relatively large number compared to many studies using regional or cross-country data and adopting the dynamic panel GMM approach.

side, dynamic panel GMM methods can be used in many cases where other strategies for identification are not available. The last essay of this thesis adopts dynamic panel GMM methods to estimate the effects of labour market dynamics on interregional migration.

1.2.4 Externalities

When individual outcomes are affected by regressors or outcomes of other individuals residing in the same region, there are region-level externalities. Presence of external effects has severe consequences for the estimation of causal effects (see also Imbens and Wooldridge, 2009; Abbring and Heckman, 2007). Estimating the direct individual-level effect using conventional methods leads to biased results. This is because the external effect changes outcomes of the individuals who should be representative of the outcomes in other levels of the explanatory variables. For example, assume that high education has both a positive direct effect and a positive external effect on the less-educated individuals. Once some individuals attain higher education, their earnings increase. But so do the earnings of the less-educated. Thus, the earnings differential between the two groups is increased by the direct effect but decreased by the external effect. The true effect of education on individuals' earnings is higher than that estimated by using the conventional approach.[6] The problem caused by externalities can be seen by looking at Equation (1) and assuming that true values of both α and β are non-zero. Including only the individual-level variable causes an omitted variable bias. Acemoglu and Angrist (2000) tackle the problem of externalities of education by including both individual and region-level education levels in their earnings regression.

In the education example above, using individual-level data underestimates the total (direct plus external) effect of education on earnings. This is both because the contribution of the external effect is ignored and because the comparison group's earnings level is increased by the external effect. Estimating the relationship from regional data would give an unbiased estimate of the total effect. Both the direct and external effect can be estimated by combining regional and individual-level data. If the external effect and the direct effect are of different sign, using individual-level data overestimates the total effect. For example, consider a situation with two workers and one job so that one of the workers is working and the other one is unemployed.[7] Assume that the unemployed worker faces some change that makes her lower her reservation wage dramatically or search for a job very actively. It is likely that, at least in the course of time, she will eventually take the only job. Estimating the effect of the change that

---

6   It should be noted that the estimated effect 'includes' the negative external effect on the earnings differential and, in a sense, gives an estimate of the effect on the earnings *differential* of the increase in education. It is clear that estimating such a parameter is rarely, if ever, the aim of analysis.

7   Imbens and Wooldridge (2011) have presented a similar 'fixed-number-of-jobs' example.

occurred on employment probability gives a large positive effect. Notice that the other worker who lost her job (due to the external effect from job competition) is now in the comparison group. Due to the external effect, the estimate is far from the real total effect which equals zero because the total employment did not change. The bias can be avoided, again, by using regional-level data. Blundell *et al*. (2004) solve the externality problem in estimating the effect of a labour market programme by using a control group from regions where there was no labour market programme. Crépon *et al*. (2013) estimate various models with differing degrees of flexibility. The most flexible model allows the share of treated individuals in the region of residence to affect employment of treated and non-treated individuals differently. The above discussion and examples illustrate that regional-level information should be used, and more and more often is used, if external effects are expected to exist.[8] The idea of externalities is central in the second essay of this thesis, which explores the regional employment externalities caused by home-ownership. In the third essay, the key regressor, health care spending, is measured in the municipality-level which means that possible externalities on individuals residing in the municipality are captured by the estimates.[9]

1.2.5 Clustering of standard errors

The final point to be made is related to any model in which the outcome variable is measured at a more disaggregated level than the explanatory variable of interest. This is often the case when regional variation in the explanatory variable is exploited. The most common example in empirical public and labour economics is individual-level data augmented with a regional level regressor. Since Moulton (1986), it has been acknowledged that, in the presence of region-level stochastic errors, the standard errors of the coefficient estimate of aggregate explanatory variables are not correctly calculated by the standard procedure. Thus, in estimation of any model that includes a regional-level regressor, clustering of errors needs to be taken into account. Nowadays, this can be easily done in most statistical analysis softwares. Analyses that estimate effects of policy reforms or use policy reforms as instruments often need to take clustering into account, as pointed out by Bertrand, Duflo and Mullainathan (2004). Angrist and Pischke (2008) end their discussion on clustering by pointing out that (macroeconomic) studies using data at more aggregated level do not face the problem of clustering.

---

8   Heckman *et al*. (1999) adopt a different approach and use a general equilibrium model to illustrate the externalities of education.

9   In a similar fashion, coefficient estimates in the fourth essay capture any externalities at the province level although we do not suspect externalities in that context.

1.3 Summaries of the essays

1.3.1 The effects of pre-populated tax return on deductions: Evidence from a natural experiment

In the first essay, which is joint work with Kaisa Kotakorpi, we estimate the effect of the details of the income tax filing system on the behaviour of taxpayers, utilising data from a policy experiment that took place in Finland in the 1990's. The policy experiment was conducted by adopting a pre-filled tax return instead of the previously used blank form in some regions. We find that adopting a pre-filled income tax return leads to a significant reduction in the number of individuals claiming for deductions. This finding may be due to complexity or salience effects.

Recent research on the effects of taxation has paid attention to features of the tax system that are irrelevant from the point of view of traditional public economics theory. The literature considers aspects such as the way of collecting taxes, the information that taxpayers possess on taxes and the visibility of the financial incentives created by taxes (for a brief review, see Chetty, 2011). A recent contribution by Kleven *et al.* (2011) finds that whether particular (positive or negative) type of income is reported to the tax authority by a third party and pre-filled in the tax return is of importance for the degree of tax evasion related to that type of income. According to the results, there is much more evasion in the self-reported income. The first essay is closely related to the research of Kleven *et al.* (2011). We allow receiving a pre-filled return to have an effect on the self-reports as well. More specifically, we look at the effect on deductions, which is the most important category of self-reported information in the current Finnish system. The Finnish experiment with a pre-filled tax return in the late 1990's is suited for examining the effect. We argue that there are several theoretical reasons for the effect including complexity, salience, opportunity cost of time and fixed costs of tax compliance. We estimate large negative effects for the pre-filled return on self-reported deductions. Over 10% less of the recipients of the pre-filled return did claim any discretionary deductions, compared to the comparison group of recipients of the blank form.

Our identification strategy of the causal impact is based on the region-specific adoption of the pre-filled tax return. The experiment was conducted in different municipalities in different years, thus creating cross-municipality variation in the treatment (the pre-filled return). Thus, our key assumption is that the choice of the experiment municipalities is not related to the coming change in the deduction claims in these municipalities. Thus, by choosing an appropriate control group from the non-experiment municipalities, we are able to identify the causal effect.

Jani-Petri Laamanen

1.3.2 Home-ownership and the labour market: Evidence from rental housing market deregulation

In the second essay, I use a regional policy reform to identify the labour market effects of the regional rate of home-ownership. Finnish rental housing market deregulation took place in the early 1990's and started as a regional experiment: New dwellings were deregulated in the northern and central counties before the south. Deregulation markedly changed the incentives to supply rental housing and thus created identifying variation in home-ownership across regions. I find that while home-owners are less likely to experience unemployment, as documented by many earlier studies, an increase in the rate of home-ownership causes regional unemployment to rise.

Since Oswald (1996), several studies have found that high rates of home-ownership are associated with higher unemployment. However, more recent micro-evidence on the link between housing tenure and unemployment suggests that home-owners have relatively favourable labour market outcomes (Coulson and Fisher, 2009; Munch *et al.*, 2008; Munch *et al.*, 2006; van Leuvensteijn and Koning, 2004; Flatau *et al.*, 2003). There thus seems to be a contrast between the micro- and macro-level results. I argue that, given that home-ownership creates externalities in the labour market, the two effects that create the results may co-exist. Using identifying variation in regional home-ownership rates, I estimate the external effects of regional home-ownership on the labour market and find that they are statistically significant and negative. At the same time, individual owner occupiers are less likely to be unemployed than non-owners.

Oswald (1999) and more recently Blanchflower and Oswald (2013) have suggested various mechanisms through which the externalities may operate.[10] My data permits me to test two hypotheses on the mechanisms that have not been previously discussed in the literature on home-ownership and unemployment. The hypotheses are based on recent macroeconomic and labour economics research. I argue that home purchases increase households' indebtedness, which may be negatively reflected in their consumption expenditures and positively in their labour supply. If households decrease consumption as a result of buying their homes, negative labour market consequences follow through a drop in the aggregate demand. The second mechanism is related to labour market displacement. If home purchases are associated with increases in labour supply, the employment consequences are likely to be positive. However, there may be also negative employment consequences on other workers (externalities) through displacement effects (see Crépon *et al.*, 2013; Blundell *et al.*, 2004).

My results suggest that home-ownership may create negative labour market externalities by holding back home-owners' consumption. Unfortunately, as far as I am concerned, there are no studies on the effect of buying a home or taking out a

---

10   It should be noted that some of the mechanisms operate in the long-run and some also in the short-run. My study examines the short-run because I estimate immediate effects of home-ownership.

mortgage on consumption expenditures. I find indirect evidence on the consumption mechanism. I find that increases in home-ownership are reflected negatively on local employment in sectors that produce non-tradable goods and services but not in sectors that produce tradable goods. My way of interpreting this result is similar to that of Mian and Sufi (2012), who look at the effects of local house price crashes on employment. They argue that, since regional decline in house prices during the recent financial crisis is reflected in local employment in the non-tradable but not in the tradable sector, the effect is likely to operate through consumption demand. Since I also find results that are in line with the displacement hypothesis, I conclude that both the consumption and the displacement mechanism may contribute to the commonly observed positive correlation between home-ownership and unemployment.

In this essay, a policy reform has created regional variation in the key regressor. As in the first essay, it is crucial that the reform regions were not selected based on the anticipated change in the outcome variable. This is my key identifying assumption and I describe the political process that led to the choice of reform regions to show that the decision makers were not concerned about the labour markets of the reform and non-reform regions and the choice was made for other reasons. Mentions of the labour market were virtually absent from the official documents related to the reform. I thus argue that the regional variation in home-ownership rate created by the reform can be used to identify the causal effects. Because there are externalities caused by the explanatory variable, the use of cross-regional variation in addition to individual-level variation is essential to capture the impacts correctly.

1.3.3 Welfare state and life satisfaction: Evidence from public health care

In the third essay, which is joint work with Kaisa Kotakorpi, we examine the link between the welfare state and citizens' life satisfaction by using evidence from publicly provided health care services. By combining local level data on public health care, and individual level data on life satisfaction, we show that relatively high expenditures in health care have a positive effect on individuals' life satisfaction in our data. We find some evidence for an 'ends-against-the-middle' equilibrium (Epple and Romano, 1996) in the provision of public health care, where middle-income individuals prefer higher public expenditure than low-income or high-income individuals. Further, our results indicate that valuation for health care depends on individual political orientation.

The value of many of the services typically associated with the welfare state is difficult to measure, as many of these services have public goods characteristics. Further, even if some publicly provided services are essentially private in nature (such as in the case of health care and education), individuals' valuation for them cannot in most cases be observed directly as they are often offered free of charge or at heavily subsidised prices by the government. Studying the effects of the services on subjective well-being

of citizens is one of the ways to tackle the evaluation problem (for discussion on the alternative methods, see Frey and Stutzer, 2005).

Finnish municipalities are responsible for providing health care services for their residents but have a high degree of autonomy over the specific aspects of care. Thus, there is regional variation in health care expenditures on top of the variation caused by the need for the services. Since higher need for services causes costs to rise and is likely to be negatively related to subjective well-being, we construct an expenditure measure more closely related to the level of service provision or service quality. The relevant policy instrument from the point of view of the local government is expenditures conditional on service volume. We estimate the amount of excess resources spent on public health care, given the number of treatment days and visits to public health centres and hospitals. Given the fact that the variation in excess spending is not due to any reforms or any other exogenous factor, we need to discuss the cause for the variation and its possible relationship with subjective well-being.[11] We argue that the variation in health expenditures reflects municipal decision making which, in turn, is dependent on the political opinions and voting behaviour of the inhabitants. To avoid omitted variable bias due to correlation between political opinions and subjective well-being, we conduct robustness checks by including various control variables that reflect the political opinions of our sample individuals and variables measuring the municipal politics. Thus, our strategy is to control for the possibly endogenous component in health care spending. This strategy is based on knowledge of the mechanism that determines public spending.

1.3.4 Worker turnover, structural change and interregional migration: Evidence from Finland

In the fourth essay, I study the effects of regional labour market conditions on interregional migration using province-level panel data on bilateral migration flows and disaggregated labour market flows. The results indicate that hires from unemployment and job separations leading to unemployment have sizeable effects on migration. The effects of hires from and separations to other labour market states, while statistically significant, appear smaller. Further, the results suggest that inter-industry and inter-firm shifts in employment are immaterial for migration. Taken together, interregional

---

11   In principle, with data on individual service use, it would be possible to construct a measure of health care spending that varies across individuals within municipalities. Such data, however, would only be available to those who actually have used the services. Using municipality-level aggregates can also be seen as a way to exclude endogenous individual level variation from the explanatory variable. Notice that we do not exclude the possibility of externalities in this paper: High quality health care may improve well-being of service users' relatives and friends as well. These effects would be captured by our estimate as long as the relatives and friends live in the same municipality.

migration is largely affected by regional differences in unemployment and the employment opportunities available for unemployed workers.

It has been long hypothesised and documented by various studies that interregional migration is linked to labour market conditions of regions. Many studies use the unemployment rate as the key regressor to characterise the regional labour market conditions. Some of the studies have failed to find statistically significant effects for this variable (e.g. Furceri, 2006; Hatton and Tani, 2005; McCormick and Wahba, 2005; Parikh and Leuvensteijn, 2003). Fields, in his 1976 and 1979 papers, argues that the unemployment rate may not be sufficient indicator for the regional labour market conditions that are relevant for potential (in- or out-) migrants. Some recent studies, following the argumentation, have included alternative indicators of the labour market conditions in their migration equation (e.g. Carlsen *et al.*, 2006; Hämäläinen and Böckerman, 2004). I contribute to this literature by examining the effects of labour market indicators based on linked employer-employee data that previously were not used in the interregional migration literature. More specifically, I explore the roles of worker turnover (hires and separations) and change in the structure of employment as well as 'churning' as determinants of migration. The results are mostly in line with earlier micro-level and regional-level studies and include new results on the role of labour market conditions on migration.

The key variables in the data include regional-level information on gross migration flows and labour market flows. The data is a region-level panel and, thus, panel data methods are used. Since the possibility of autocorrelation in migration cannot be ruled out *a priori*, I use a dynamic panel GMM method in estimations. Lagged differences of the key explanatory variables are used as the instumental variables. In addition to these instruments, two 'genuine' instruments (investments and exports) are used.

References

Abbring, J. H. and Heckman, J. J. (2007). Econometric Evaluation of Social Programs, Part III: Distributional Treatment Effects, Dynamic Treatment Effects, Dynamic Discrete Choice, and General Equilibrium Policy Evaluation. In: J.J. Heckman and E.E. Leamer (eds.), *Handbook of Econometrics*, Edition 1, Vol. 6B. Elsevier.

Acemoglu, D., and J. Angrist (2000). How Large Are the Social Returns to Education: Evidence from Compulsory Schooling Laws. In *NBER Macroeconomics Annual 2000*, 9–74.

Angrist, J. D. and Pischke, J.-S. (2008). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Arellano, M. and Bond, S. (1991) Some tests of specifications for panel data: Monte Carlo evidence and an application to employment equations, *Review of Economic Studies*, **58**, 227–297.

Arellano, M. and Bover, O. (1995) Another look at the instrumental-variable estimation of error-component models, *Journal of Econometrics*, **68**, 29–52.

Bertrand, M., Duflo, E. and Mullainathan, S. (2004). How Much Should We Trust Differences-in-Differences Estimates?. *The Quarterly Journal of Economics*, 119(1), 249–275.

Blanchflower, D. G. and Oswald, A. (2013). *Does High Home-Ownership Impair the Labor Market?* NBER Working Paper No. 19079.

Blundell, R. and Bond, S. (1998) Initial conditions and moment restrictions in dynamic panel data models, *Journal of Econometrics*, **87**, 115–143.

Blundell, R., Costa Dias, M., Meghir, C. and Van Reenen, J. (2004). Evaluating the Employment Impact of a Mandatory Job Search Program. *Journal of the European Economic Association*, 2(4), 569–606.

Blundell, R., Reed, H. and Stoker, T. M. (2003). Interpreting Aggregate Wage Growth: The Role of Labor Market Participation. *The American Economic Review*, 93(4), 1114–1131.

Carlsen, F., Johansen, K. and Røed K. (2006) Wage formation, regional migration and local labour market tightness, *Oxford Bulletin of Economics and Statistics*, **68**, 423–444.

Chetty, R. (2011). *Salience and Taxation: Evidence and Policy Implications*. Testimony for the U.S. Senate Finance Committee, March 2011.

Coulson, N. E. and Fisher, L. M. (2009). Housing Tenure and Labor Market Impacts: The Search Goes on. *Journal of Urban Economics*, 65, 252–64.

Crépon, B., Duflo, E., Gurgand, M., Rathelot, R. and Zamora, P. (2013). Do Labor Market Policies have Displacement Effects? Evidence from a Clustered Randomized Experiment. *Quarterly Journal of Economics*, 128(2), 531–580.

Dahlberg, M., Mörk, E., Rattsø, J. and Ågren, H. (2008). Using a Discontinuous Grant Rule to Identify the Effect of Grants on Local Taxes and Spending. *Journal of Public Economics*, 92(12), 2320–2335.

Di Tella, R., MacCulloch, R. J. and Oswald, A. J. (2003). The Macroeconomics of Happiness. *Review of Economics and Statistics*, 85(4), 809–827.

Fields, G. S. (1976) Labor force migration, unemployment and job turnover, *The Review of Economics and Statistics*, **58**, 407–415.

Fields, G. S. (1979) Place-to-place Migration. Some new evidence, *The Review of Economics and Statistics*, **61**, 21–32.

Flatau, P., Forbes, M., Hendershott, P. H. and Wood, G. A. (2003). *Homeownership and Unemployment: The Roles of Leverage and Public Housing*. NBER Working Paper No. 10021.

Furceri, D. (2006) Does labour respond to cyclical fluctuations? The case of Italy, *Applied Economics Letters*, **13**, 135–139.

Frey, B. S. and Stutzer, A. (2005). Happiness Research: State and Prospects. *Review of Social Economy*, 62, 207–228.

Gordon, N. (2004). Do Federal Grants Boost School Spending? Evidence from Title I. *Journal of Public Economics*, 88(9), 1771–1792.

Hatton, T. and Tani, M. (2005) International migration and inter-regional mobility in the UK: 1982–2000, *The Economic Journal*, **155**, 342–358.

Heckman, J. J., Lochner, L. and Taber, C. (1999). Human Capital Formation and General Equilibrium Treatment Effects: A Study of Tax and Tuition Policy. *Fiscal Studies*, 20(1), 25–40.

Hämäläinen, K. and Böckerman, P. (2004) Regional labor market dynamics, housing and migration, *Journal of Regional Science*, **44**, 543–568.

Imbens, G. W., and Wooldridge, J. M. (2009). Recent Developments in the Econometrics of Program Evaluation. *Journal of Economic Literature*, 47(1), 5–86.

Kleven, H. J., Knudsen, M. B., Kreiner, C. T., Pedersen, S. and Saez, E. (2011). Unwilling or Unable to Cheat? Evidence from a Tax Audit Experiment in Denmark. *Econometrica*, 79(3), 651–692.

Kotakorpi, K. and Laamanen, J.-P. (2010). Welfare State and Life Satisfaction: Evidence from Public Health Care. *Economica*, 77, 565–583.

Van Leuvensteijn, M. and Koning, P. (2004). The Effects of Home-ownership on Labour Mobility in the Netherlands. *Journal of Urban Economics*, 55, 580–596.

McCormick, B. and Wahba, J. (2005) Why do the young and educated in LDCs concentrate in large cities? Evidence from migration data, *Economica*, **72**, 39–67.

Mian, A. R. and Sufi, A. (2012). *What Explains High Unemployment? The Aggregate Demand Channel*. Chicago Booth Research Paper No. 13–43.

Moulton, B. R. (1986). Random Group Effects and the Precision of Regression Estimates. *Journal of Econometrics*, 32, 385–397.

Munch, J. R., Rosholm, M. and Svarer, M. (2006). Are Homeowners Really More Unemployed? *The Economic Journal*, 116, 991–1013.

Munch, J. R., Rosholm, M. and Svarer, M. (2008). Home Ownership, Job Duration, and Wages. *Journal of Urban Economics*, 63, 130–145.

Nickell, S. (1981). Biases in Dynamic Models with Fixed Effects. *Econometrica*, 49(6), 1417–1426.

Oswald, A. J. (1996). A Conjecture on the Explanation for High Unemployment in the Industrialized Nations: Part I. Working Paper No. 475, University of Warwick.

Oswald, A. J. (1999). The Housing Market and Europe's Unemployment: A Non-technical Paper. University of Warwick.

Parikh, A. and van Leuvensteijn, M. (2003) Interregional labour mobility, inequality and wage convergence, *Applied Economics*, **35**, 931–941.

Stevenson, B. and Wolfers, J. (2008). Economic Growth and Subjective Well-Being: Reassessing the Easterlin Paradox. *Brookings Papers on Economic Activity*, 39(1), 1–102.

Chapter 2
The Effects of Pre-populated Tax Return on Deductions: Evidence from a Natural Experiment[*]

Kaisa Kotakorpi[a] and Jani-Petri Laamanen[b†]
[a] University of Turku and CESifo
[b] University of Tampere

Abstract

We estimate the effect of the details of the income tax filing system on the behaviour of taxpayers, utilising data from a policy experiment that took place in Finland in the 1990s. We find that receiving a pre-filled income tax return lead to a reduction of over one-fourth in the number of individuals claiming for deductions. We discuss the possible theoretical interpretations of our results that are related to complexity and salience effects, opportunity cost of time as well as tax evasion.

Keywords: income tax filing, complexity, salience, natural experiment
JEL: H24, H31

[†] E-mail addresses: Kaisa.Kotakorpi@utu.fi; Jani-Petri.Laamanen@uta.fi

2.1 Introduction

Recent evidence suggests that factors beyond actual tax rates and benefit levels affect how individuals react to the tax and benefit system. An important insight is that the particular design features of the tax and benefit system – affecting for example the complexity of the system and the visibility of the incentives involved – affect individual behaviour. For example, empirical evidence shows that program complexity affects the take-up of social benefits (see e.g. Currie 2006). In the context of tax systems, recent evidence on tax salience indicates that the way in which taxes are collected or represented affect consumer behaviour (see e.g. Chetty, Looney and Kroft 2009, Finkelstein 2009). Further, the evidence presented in Saez (2010) indicates that taxpayers may fail to understand the features of a complex income tax system, while Shapiro and Slemrod (1995) showed that a change in the income tax withholding system changed consumption behaviour, even though there was no change in the tax rate. Kleven *et al.* (2011) stress the importance of third-party information (in a system using pre-filled tax returns) for tax compliance and self-reporting behaviour.

The purpose of this paper is to examine how the income tax filing system affects taxpayers' filing behaviour. We utilise a policy experiment that took place in Finland in the mid-1990s, whereby a proportion of taxpayers received a pre-filled income tax return, whereas other taxpayers had to file a full return (the normal practice). For those who received a pre-filled return, items such as labour income and benefits received had been pre-filled based on information received from third parties such as employers and benefits agencies. These individuals were required to file a final return only if some income information was incorrect or missing and had the *option* of filing a return e.g. in case they wanted to file for some discretionary deductions. The experiment was gradually expanded over the years, and currently most Finnish taxpayers receive a pre-filled income tax return. Due to the gradual expansion of the experiment and the nature of the rules that determined inclusion of individuals into the experiment, we will be able to estimate the causal effect of the income tax filing system on taxpayers' filing behaviour.

Beyond evaluating the effects of this particular reform, the evidence that we find is likely to be of interest for policy in other countries. For example, a similar system has been proposed in the U.S.[1], where the proposed shift to using pre-filled income tax returns has been primarily justified by savings in tax preparation costs (Goolsbee 2006).[2] In addition to such direct cost savings, it is interesting and important to know

---

[1]   See http://obama.3cdn.net/8f478c5e1bb07ca0b1_sh1umv2zy.pdf (Barack Obama's Economic Agenda '08), accessed on 7 Dec 2012.

[2]   Prefilled returns are used also at least and to varying degrees in Sweden, Denmark, Norway Iceland, Estonia, Chile and Spain (OECD 2006). A similar system has been experimented with in California (Bankman 2005).

whether and how taxpayers' filing behaviour would be affected. Filing behaviour will affect tax revenues directly. Further, if a simpler return affects perceived tax complexity and tax salience, it may affect how individuals view the incentive effects of labour taxation and hence even labour supply. The paper therefore also contributes to the literature on how individual behaviour is affected by tax complexity and salience. Our findings indicate that receiving a pre-filled return significantly reduced filing for discretionary deductions such as travel expenses and other costs of acquiring income.

The paper is organised as follows. In Section 2.2, we discuss some theoretical considerations related to why and how the details of the income tax filing system might affect filing behaviour. In Section 2.3, we describe the Finnish experiment involving pre-filled income tax returns. The data and methodology are described in Section 2.4. The results are presented in Section 2.5. Section 2.6 concludes.

2.2   Theoretical considerations

Why would the details of the income tax filing system matter for reporting behaviour? One can think of at least three broad lines of argument. First, there are very standard economics arguments related to the cost of reporting. A second set of explanations is derived from findings in behavioural economics, and is mainly related to tax salience and the complexity of the tax system. Third, the opportunities and incentives for tax evasion (willful underreporting) may be affected by the details of the tax filing system.

Let us first consider conventional economics arguments related to the costs of reporting. The main component here is likely to be the opportunity cost of time spent on activities related to reporting, in particular filling out forms.[3] The opportunity cost of time is likely to be convex: it is fairly easy to find a small amount of time for any particular activity, but since there is only a fixed number of hours in a day, additional time units become increasingly costly. Many items were pre-filled for those individuals who received a pre-populated tax return. Therefore, considering possible effects on reporting any discretionary items such as deductions, a convex opportunity cost of time would point towards individuals who received a pre-populated return filing for more such items, as compared to individuals who had to file a complete return.

Second, let us turn to arguments related to findings from behavioural economics. From a standard economics point of view, individuals should be fully attentive to the details of the tax system, and should be able to carry out the complex calculations that

---

3   There were no significant differences in other physical costs incurred by individuals subject to the two different systems of tax filing. Making adjustments to the pre-filled tax return was in practice identical to reporting similar items in the old system. Also other rules e.g. regarding retaining receipts related to deduction claims were identical across the two systems. Regardless of the system that was applied to a given individual, a pre-paid envelope for posting the return or the corrected pre-filled form to the tax authorities was provided.

may be necessary for correct reporting. Details of the tax system should be of minor importance. However, recent literature shows that particular design features of the tax system do affect behaviour in important ways (for a recent review, see Congdon, Kling and Mullainathan 2011). Two notions appear to be particularly relevant here: salience and complexity. Salience refers to the degree of visibility of taxation (see e.g. Chetty, Looney and Kroft 2009). In a similar vein as previous literature on tax salience is concerned with the visibility of tax rates (or tax inclusive prices), we refer to the visibility of other parts of the tax code.[4] In our case, some individuals were required by law to fill out a tax return, whereas those individuals who received a pre-filled tax return in most cases did not have to take any action to comply with their legal duties as taxpayers. This may have made income taxation less visible for them, and they may simply have forgotten to file for any discretionary items such as deductions.

The argument related to complexity goes as follows: The tax system is highly complex, and there may be significant fixed costs related to learning the details of the system in order to report one's income and other items correctly for tax purposes. Those who are required to file a complete income tax return have no choice but to incur these fixed costs. On the other hand, those who receive a pre-filled return, may choose to avoid these costs. If the expected gain from filing any discretionary items is smaller than the fixed cost, the individual will not file for any discretionary items.[5]

Taking together the two lines of argument related to (conventional) costs of reporting and to complexity costs, a taxfiler's cost function (as a function of the number of items reported) may first be concave and then convex: there may be a fixed cost related to learning the details of a complex tax system, but thereafter the function may be convex due to convex opportunity costs of time. Therefore the effect of receiving a pre-filled tax return on reporting of any discretionary items such as deductions, based on these arguments, is unclear a priori.

A third line of argument is related to the opportunities and incentives for tax evasion. According to the deterrence theory of tax evasion (Allingham and Sandmo, 1972), the perceived probability of detection is a key determinant of the extent of tax evasion. Kleven *et al.* (2011) and Slemrod (2007) stress the importance of third-party reporting for the extent of tax evasion: as the authorities receive information on certain income items also from third parties (e.g. employers), the probability of detection is very high and the opportunities for evasion correspondingly very low on these items. In our case, many important items (e.g. regular wages) were indeed subject to third-party reporting in both systems (regardless of whether the individual received a pre-filled return or filed

---

4   Whereas the previous literature is mostly concerned with the effects of salience on real behaviour such as consumption decisions, we are in the current paper interested in the effects of salience on reporting behaviour.

5   The fixed cost may also be related to the negative psychological emotions associated with having to think about taxation or complex tasks in general. It is however hard to think of any "conventional" type of fixed costs associated with filing a tax return.

Jani-Petri Laamanen

a complete return). For these items, computerised matching of information is possible and the actual probability that evasion is detected is high, and does not depend on whether the third party information is pre-filled or not. Nevertheless, the *perceived* probability of detection may be higher in the case of receiving a pre-filled tax form: when receiving an empty form, individuals may not be fully aware what information is available to the authorities from third parties. Hence underreporting may be lower for those receiving a pre-filled form. This argument is however likely to be most relevant for items such as wage income that are printed on the pre-filled form (and where third party information actually exists). For other items such as discretionary deductions, for this argument to play a role one would need to rely on the existence of some type of spillover effects and even misperception: the fact that the tax authorities have some third party information on the taxpayer's income is brought to the individual's attention, and she may wonder whether the tax authorities actually know more than is printed on the form.

Putting together the three lines of argument outlined above, one may think of a broad notion of tax compliance, defined in terms of the accuracy of reporting and taking into account both under- and over-reporting. Such a notion would then encompass all the three mechanisms discussed above: First, costs of reporting may naturally reduce the accuracy of reporting (see e.g. Pitt and Slemrod (1989) on the compliance cost of itemising deductions). Second, behavioural explanations point to the possibility of honest mistakes and omissions in reporting. Third, tax evasion refers to a particular type of non-compliance, namely the case of willful underreporting.

2.3   The experiment

The paper utilises data from the so-called Finnish tax proposal experiment implemented between 1995 and 2004. During the experiment, a proportion of non-entrepreneur taxpayers did not have to file a return, but instead received a pre-filled return ("tax proposal") where their tax liability had been calculated based on third-party reports (e.g. wage income reported by employers and transfers reported by benefit agencies) and other relevant information available to the government. These taxpayers were required to file a final return only if some income information was incorrect or missing (which usually was not the case), and had the option of making adjustments and filing a final return for example if they wished to claim for some discretionary deductions; if they did not do so, the tax authorities' original proposal was implemented.

The main objective of moving to a system with pre-filled income tax returns was to save on tax filing expenses, even though it was also mentioned that the accuracy of reporting should not be compromised. A key piece of information for our purposes are the rules that determined a taxpayer's inclusion into the experiment. There were differences in inclusion criteria between groups of municipalities: different

municipalities used different selection models, whereby both the number and types of individuals included into the experiment varied across municipalities. Within each model, individuals were selected according to variables such as age, income, lagged deductions and wealth holdings. It is noteworthy that these criteria were not known to taxpayers at the time, and therefore self-selection into the experiment is not an issue. The rules created variation in treatment status across individuals, depending on where the individual lived.

The experiment started in 1995 with approximately 300,000 individuals (7% of taxpayers) receiving a pre-filled return. In the first year of the experiment, there were four different selection models used: the most selective model was one where only pensioners and young individuals with low earned income in a given municipality were included in the experiment; and the most inclusive model involved all regular taxpayers in a given municipality. The other two selection models were between these two extremes in the share of individuals selected: the models involved regular wage earners in a given municipality, and differed in the exact details of the selection rules. One of the above four models was implemented in 9.8% out of the 439 municipalities and therefore in the first year, most municipalities did not participate in the experiment at all[6].

In subsequent years, the number of participating individuals increased each year, through two channels: The experiment was extended into new municipalities, and the selection models were changed (usually) to include more taxpayers in a given municipality. The number of individuals who received a pre-filled return each year is reported in Table 1. Today most regular, non-entrepreneur taxpayers receive a pre-filled return.

Table 1. Number of taxpayers receiving a pre-filled return.

Year	Pre-filled return	Normal retur	Total	% pre-filled
1995	303,981	3,819,187	4,123,168	7.4
1996	1,325,855	2,813,613	4,139,468	32.0
1997	2,367,027	1,797,996	4,165,023	56.8
1998	2,299,431	1,884,912	4,184,343	55.0
1999	2,850,301	1,359,241	4,209,542	67.7
2000	2,864,301	1,369,787	4,234,088	67.6
2001	2,875,631	1,378,058	4,253,689	67.6
2002	2,920,112	1,359,233	4,279,345	68.2
2003	2,947,364	1,352,664	4,300,028	68.5
2004	2,950,826	1,365,848	4,316,674	68.4

---

6   Our data exclude the 16 municipalities in the autonomous island of Åland.

Let us next discuss how decisions were made regarding which selection model to use at the municipal level. At the time, Finland was divided into 10 tax office areas, which were in turn divided into a total of 93 tax districts. There were altogether 455 municipalities in 1995, and each municipality belonged to one tax district. When the experiment was launched, each tax office chose one or sometimes two default selection models that were applied to the tax districts within that tax office area. Before making a final decision on the models to be applied to each district and municipality, the tax offices were able to apply the different models to data from previous years' tax collection, to see for example how many taxpayers would be selected from each municipality under any given model. Some municipalities and districts were then subject to exceptions, meaning that some other model than the default model of their tax office area was finally applied to that municipality/district. The main reason for such exceptions was to include a sufficient number of taxpayers into the experiment. Therefore, while not completely random, the selection of municipalities into the experiment appears unrelated to our outcomes of interest (e.g. trends in deductions claims), and therefore does not interfere with our identification strategy. Identification will be discussed in detail on Section 2.4.2.

2.4   Data and methodology

2.4.1   Data

Our analysis is based on two data sources. First, we have obtained detailed documentation on the experiment from the tax authorities. Notably, this includes detailed information on the criteria based on which individuals were selected into the experiment each year. Second, for carrying out the main part of the analysis, we use individual level data on a representative panel (20%) of Finnish non-entrepreneur taxpayers (and their spouses) excluding those in the top 1% of the income distribution from Statistics Finland's tax register for 1993–1997. Altogether, there are approximately 1 million individuals in the data, and we follow them over the entire period. The data cover all variables contained in the tax register: there is information on various forms of income as well as tax payments (for example labour income, capital income, transfers, asset holdings, debt, taxes paid, very detailed information on various types of deductions), as well as background information (age, gender, marital status, place of residence, level of education).

We anticipate that the experiment may have affected the amount of deductions that individuals filed for.[7] Discretionary decutions provide a key example of pure self-reporting. On the other hand, regardless of whether individuals received a pre-filled return or not, information on individual incomes was filed also by employers and/or benefits agencies, so there was probably less scope for the experiment to have an effect on reported income. We therefore concentrate on deductions, and examine whether the experiment affected the number of individuals who filed for any discretionary deductions, as well as the amount filed for. Discretionary deductions include the costs of travel to work, interest payments for mortgages and student loans, costs for acquiring income and other smaller items.

In what follows, we look at the effect on deductions claims of receiving a pre-filled return for the first time[8]. We utilise data from the first three years of the experiment (1995–7), when the experiment expanded considerably from covering 7% of regular taxpayers in 1995 to covering 57% of taxpayers in 1997.[9]

2.4.2 Methodology

To examine the effects of receiving a pre-filled tax form, we estimate models of the following type:

$$y_{it} = \alpha_i + \gamma_t + D_{it}\delta + \beta' X_{it} + \epsilon_{it}, \tag{1}$$

where y_{it} is the outcome variable of interest for individual i in year t. D is a dummy variable indicating if an individual received a pre-filled tax form and δ is the effect on the outcome variable. We control for year dummies γ_t as well as for individual fixed effects α_i (through estimating the equation in first differences) and time-varying individual characteristics X_{it}.[10] ε is the error term.

---

7   In the Finnish system, there are a number of deductions / allowances that taxpayers receive automatically, as well as a number of discretionary deductions that they can apply for. The latter is similar to itemisation in the U.S. case. The difference is that the taxpayers cannot choose between a basic allowance and itemisation. Rather, all taxpayers receive a basic allowance, and they can choose whether to file for additional deductions on top of it. Itemisation behaviour in the U.S. has been studied by Pitt and Slemrod (1989) and Slemrod (1989).

8   We only include individuals who are present in the panel for the whole period of investigation (1995–1997). This means also excluding individuals who were in the top 1% of the income distribution in any of the years. Altogether, some 10% of the sample individuals are dropped due to entry, exit or having high income during the period.

9   Analysing later years when more and more individuals are treated becomes problematic, since it is increasingly difficult to find suitable non-treated controls for the treated individuals.

10   We control for education (six categories), number of children (a dummy for each number), marital status (four categories) and province of residence (20 provinces) in all regressions.

Given the nature of the rules that determined whether an individual will receive a pre-filled income tax return or not, and the fact that these rules varied both across individuals according to their place of recidence, as well as over time, we are able to identify the causal effect of the method of income tax reporting on taxpayers' behaviour. Information on the selection rules is used to form an appropriate control group for treated individuals. As was explained above, individuals were selected into the experiment according to e.g. their income histories and other characteristics. We require that individuals in the control group have these same characteristics; the reason why the control individuals were not treated was that they lived in a different municipality that applied some other selection model.[11]

We feel that the reform provides an unusually attractive set-up, since selection of individuals into the experiment was based only on observables; and since we have information on those observables, we can select the control group by those very same criteria. The identifying assumption is that the year-to-year changes in the outcome variable for the treatment and control groups would have been similar in the absence of the treatment. As usual in a difference-in-difference set-up, we can provide support for this assumption through examining the pre-treatment trends in the outcome variable in the two groups. In the analysis, we also exploit the panel nature of the data and control for individual fixed effects.

In what follows, the treatment group consists of individuals who received the pre-filled return for the first time in 1995, 1996 or 1997, and we pool the treatment and control groups for the different years. Figure 1 shows trends in the proportion of individuals claiming for discretionary deductions in the treatment and control groups. To be able to pool the analysis for the different years, the x-axis measures years until the individual entered the experiment. That is, year 0 in the figure is the year in which the individuals in the treatment group received a pre-filled income tax return for the first time.[12]

Figure 1 shows that before the experiment, the proportions of individuals claiming for discretionary deductions in the treatment and control groups follow each other

---

11   More details of how we have constructed the control group are given in the Appendix. Note that there are in general many individuals who could serve as a control for any treated individual (i.e. who fulfil the selection criteria in the way explained in the Appendix). We select one control individual for each treated individual randomly from this pool of eligible controls.

12   Notice that since our data begin in the year 1993, period -4 only includes individuals who were treated in 1997 (and their controls). Period -3 includes individuals treated in 1996 or 1997 (and their controls) and the later periods include all treated individuals and their controls.
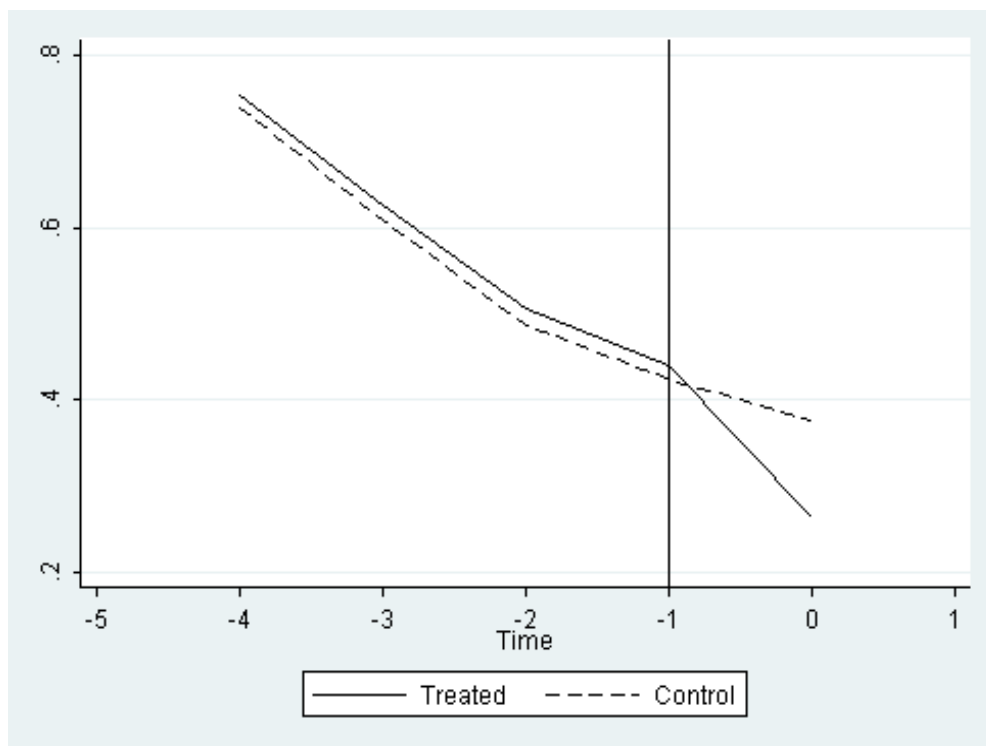
Figure 1: Trends in proportions of individuals claiming for deductions.

very closely. There is a downward trend in both groups before the experiment.[13] In the year when the treated individuals received a pre-filled tax return, there appears to be a marked drop in the proportion of individuals claiming for deductions in the treatment group, and the trends in the proportions of individuals claiming for deductions in the two groups diverge. This graphical analysis provides some preliminary evidence that receiving a pre-filled income tax return may have had an effect on deductions claims.[14] We next turn to an econometric analysis of this effect.

---

13   There are three possible reasons for this downward trend. First, the fraction of individuals claiming for deductions may change over time for some very natural reasons e.g. because of changes in economic conditions; that is, there is no reason to assume that this fraction should remain constant over time. (Indeed, this is the primary reason why we need a control group in the analysis.) Second, in the first three periods, there are changes in the composition of our sample (both treatment and control groups): Later years of the experiment included individuals with a larger average propensity to claim deductions, and these are the individuals who show up in the first years in Figure 1, as explained in footnote 10. Third, there were some changes in the menu of self-reported deductions during the sample period. Most of these had only small impacts, but in 1997, interest payments for loans were pre-filled for the first time. Excluding these items from our dependent variable causes a decline in self-reported deductions.

14   When the graphical analysis is conducted separately for individuals treated in different years (and their control individuals), the conclusions do not change.

Jani-Petri Laamanen

2.5 Results

Let us first turn to results regarding the fraction of individuals who claimed for discretionary deductions: The dependent variable is a dummy for whether the individual claimed for any discretionary deductions.[15] The results are reported in column (1) of Table 2. The results indicate that receiving a pre-filled return reduced the fraction of individuals claiming for discretionary deductions by 12.7%-points. This is a fairly large effect: in the year preceding their inclusion into the experiment, approximately 44% of the individuals in our treatment group filed for deductions. Receiving a pre-filled return therefore reduced the number of individuals claiming for deductions by over 1/4.

Table 2. Effect of receiving a pre-filled tax return on deduction claims.

	(1) Dep. var.: Deduction-dummy	(2) Dep. var.: Deduction amount (FIM)
Pre-filled return	-0.1274** (0.0092)	-659.01** (213.94)
N	370,748	370,748
Individual controls and fixed effects	Yes	Yes

Robust and clustered (municipality) standard errors in parentheses.
** p< 0.01, * p<0.05

To provide support for our identifying assumption, we have also examined a corresponding pseudo-treatment, where we pretend that the individuals in the treatment group received a pre-filled return one year earlier than they actually did. The estimated effect of the pseudo-treatment is an insignificant -0.4%-points (standard error 1.0%-points), which provides support for our set-up.

Let us next analyse how receiving a pre-filled tax return affected the amount of discretionary deductions that individuals filed for on average. The dependent variable is now the amount of discretionary deductions filed for, that is, a total sum of money associated with all relevant items on the tax return. The set-up is otherwise identical to the above analysis. The results are presented in column (2) of Table 2. The result shows that the amount of deductions filed for decreased on average by about FIM 660, which corresponds to about € 194 in 2013 (inflated by the average earnings index). This reduction reflects the fact that the fraction of zeros increased, as was shown above.

The average amount of deductions (conditional on filing a positive amount) increased by an extra FIM 872 in the treatment group compared to the recipients of

---

15   That is, this dummy equals 0 if *all* items associated with discretionary deductions in the individuals tax record were zero; and 1 otherwise.

the conventional return. This figure includes both the change in the composition of the group which claimed deductions and the changes in the claimed amounts. In Figure 2, we take a closer look at the changes in the distribution of deduction claims. We present the whole distribution of deductions (excluding zeros and the top 1% to make the changes elsewhere visible) before and after the reform, for the treatment and control groups separately. In the treatment group there was a significant drop in frequencies at low deductions. In the control group, on the other hand, there was a much smaller or no drop. For example, the frequency of the lowest bin (FIM 1-800) dropped by about 6,000 to slightly over 4,000 individuals in the treatment group. The control group, where the frequency was similar to the treatment group before the reform, ended up having around 9,000 individuals in this bin. Formal analyses confirm the visual observation that the reform caused a significant drop in frequencies at low deductions, but that there was no such drop in frequencies at high deductions. We tried estimating models in which the dependent variable was a dummy indicating deductions larger than some amount K. It appears that the p-value of the treatment effect falls below 0.05 as K is increased to approximately FIM 12,000. The magnitude of the effect falls
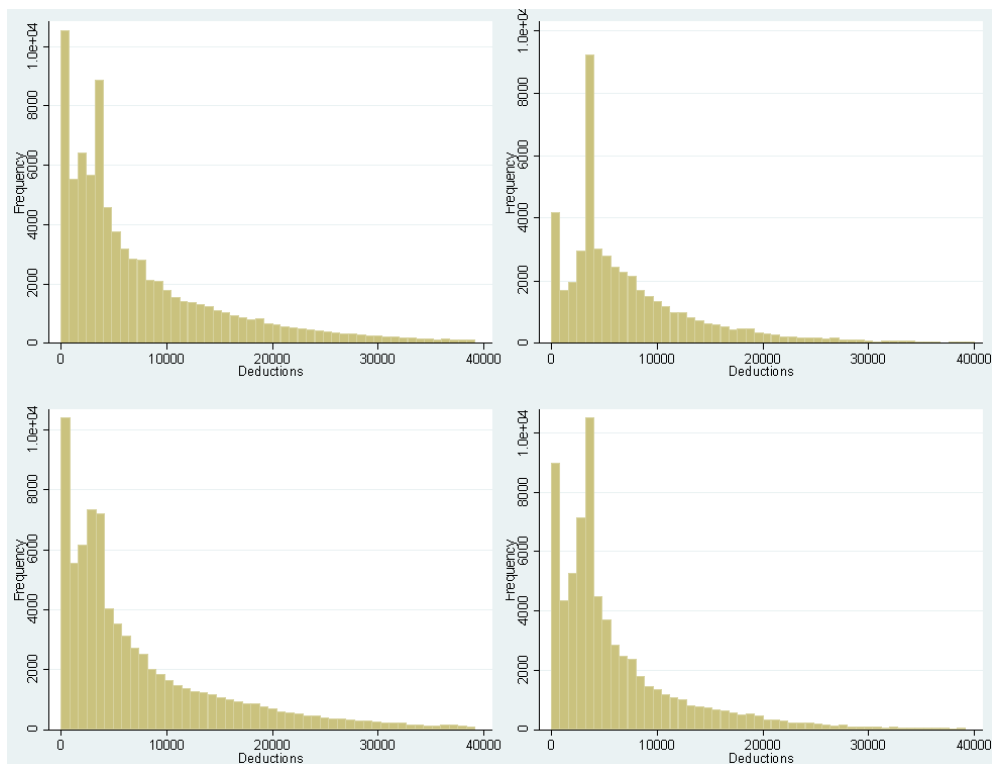

Figure 2: Histograms of deductions claims before and after the reform, in treatment group (top panel) and control group (bottom panel)

Jani-Petri Laamanen

to about 0.8%-points. The disappearance of predominantly low deductions from the treatment group thus increased the average deductions (conditional on being positive) in this group.

We have also examined the effect of receiving a pre-filled income tax return on the probability of filing for deductions separately for different income and educational groups (see Table 3). First, the sample is divided into four groups according to the annual income quartiles. We define the quartiles separately for each year and an individual's quartile is determined by her previous year income level. We obtain statistically significant effects for all income groups. Only the effect on low-income individuals is (slightly) outside the 95% confidence level of the average effect. We also estimated the model separately for the highest and the lowest income decile. The results of this exercise confirmed the conclusion that the effect of the reform is not very heterogeneous with respect to income. Estimating the effect separately for educational groups also points towards a fairly homogeneous effect. The effects on those with the lowest level of education is slightly below the confidence interval of the average effect and the effect on those in the second educational group, in turn, is just slightly above.

However, when interpreting the above results, it should be noted that individuals with high income or education have a much higher propensity to claim for deductions to start with. Therefore, even though the absolute effects reported above are fairly homogeneous across groups, the relative effects are much higher for individuals with low income or low education. For example, in the lowest income quartile, 21% of individuals claimed for deductions to start with, and therefore the estimated 11%-point reduction implies that receiving a pre-filled income tax return caused more than 1/2 of those individuals not to file for deductions. For individuals in the highest income quartile, on the other hand, the fraction of individuals claiming for deductions dropped 13%-points from a starting level of 63%, amounting to a drop of 1/5 caused by receiving a pre-filled return. These relative effects are reported for each income quartile and education group on the last row of the corresponding panel of Table 3, which shows that the relative effect is decreasing in both income and the level of education.

Table 3. Effect of receiving a pre-filled tax return on deduction claims

	Income quartile			
	(1)	**(2)**	**(3)**	**(4)**
Pre-filled return	-0.1065**	-0.1321**	-0.1242**	-0.1308**
	(0.0082)	(0.0077)	(0.0109)	(0.0151)
N	92,671	92,701	92,687	92,689
Individual controls and fixed effects	Yes	Yes	Yes	Yes
Relative effect	-50%	-33%	-25%	-21%
	Education group			
	(1)	**(2)**	**(3)**	**(4)**
Pre-filled return	-0.0911**	-0.1558**	-0.1391**	-0.1403**
	(0.0092)	(0.0100)	(0.0173)	(0.0197)
N	153,555	142,642	44,153	30,398
Individual controls and fixed effects	Yes	Yes	Yes	Yes
Relative effect	-33%	-33%	-21%	-18%

Robust and clustered (municipality) standard errors in parentheses.

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

Let us next turn back to the theoretical arguments on why the details of the tax filing system might affect reporting behaviour, and to a discussion of which of the explanations outlined in Section 2 are most likely to be the driving forces behind our results. In sum, we find a large reducing effect of the pre-filled return on the proportion of individuals claiming for deductions. Hence our result cannot be explained by arguments related to the opportunity cost of time of filing taxes and deductions: as we argued above, the opportunity cost of time is likely to be convex, and would therefore point towards individuals receiving a pre-filled tax return claiming for more, not less, discretionary deductions than individuals filing a complete return.

Turning next to the explanations related to complexity and salience, an important finding is that the effect that we observe occurs predominantly through a decline in the number of individuals claiming for small deductions. This finding is consistent with individuals having a fixed complexity cost of tax filing: Individuals who received a pre-filled return and expected to be entitled to only small deductions decided not to file, as by doing so they could *choose* to avoid the fixed complexity cost of filing. Avoiding the fixed cost was not an option for individuals who had to file a complete return. On the other hand, our results appear to be less in line with the idea that receiving a pre-filled return made taxation less salient (visible). A simple visibility effect independent of the expected sum of deductions should have reduced claiming similarly across the distribution of deductions.

Finally, let us turn to the argument related to tax evasion, and compare our results to the predictions of the tax evasion model of Kleven *et al*. (2011). As was noted in

Section 2, the actual probability of detection for any given item was similar regardless of whether the individual received a pre-filled return or not. Considerations related to tax evasion may nevertheless be behind our results if the reform affected individuals' *perceived* probability of detection. For this to be the case, we would need to rely on the existence of a kind of a "spillover" effect: individuals receiving the pre-filled form realised that the tax authorities had information on e.g. their earnings, and may have become worried that the authorities may also have information on other items that were not printed on the form. If this is the mechanism behind our results, it would imply that in the new system, some individuals who would otherwise have claimed for some deductions that they were not entitled to, did not do so due to receiving a pre-filled return.

2.6   Conclusions

Receiving a pre-filled income tax return lead to a significant reduction in the propensity to claim for deductions. Since entitlement rules did not change, this finding suggests that some tax-payers failed to claim for deductions that they would have been entitled to. This effect may be due to salience effects (since filing was no longer compulsory, some individuals may simply have forgotten to claim for deductions), due to fixed complexity costs associated with income tax filing (some individuals may have rationally chosen to avoid those costs now that they were given the opportunity to do so) or due to increased perceived probability of evasion detection. The fact that it was mostly individuals with low deductions who stopped claiming tends to favour the latter two interpretations. In absolute terms, the effects of the reform appear fairly homogeneous across income and education groups. In relative terms, however, low income individuals experienced the highest drop in the fraction of individuals claiming for deductions. In this sense, the reform may seem regressive. However, this effect should be set against the lower cost of tax compliance in the new system where the duty to fill a tax return was replaced by a system involving pre-filled tax returns.

Appendix: Selecting the control group

We form the control group as follows: For each treated individual, we find a control individual by applying exactly the same criteria that caused the treated individual to be included in the experiment, but who was not treated because she lived in another municipality that either did not participate in the experiment, or used a different selection model. In selecting the control group for individuals living in, say, municipality A, from individuals living in another municipality B, we take into account the selection models used in both municipalities. The examples below clarify the procedure.

For example, assume that an individual living in municipality A, using selection model 1 (e.g. selecting pensioners only), was selected into the experiment. Consider another individual living in municipality B, who also fulfils the criteria of model 1 (i.e. is a pensioner), but was not selected into the experiment because his municipality did not participate. This individual in municipality B can now serve as a control for the treated individual in municipality A.

A more general case is the one where municipality B also participated in the experiment but used a different set of selection criteria than municipality A. In this case, when selecting a control from municipality B, we also utilise the selection criteria that were applied in municipality B: e.g. if municipality B used a model that selected low-income individuals only, we use untreated individuals in municipality B as potential controls only for high-income individuals in other municipalities.

Further, in later years of the experiment (1996→), we also require that the control individuals not only satisfy the same selection (or exclusion) criteria as the treated individual for that year, but also for the previous years of the experiment. For example, assume that an individual lives in a municipality where only pensioners were treated in 1995, and we observe that the individual is treated for the first time in 1996. We then use the selection criteria for 1996 to find a control for this individual, as explained above, but further restrict the pool of eligible controls to individuals who were non-pensioners in 1995.

References

Allingham, Michael G. and Angnar Sandmo. (1972). Income Tax Evasion: A Theoretical Analysis. *Journal of Public Economics*, 1(3–4), 323–38.

Bankman, Joseph (2005). Simple Filing for Average Citizens: The California ReadyReturn. *Tax Notes*, June 13, 2005, 1431–1439.

Chetty, Raj, Adam Looney and Kory Kroft (2009). Salience and Taxation: Theory and Evidence. *American Economic Review*, 99(4), 1145–1177.

Congdon, William J., Jeffrey R. Kling, and Sendhil Mullainathan (2011). *Policy and Choice: Public Finance through the Lens of Behavioral Economics*. Brookings Institution Press, Washington DC.

Currie, Janet (2006). The Take-up of Social Benefits. In Alan J. Auerbach, David Card, and John M. Quigley (Eds.), *Public Policy and the Income Distribution*, 80–148. Russell Sage Foundation Publications, New York.

Finkelstein, Amy (2009). EZ-Tax: Tax Salience and Tax Rates. *Quarterly Journal of Economics*, 124(3): 969–1010 .

Goolsbee, Austan (2006). The Simple Return: Reducing America's Tax Burden Through Return-Free Filing. Brookings Institution Hamilton Project Discussion Paper 2006-04.

Kleven, Henrik Jacobsen, Martin B. Knudsen, Claus Thustrup Kreiner, Sören Pedersen and Emmanuel Saez (2011). Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark. *Econometrica*, 79(3), 651–692.

OECD (2006). Using Third Party Information Reports to Assist Taxpayers Meet their Return Filing Obligations – Country Experiences with the Use of Pre-Populated Tax Returns. OECD, Paris.

Pitt, Mark M. and Joel Slemrod (1989). The Compliance Cost of Itemizing Deductions: Evidence from Individual Tax Returns. *American Economic Review*, 79(5), 1224–1232.

Saez, Emmanuel (2010). Do Taxpayers Bunch at Kink Points? *American Economic Journal: Economic Policy*, 2(3), 180–212.

Shapiro, Matthew D. and Joel Slemrod (1995). Consumer Response to the Timing of Income: Evidence from a Change in Tax Withholding. *American Economic Review*, 85(1), 274–283.

Slemrod, Joel (2007). Cheating Ourselves: The Economics of Tax Evasion. *Journal of Economic Perspectives*, 21(1), 25–48.

Slemrod, Joel (1989). The Return to Tax Simplification: An Econometric Analysis. *Public Finance Quarterly*, 17(1), 3–28.

Chapter 3
Home-ownership and the Labour Market: Evidence from Rental Housing Market Deregulation[*]

Jani-Petri Laamanen[†]

Abstract

Perhaps the most common finding relating housing to the labour market is that high home-ownership rates are associated with higher unemployment. In contrast, recent micro-evidence suggests that home-owners have relatively favourable labour market outcomes. We explore the effect of home-ownership on unemployment using a rental housing market deregulation reform which created exogenous variation in home-ownership across regions, allowing us to avoid the endogeneity problem in earlier studies. Although home-owners are less likely to experience unemployment, an increase in the home-ownership rate causes regional unemployment to rise. Externalities arising from consumption reductions and increased job competition may explain the conflicting evidence.

†   FI-33014 University of Tampere, Finland, E-mail: jani-petri.laamanen@uta.fi,
Phone: +358 50 318 5993, Fax: +358 3 2599 008

3.1   Introduction

The roles of housing markets and household credit in the economy have received increased attention after the onset of the recent economic crisis. It seems that the housing and mortgage markets may play a more important role in macroeconomic fluctuations than previously thought (see e.g. the discussion in Mian and Sufi, 2010). However, not much is known about the relationships between the housing and mortgage markets and macroeconomic outcomes. Even less is known about the mechanisms underlying these relationships. One of the earliest and most often observed relationships is the positive correlation between the rate of home-ownership and unemployment. Since Oswald's (1996) influential paper documenting this relationship, several studies have either replicated Oswald's empirical analyses with other data sets or tested the theoretical hypotheses using microeconomic data. Several studies using regional or cross-country data lend some support to the claim that a higher regional home-ownership rate leads to a higher rate of unemployment (Blanchflower and Oswald, 2013; Coulson and Fisher, 2009; Costain and Reiter, 2008; Munch *et al.*, 2006; Di Tella and MacCulloch, 2005; Green and Hendershott, 2001; Nickell, 1998). Oswald (1996) hypothesises that this is caused by lesser geographical mobility of homeowners relative to that of renters. Indeed, Battu, Ma and Phimister (2008) find that homeowners in the United Kingdom are less likely to experience a job change associated with a non-local residential move than renters. Munch *et al.* (2008) find that Danish homeowners have less local and non-local job-to-job changes than renters.

The evidence of negative mobility effects of home-ownership is in line with Oswald's (1996) hypothesis. However, several studies show that despite being less mobile, homeowners have more favourable labour market outcomes than renters. Owning one's home is found to be associated with a lower unemployment probability (Coulson and Fisher, 2009), smaller risk of becoming unemployed (van Leuvensteijn and Koning, 2004; Munch *et al.*, 2008), shorter unemployment durations (Munch *et al.*, 2006; Flatau *et al.*, 2003) and higher wages (Munch *et al.*, 2008). All of the aforementioned individual-level results are obtained when correcting for the presumed endogeneity of housing tenure status. Therefore, the findings of the micro-level studies seem to be in conflict with Oswald's empirical results and the results of the other papers that use aggregate data. Since the labour market outcomes of homeowners are generally more favourable than those of renters, regions with higher home-ownership rates should experience lower unemployment rates. This is generally not true, which means there might be some other mechanisms at work than those found by the studies thus far.

In this paper, we use Finnish individual-level data to study the effects of home-ownership on unemployment and, more generally, the labour market. We allow home-ownership to have external labour market effects. More specifically, we allow labour market outcomes of individuals to be affected by the overall home-ownership

rate in their region. Tests based on recent research are used to test new hypotheses on the mechanisms through which the externalities may work. To identify the causal effect of regional home-ownership on individual labour market outcomes, we exploit a rental housing market deregulation reform in the early 1990's. The reform produced a natural experiment that provides regional and time variation in home-ownership. Our results show that home-ownership has a significant positive external effect on unemployment, whereas, at the same time, homeowners are less likely to be unemployed than non-owners. Our results are, thus, consistent with both the harmful and the beneficial labour market effects of home-ownership found in the earlier literature. In the light of the additional analyses, it seems probable that debt-financed home-ownership hurts the local labour market by causing reductions in consumption demand. Although home-ownership has the potential to boost the labour supply of homeowners, the positive effects may be at least partly offset via displacement effects in the short-run.

3.2   Econometric Model and Data

Most of the earlier studies on the association between home-ownership and unemployment have used either aggregate or individual-level data. We combine an individual-level data set with region-level information on home-ownership to estimate probit models for whether an individual experienced unemployment during the year. The key explanatory variable is the regional rate of home-ownership. Further, we control for various individual characteristics, including a dummy variable for living in an owner-occupied dwelling and a dummy for a mortgage loan to capture the impact of individual housing tenure. We also include year dummies as well as county dummies.

By including regional-level home-ownership in our model, we allow regional home-ownership to have an effect on unemployment probability of an individual, given her own housing tenure. Although the origin of this external effect is unclear, there may be several different reasons for it. After first identifying the externality, we discuss the possible interpretations of it and perform analyses that shed some light on the mechanisms involved[1]. Since we are interested in the causal effect of regional home-ownership on unemployment, we need to take into account possible endogeneity of regional home-ownership. The results in Oswald (1996) come from simple regional-level regressions, and the author argues that his coefficient estimates may understate the positive causal effects. Assuming exogenous regional home-ownership would yield similarly biased estimates in our study as well. Theoretically, regional home-ownership depends on the supply of and demand for owner-occupied housing. The endogeneity of the home-ownership rate probably arises from the fact that the regional demand for

---

1   Most attempts to empirically identify externalities have been made in the literature on the effects of education (see e.g. Acemoglu and Angrist 2000). Externalities in the labour market have been discussed and estimated by Crépon *et al.* (2013) and Blundell *et al.* (2004).

JANI-PETRI LAAMANEN

owner-occupied housing depends positively on the employment of individuals residing in the region. Regional labour supply and demand shocks are, thus, likely to induce a negative association between home-ownership and unemployment. Controlling for labour supply and demand factors would alleviate the endogeneity problem and reduce the downward bias in the coefficient of regional home-ownership. Thus, it is not surprising that many of the earlier studies that include a broad range of regressors in their unemployment equation estimate a positive coefficient on the home-ownership variable (Costain and Reiter, 2008; Di Tella and MacCulloch, 2005; Nickell, 1998). However, as Oswald (1996) points out, instrumental variables are needed to obtain an unbiased estimate of the causal effect. Appropriate instruments for home-ownership are rare, and, therefore, it is likely that the earlier literature has been unable to identify the causal effect reliably[2]. To our knowledge, our paper is the first to use a policy reform to identify the causal effect of regional home-ownership on unemployment. We first estimate a model assuming exogenous regional home-ownership rates. We then proceed to relax this assumption and use the rental housing market deregulation reform to construct instrumental variables to deal with endogeneity.

We employ a Finnish register-based data set augmented with information on regional housing markets and the rental housing market deregulation reform. The individual-level data set used is a service file of annual Income Distribution Statistics (IDS) for years 1990–1992, which includes a rich set of register and survey variables on more than 30,000 individuals in more than 11,000 households per year. IDS contains information on individuals' labour market outcomes during the year. Specifically, we know the number of unemployment and employment months reported by the individual. The data includes information on individuals' housing tenure and important control variables, such as sex, education, age and household composition. Information on the place of residence is included for each respondent, which allows us to match regional-level home-ownership rates to the data and to divide the sample individuals into those affected by the deregulation reform and those not affected.

3.3   Analysis and Results

3.3.1   Descriptive Analysis

To conduct a descriptive analysis, we first estimate the unemployment model for the years 1990–1992 using county-level home-ownership as an explanatory variable. To reflect the average home-ownership, the rate of home-ownership is calculated as the average of the year-end value and the previous year-end value. We have restricted the

---

2   To our knowledge, the paper by Coulson and Fisher (2009) using data from the United States is the only study that uses instrumental variables. Their instruments are the state marginal tax rate applied to mortgage interest deduction and the percentage of households living in multifamily housing.

sample to include individuals of working age (from 15 to 64 years old) only. County dummies and year dummies as well as individual-level control variables are included in the model. Results of this exercise are presented in Table 1. Because the model includes an aggregate-level regressor, we need to account for the possibility of a correlation between the error terms of individuals within a county. We use robust standard errors that allow for such correlation (see Moulton, 1986).

The coefficient for regional home-ownership is positive and statistically significant, and, thus, we are able to replicate Oswald's (1996) results using individual-level data and controlling for characteristics of individuals. It appears that regional home-ownership is positively associated with unemployment, while owner-occupiers (especially those with mortgage loans) are less likely to be unemployed than other individuals[3].

At this point, it is worth noting that there is a potential danger in including the regional rate of home-ownership and individual home-ownership variable (as well as the mortgage variable) in the same model[4]. Namely, variation in the regional rate of home-ownership captures part of the variation in individual home-ownership and can be thought of as a result from a first stage regression of individual home-ownership on region-year interaction dummies. Angrist (2013), and, earlier, Angrist and Pischke (2009) note that running a model with this kind of instrumented variable and the original variable may lead one to conclude that there are externalities in cases where they are, in fact, absent. This happens if the original variable is endogenous for any reason. In turn, excluding the group-average variable from the regression leads to omitted variable bias in the estimate of the original variable if there are externalities. Thus, including both variables is a correct strategy when externalities are suspected (as in our case) but may lead to false conclusions if the individual-level variable is endogenous.[5]

The aforementioned logic applies to all analyses in this paper, and we want to make sure that our conclusions are not misled by the aforementioned issues. Thus, we estimate alternative versions of our key unemployment and employment models in which the individual home-ownership variable (as well as the mortgage variable, which is likely to be endogenous) is excluded.[6] The results of these estimations and comparisons to the estimates in the tables of results can be found in the Appendix. The problem with these models is that we do not get separate estimates of direct effects and externalities but, instead, the estimated coefficient of the regional-level regressor captures them both. It

---

3   Our mortgage variable is at individual level and the owner-occupancy variable at the household level. We also constructed a variable indicating a mortgage by another household member but it was not statistically significant in the regressions.

4   We are grateful to Bas van der Klaauw for pointing this out to us.

5   For more discussion on these issues, see Angrist (2013).

6   We also present results from models in which the regional home-ownership variable is excluded.

Table 1. Model of Unemployment Experience[1]

Regional home-ownership		
County home-ownership rate	4.17***	(1.53)
Personal characteristics		
Mortgage	-0.0443***	(.006)
Owner-occupier	-0.0393***	(.004)
Male	0.0402***	(.005)
Age	0.0094***	(.001)
Age squared	-0.0001***	(.000)
Marital status		
Single	-0.0294***	(.010)
Married	-0.0616***	(.009)
Separated	-0.0172	(.026)
Divorced	-0.0236	(.015)
Unknown	-0.0151	(.013)
Household size	0.0047*	(.002)
Number of children	-0.0198***	(.003)
Education	0.0620***	(.015)
Basic or no degree	0.0926***	(.015)
Lower secondary	0.0581***	(.014)
Higher secondary	0.0218	(.019)
Vocational college	-0.0170	(.018)
Higher University	-0.1102*	(.057)
Graduate school		
Type of municipality	-0.0284	(.018)
Capital region	-0.0084	(.006)
Urban	-0.0047	(.007)
Semi-urban		

[1] Marginal effects calculated at sample means from a probit regression including small county dummies and year dummies. N = 30,316. Robust clustered (county-year) standard errors in parentheses. (*) denotes significance at 10% level, (**) at 5% level and (***) at 1% level. The omitted category is widowed females with a lower university degree residing in a rural region.

can be seen that the differences in results are negligible[7]. Thus, our main finding that regional home-ownership increases unemployment is robust to this modification. We also argue that the direct individual-level effect of home-ownership is almost surely negative or zero, as has been found in earlier studies. It is very difficult to imagine a theoretical model in which buying one's home causes one to become unemployed. Thus, we interpret the statistically significant positive coefficient of regional home-ownership as a sign of externalities.

The estimated coefficient of regional home-ownership might seem to be relatively large, because a one percentage point increase in the rate of home-ownership is associated

---

7   For example, the marginal effect of regional home-ownership is 4.17 when home-ownership and mortgage dummies are included (as in Table 1) and 4.29 when they are excluded.

with about a 4 percentage point increase in unemployment. However, our dependent variable is a measure of the unemployment experience (from 1 to 12 months) during the year, and, thus, the coefficient estimate cannot be interpreted as the effect on the unemployment rate. We have tried estimating our models using different definitions of the dependent variable. It appears that the overall effect mostly comes from short-term or occasional unemployment i.e. probability of experiencing small number of months of unemployment. The result on the association between having a mortgage loan and being less likely to experience unemployment is in line with results by Flatau *et al.* (2003), who find that homeowners with mortgage experience shorter unemployment spells than homeowners without mortgage and individuals in other tenures.

3.3.2 The Rental Housing Market Deregulation Reform

To investigate whether causality runs from regional home-ownership to unemployment, we use a rental housing market deregulation reform to estimate instrumental variables models of unemployment. The Finnish rental housing market was deregulated in the early 1990's. The most important feature of the reform was that rent ceilings and exact limits on rent increases were removed. In addition, eviction without specifying its grounds was made easier. There was a serious shortage of rental housing in the country, and the goal of the reform was to encourage supply in the private rental sector. The opposition was worried that the reform would lead to significant price increases, and, therefore, the government wanted to experiment with deregulation in parts of the country. Specifically, the bill was a proposal to relax some of the existing rental housing market regulations in regions where 'demand and supply of rental housing are in approximate balance'. The seven counties of northern and central Finland were chosen as the region for the experiment. To exclude regions of large supply deficiency in the rental housing market, all university cities were left out of the experiment. However, the opposition argued that there were at least some localities that suffered from a shortage of rental housing in the region of the experiment and that the regulations were, therefore, at least partly binding. Theoretically the reform had the effect of increasing the supply of rental housing and, thus, decreasing the rate of home-ownership. This is the case because the regulations had constrained the income accrued to landlords and made eviction relatively difficult. Thus, deregulation increased the supply of rental dwellings by making renting more profitable for landlords. As we later demonstrate, the reform had a negative effect on home-ownership rates of the reform regions during our sample period.

The reform was implemented during the 1991–1994 period, and the market was deregulated gradually. Different buildings and rental contracts were subject to deregulation in different years and, important to our identification strategy, the timing of the first phase of the reform differed between geographical areas. In the 'target

regions', the seven counties in northern and central Finland (excluding the six largest cities), markets were deregulated earlier than those in the five remaining counties. In the first phase, in the beginning of 1991, new apartments and houses in the target region were subject to deregulation. The second change in the legislation, effective on the 1st of February 1992, deregulated all new contracts for private rental dwellings in the whole country. Finally, in the beginning of the year 1994, all rental contracts were deregulated. Thus, dwellings constructed and rented between the 1st of January 1991 and the 1st of February 1992 were free from regulation in the target region but not in other regions. This means that in the target region only, some dwellings were free from regulation for 13 months. At the end of 1991, dwellings constructed during that year constituted approximately 2.3 percent of the housing stock in the country.

To serve as a relevant instrument, the reform needs to have an effect on regional home-ownership. To give and idea of the effect, we have calculated home-ownership rates over time from aggregate housing data. Comparison between the development of home-ownership in the target regions and elsewhere reveals that the reform decreased home-ownership. Figure 1 plots aggregate home-ownership rates in reform regions and other regions four years prior to and four years after the reform. The vertical line indicates the reform in the beginning of 1991. Since the target regions comprised of municipalities with fewer than 15,415 dwelling units, we have excluded municipalities larger than this from the calculation of the home-ownership rate of the other regions as well. In 1991, when the reform was implemented in the target regions only, the home-ownership rate of these regions decreased relative to home-ownership rate elsewhere. Both before and after 1991, the home-ownership rates of the two regions followed an approximately similar trend.

Although the reform seems to have had an effect on home-ownership, this effect is, admittedly, rather small. A closer look at the home-ownership rates of the regions reveals that the effect varies between municipalities of different types and sizes[8]. Figure 2 plots the home-ownership rates by municipality type. It can be seen that the effect is most pronounced in rural municipalities, although there is a visible effect in all municipality groups. During the 1991–1992 period, rural reform regions experienced a 0.6 percentage point fall in home-ownership, whereas in the other regions, the home-ownership rate increased by 0.5 percentage points. Clearly, the reform had an effect on the home-ownership differential between the two groups of regions. It seems that the change in the differential was persistent. These findings increase confidence in the identifying variation in home-ownership created by the reform.

As can be seen in the aforementioned figures, the trends in home-ownership in target regions and elsewhere before the reform were very similar. However, it is also important to our analysis that the choice of the target regions was independent of the labour market

---

8   In our instrumental variables analysis, we exploit the fact that it was only the new housing stock that was subject to the reform, and, thus, there were larger changes in the areas with more new dwellings.
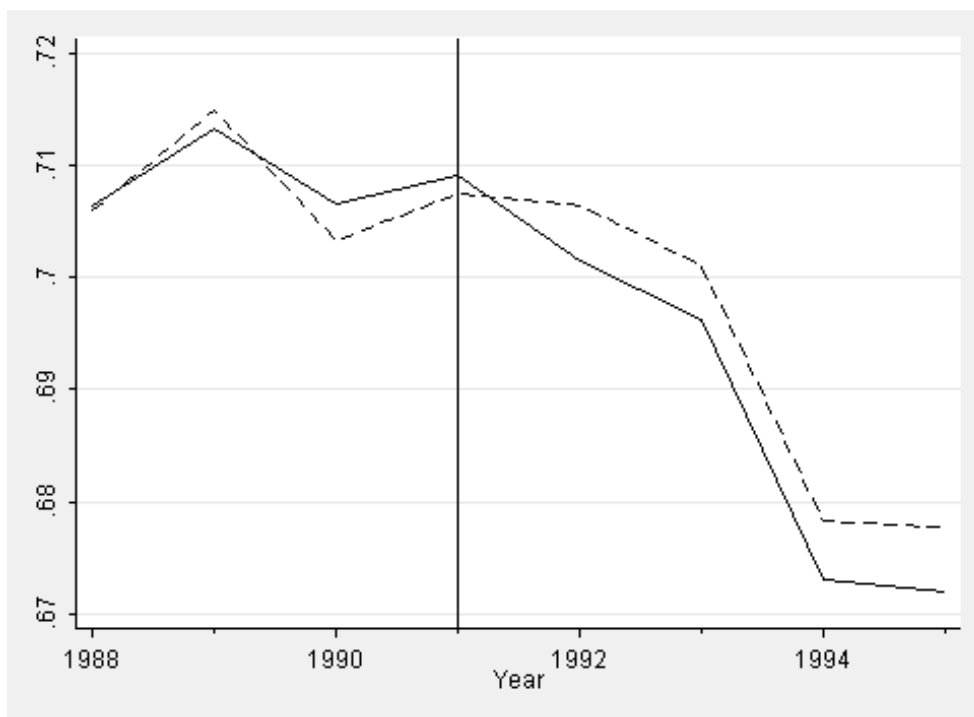
Figure 1: Trends in home-ownership in reform (solid line) and other (dashed line) regions

prospects of the regions. It appears from the bill and the preceding committee report that the choice of regions was based on rental housing market conditions only. References to labour markets are virtually absent from these documents. This absence might be because the labour market was not a concern, because of very low unemployment after the strong economic growth in the 1980's. As mentioned previously, the government wanted to experiment with deregulation in regions where rental housing markets were in approximate balance. This strategy was probably chosen because it would reduce the magnitude of unwanted consequences (rental price increases). However, neither a systematic analysis of regional markets was done nor did the committee suggest choosing any specific regions or any criteria. The government may have wanted to choose a set of administrative regions (counties) that constitute a single contiguous area. Based on the presumption that the housing markets were closer to equilibrium in the north than in the south, northern and central counties (excluding the university cities) were chosen. In light of the previous discussion, it seems clear that the choice of the reform regions was not related to labour market concerns.

Figure 2: Trends in home-ownership in reform (solid line) and other (dashed line) regions by municipality type (top left: urban, top right: semi-urban, bottom: rural)

3.3.3 Instrumental Variables Analyses

In constructing the instruments for the regional home-ownership variable, we take into account that two separate phases (1991 and 1992) of the rental housing market reform took place during the period of investigation. Furthermore, in six counties, the reform excluded the largest cities, whereas in one county, all municipalities were included. We construct four dummies that reflect the two phases and two different scales of the reform. The first dummy (A) is for the county that was fully exposed to the reform in 1991. The second dummy (B) is for this county in 1992. The third dummy (C) is for the other six reform counties where largest cities were excluded in 1991. The fourth dummy (D) is for these counties in 1992. Separate dummies are constructed for years 1991 and 1992 because of different phases of the reform but also because we wish to allow the effects of the reform on home-ownership to evolve over time (see the previous section). Since only new dwellings were deregulated in the reform, the dummies are

interacted with the share of deregulated multifamily dwellings in the county[9]. Since we use a midyear measure of the home-ownership rate, the relevant new housing stocks are half the number of dwellings constructed in 1991 and dwellings constructed in January 1992 for the instruments of 1991 and 1992, respectively[10]. We expect that each of the instruments has a negative effect on regional home-ownership.

Table 2 presents the results from estimating an instrumental variables probit model of unemployment. The first-stage results indicate that the reform indeed had an effect of decreasing home-ownership in the target regions. The effect of regional home-ownership on unemployment is positive and statistically significant. This finding confirms the result of the descriptive analysis in Table 1. As expected, the estimated coefficient is much larger when instrumental variables are used. This is because the estimate in Table 1 is contaminated by the endogenous variation in the rate of home-ownership. Our results are consistent both with earlier results obtained by using microdata and the results of Oswald (1996) and other studies using aggregate data. Home-ownership seems to be positively associated with individual owner-occupiers' labour market outcomes *and*, at the same time, increase overall unemployment.

Although the main aim of our analysis is to study the causal effect of home-ownership on unemployment, our data also permits us to explore some possible mechanisms through which the effect might work. Oswald (1999) and Blanchflower and Oswald (2013) discuss various potential labour market effects of home-ownership, some of which involve external impacts. The externalities are related to labour immobility and not-in-my-backyard (NIMBY) behaviour. Immobility of homeowners makes it more difficult for other people to move close to ideal jobs as well, and the labour market, as a whole, becomes less efficient. NIMBYism and less job creation may follow from home-ownership, if homeowners try to prevent businesses from locating to their home region. However, these mechanisms are more likely to be relevant in the longer run, and we focus on the immediate effects of home-ownership. In what follows, we conduct analyses that shed light on the relevance of two new hypotheses based on the idea that households' decisions about labour supply and consumption may be linked to changes in their housing tenure. This happens because most home purchases, typically financed

---

9   Interacting with the share of multifamily (rather than all) dwellings improves the strength the of instruments because renting single-family houses is rare. We also estimated models with dummy instruments and dummies interacted with all new dwellings. These alternative instruments appeared to have satisfactory explanatory power as well, and the qualitative results did not differ from those presented here.

10   Our dwelling data is annual but we have quarterly data on buildings by building type. The number of dwellings constructed in January 1992 is measured as one-third of the share of buildings constructed in the first quarter multiplied by the number of dwellings constructed during the whole year. The exact way of measuring the new housing stock has only minor impacts on the estimates and does not influence the qualitative results.

Jani-Petri Laamanen

Table 2. IV Model of Unemployment Experience[1]

Regional home-ownership		
County home-ownership rate	9.51***	(3.14)
Personal characteristics		
Mortgage	-0.0443***	(.006)
Owner-occupier	-0.0393***	(.004)
Male	0.0403***	(.005)
Age	0.0094***	(.001)
Age squared	-0.0001***	(.000)
Marital status		
Single	-0.0296***	(.010)
Married	-0.0616***	(.009)
Separated	-0.0171	(.026)
Divorced	-0.0228	(.015)
Unknown	-0.0153	(.013)
Household size	0.0048*	(.002)
Number of children	-0.0198***	(.003)
Education		
Basic or no degree	0.0624***	(.015)
Lower secondary	0.0928***	(.016)
Higher secondary	0.0582***	(.014)
Vocational college	0.0222	(.019)
Higher University	-0.0166	(.018)
Graduate school	-0.1104*	(.057)
Type of municipality		
Capital region	-0.0282	(.018)
Urban	-0.0081	(.006)
Semi-urban	-0.0046	(.007)
First-stage results: Excluded instruments		
Instrument A	-0.1661***	(.043)
Instrument B	-0.6135***	(.188)
Instrument C	-0.1821***	(.021)
Instrument D	-0.4877***	(.119)

[1] Marginal effects calculated at sample means from a probit regression including county dummies and year dummies. N = 30,316. Robust clustered (county-year) standard errors in parentheses. (*) denotes significance at 10% level, (**) at 5% level and (***) at 1% level. The omitted category is widowed females with a lower university degree residing in a rural region.

by mortgage loans, are associated with increased housing expenditures[11]. Our two hypotheses are the ones that might be relevant in the light of recent research and can be tested using our data, but we do not claim that there are no other (positive or negative, short-run or long-run) impacts of home-ownership on the labour market.

---

11   Scanlon and Whitehead (2004) have documented that the average expenditures on housing have been higher for owners than for renters in most countries both in the 1990's and in the 2000's.

It has been argued by some authors that homeowners' favourable labour market outcomes are caused by homeowners' lower reservation wages and relatively active local job search, which arise from the need to meet high mortgage payments and the reluctance to move to another region (e.g. Munch *et al.*, 2006 and Flatau *et al.*, 2001). Recent research suggests that this may produce negative externalities in the local labour market. As one group of individuals increases their labour supply and working, other individuals may be displaced from jobs as a result. This is especially likely in the short run, since short-run labour demand may not be sufficiently elastic. Crépon *et al.* (2013) and Blundell *et al.* (2004) have studied displacement effects associated with labour market programmes. Abbring and Heckman (2007) includes a thorough discussion on displacement and various other types of external effects in the labour market. Since, in theory, displacement may be associated with any increase in labour supply, we argue that high job-search intensity and low reservation wages of new homeowners may lead to displacement of other workers in the same region.

In addition to boosting homeowners' labour supply, mortgage loans associated with home purchases may also affect consumption. Some recent studies have examined the effects of household credit on consumption behaviour. It has been argued that credit-constrained households with debt are forced to cut back on spending when house prices are declining (Mian *et al.*, 2013 and Mian and Sufi, 2010; Dynan, 2012). However, debt may, under some circumstances, be negatively linked to consumption even in the absence of unexpected changes in asset values. Stephens (2008) shows that repayment of a vehicle loan leads to an increase in nondurable consumption. Coulibaly and Li (2006) examine the effect of a final mortgage payment and find that it is associated with an increase in durable consumption (house furnishings and entertainment equipment). Although the credit-constraint explanation of these results does not imply that individuals would cut back on consumption when they become borrowers, the possibility that home buying and associated mortgage borrowing decrease consumption can not be ruled out *a priori*[12]. For example, one can imagine reasons why households do not want or are unable to save as much before they buy their home as after the purchase. Many countries have policies that favour home-ownership and make it a relatively profitable method of saving. Thus, it may be optimal for households to start saving more (and consume less) after a home purchase. Another possible mechanism that would lead to decreased consumption is related to self-control problems that prevent individuals from saving as much as they would like. Laibson (1997) has studied self-control problems and has mentioned mortgage contracts as an example of commitment devices that help

12   Notice also that labour supply and consumption choices are both part of consumer's optimisation problem. Thus, if there is a change in the amount of labour supplied due to a change in housing tenure, consumption is likely to change as well and *vice versa*. Thus, the claim that housing tenure affects consumption is closely linked to the claim that housing tenure affects labour supply.

individuals force themselves to save more[13]. A recent study by Mian and Sufi (2012) links consumption decreases caused by house price reductions to decreases in local employment. Similarly, if an increase in home-ownership and associated mortgage borrowing leads to a reduction in household spending, this may be negatively reflected in the local labour market.

Both the displacement hypothesis and the consumption hypothesis lead to empirically testable predictions about the labour market externalities of home-ownership. In Table 3, we start by presenting coefficient estimates of regional home-ownership separately for homeowners with mortgages, owner-occupiers who do not have a mortgage loan and non-owners. It appears that only homeowners with a mortgage are not affected by the external impact, and the effect is larger for non-owners than for owner-occupiers. In the light of the displacement story, non-owners may be more likely to be displaced than owners. This happens because their job search is not intensified and their reservation wages are not lowered by their housing tenure, which arguably makes non-owners more vulnerable to changes in local competition for jobs[14].

We next see if results in line with the displacement hypothesis can be found. Clearly, displacement is strongest for individuals whose labour is a close substitute for the labour of individuals who recently bought their homes and, as a result, increased their labour supply. As a simple test, we estimate the unemployment model separately for individuals whose personal characteristics are similar to those of recent buyers and for individuals who are dissimilar to them. To do this, we first estimate a probit model of having bought a home. The explanatory variables are the same personal characteristics as in our unemployment model (gender, age, age squared, marital status, household size, number of children and education). We then divide the sample in two, based on the predicted probability of being a buyer. Estimating the unemployment model for the individuals who are similar to buyers yields a large and significant estimate, whereas the estimate for the dissimilar half of the sample is much smaller and not statistically significant[15]. Although the result is in line with the displacement hypothesis, it may be that it is a mere coincidence that individuals who are similar to new homeowners are largely affected by regional home-ownership.

Next, we perform another analysis related to displacement. Crépon *et al*. (2013) note that the more there are individuals in a market whose job search is positively affected,

---

13   Ashraf *et al*. (2006) find empirical support for the claim that individuals are willing to use commitment devices to be able to save more.

14   In analyses not presented here, we also found that the effect is more pronounced on groups whose labour market attachment is likely to be weak and who therefore are more likely to be affected by job competition. These groups include students, mothers of pre-school chidren and individuals younger than 25 years old. The results are available on request.

15   We lose some observations in the sample of individuals similar to buyers because graduate schooling predicts non-unemployment perfectly in this sub-sample. We have tried combining the two highest educational dummies and including all observations and the results did not change.

the larger the displacement is likely to be among the job seekers in that market. We do not have data on job search, but we know the sector of all sample individuals who work. Applying the idea of Crépon *et al.* (2013) to our case implies that the greater the number of buyers working in a given sector, the fewer the chances for other workers finding employment in that sector. Thus, we calculate from our sample the number of buyers working in each sector and, based on this figure, divide sectors in two groups with both representing about half of the total employment. In Table 3, we first present the overall effect of regional home-ownership on the probability of working during the year and then estimate the model separately for working in the two groups of sectors. The overall effect of regional home-ownership on employment is negative and, thus, in line with the estimated effect on unemployment. We find a negative employment effect in sectors where a large number of buyers are working ('Buyers' sectors'), whereas the estimated effect is not statistically significantly different from zero for sectors with relatively few buyers ('Other sectors'). This result is in line with the displacement hypothesis, and it seems that individuals whose labour supply is increased because of a home purchase may displace other workers in sectors where they work.

The two tests that we have performed seem to support the displacement hypothesis. We next turn to the externality mechanism that works through a reduction in consumption resulting from home-ownership. Mian and Sufi (2012) argue that the employment changes that are caused by consumption changes can be identified by examining employment changes separately in non-tradable and tradable sectors. Tradable sectors produce goods that can be consumed outside the region where they are produced. Non-tradable goods, in turn, are consumed locally. Therefore, a local change in consumption demand has a local impact on producers of non-tradable goods, but the impact on tradable sectors is spread over regions. This makes it possible to infer whether employment changes are caused by changes in consumption. We use a strategy similar to that of Mian and Sufi (2012) and test whether the local employment effect of region's home-ownership differs between non-tradable and tradable sectors[16]. The results presented in Table 3 show that home-ownership has a negative effect on employment in non-tradable sectors but has no effect on employment in tradable sectors. Although it is possible that some other mechanism is behind this result, our evidence is in line with the idea that home-ownership decreases employment by decreasing consumption.

It is not possible to reliably disentangle the displacement effect and the consumption effect with our data. We perform one more test to shed light on which of the two effects is likely to be of more importance. We divide the sectors into four groups based on

---

16   The sectors classified as non-tradable include retail and wholesale, restaurants, bars, canteens and catering, taxis, motion picture and video production and distribution, motion picture projection, arts performances, concerts and artistic creation and interpretation, libraries, archives, museums and exhibitions, sporting activities and operation of sport arenas and stadiums, dance halls and dancing schools, gambling, circus, amusement parks, other recreational activities, hairdressers, beauty salons, photo portraits, day care, funerals and other personal and household services.

Jani-Petri Laamanen

Table 3. IV Models of Unemployment Experience and Employment Experience[1]

Effect on sub-group unemployment				
Homeowners with mortgage	2.49	(3.79)	N =	6,682
Owner-occupiers, no mortgage	9.89***	(3.77)	N =	17,625
Non-owners	15.99***	(3.17)	N =	5,998
Similar to buyers	14.59***	(3.99)	N =	15,140
Dissimilar to buyers	4.43	(2.88)	N =	15,157
Effect on employment				
All employment	-6.21***	(1.79)	N =	30,316
Employment in				
buyers' sectors	-4.61**	(2.02)	N =	30,316
other sectors	1.60	(3.16)	N =	30,316
Employment in				
non-tradable sectors	-4.58**	(2.02)	N =	30,316
tradable sectors	1.33	(3.13)	N =	30,316
Employment in				
non-tradable & buyers' sectors	-1.33*	(.805)	N =	30,316
non-tradable & other sectors	-3.00*	(1.60)	N =	30,316
tradable & buyers' sectors	-3.13	(1.91)	N =	30,316
tradable & other sectors	4.78	(3.70)	N =	30,316

[1] Marginal effects calculated at sample means from IV probit regressions including county dummies and year dummies. Robust clustered (county-year) standard errors in parentheses. (*) denotes significance at 10% level, (**) at 5% level and (***) at 1% level. Control variables as in Tables 1. and 2.

the two previous classifications. If there were displacement but no consumption effect, we would find an effect only on sectors with large number of buyers (non-tradable/ buyers' and tradable/buyers'). If instead there were only the consumption effect but no displacement, we would find an effect only on non-tradable sectors (non-tradable/ buyers' and non-tradable/other). Finding an effect for all groups of sectors except the tradable/other group would indicate coexistence of the two effects. The results in Table 3 indicate that there is only the consumption effect[17]. However, the effect on tradable/buyers' sectors is very close to significance. Based on these analyses, it would be safe to conclude that home-ownership might influence employment through both displacement and consumption effects.

It should be noted that the displacement effect would have to be very large for it alone to generate a positive relationship between the rate of home-ownership and the unemployment rate. It would require that more than one worker become unemployed by every newly employed home-owner. This would be possible only if homeowners worked longer hours than the displaced workers. We do not think that it is a credible claim that each homeowner displaces more than one other worker. Since our dependent

---

17   Note that the first of the four parameters loses statistical significance when individual home-ownership dummy and mortgage dummy are excluded.

variables measure any experience of unemployment and employment, from one month to 12 months, more than one-to-one displacement is more likely. It appears from analyses not presented here that the unemployment and employment effects are more pronounced in the experience of only a few months of unemployment or employment, whereas the effect is smaller or negligible on full-year unemployment and employment. At the same time, homeowners with mortgages and those who bought their homes are less likely to experience small numbers of unemployment or employment months. Thus, at the extreme, a displacement in our data may mean a displacement of 12 one-month jobs by one 12-month job.

There are no obvious bounds on the magnitude of the consumption effect, since it depends on the reduction in money spent and on its effect on employment. It should be noted that there is no guarantee that the estimated consumption effect on employment equals the total consumption effect on employment. This is the case if the consumption change due to home-ownership differs between non-tradable and tradable goods. For example, if home buying is not only associated with a reduction in spending on non-tradable goods but also with an increase in spending on tradable goods, the effect on employment is negative locally but may even be positive globally. On the other hand, the estimated effect on the local labour market may understate the global effect if there is a reduction in consumption of tradable goods that leads to a larger employment effect than the reduction in consumption of non-tradable goods. Thus, further research on the effects of housing tenure on consumption of different goods is needed to shed light on the consumption effect of home-ownership and whether it is local or has wider labour market consequences as well. A final note should be made about the consumption effect. During our period of analysis, house prices were rapidly falling and interest rates were rising in the wake of a major economic crisis. These developments were likely to cause a reduction in household spending because of unexpected household balance sheet changes and increases in mortgage payments[18]. Although higher home-ownership means that there are more households susceptible to the changes, and, thus, more households which cut back on consumption, we believe that this may have only a small (positive) effect on our estimates. This is the case because we study the immediate effects of changes in home-ownership. It is unlikely that there was sufficient time for the house price and interest rate shocks to affect the households that bought their homes.

---

18   Most mortgages in Finland were, and still are, variable-rate mortgages.

3.4   Conclusions

Various policies influence individual housing tenure choices. The earlier evidence on the labour market effects of these policies is partly mixed. In particular, many studies have found that higher prevalence of home-ownership is associated with higher aggregate unemployment, whereas studies using microdata suggest that homeowners have relatively favourable labour market outcomes. We find that, in addition to the latter result, home-ownership has effects above and beyond the direct effects on individuals. Namely, although homeowners are less likely to experience unemployment, significant externalities counteract the positive effects of home-ownership at the aggregate geographical level. The external effects may be because of consumption reductions and increased local job competition caused by home purchases, especially if the purchases are financed by debt. More theoretical and empirical research is needed to better understand the mechanisms at work.

Appendix

Table A1. Alternative Model Specifications[1]

	(a)	(b)	(c)
Exogenous regional home-ownership			
Table 1: Unemployment experience			
County home-ownership rate	4.17***	4.29***	-
Mortgage	-0.04***	-	-0.04***
Owner-occupier	-0.04***	-	-0.04***
Endogenous regional home-ownership			
Table 2: Unemployment experience			
County home-ownership rate	9.51***	9.49***	-
Mortgage	-0.04***	-	-0.04***
Owner-occupier	-0.04***	-	-0.04***
Table 3: Employment experience			
County home-ownership rate	-6.21***	-6.21***	-
Mortgage	0.11***	-	0.11***
Owner-occupier	0.00	-	0.00

[1] Marginal effects calculated at sample means from a probit regression including county dummies and year dummies. N = 30,316. Robust clustered (county-year) standard errors in parentheses. (*) denotes significance at 10% level, (**) at 5% level and (***) at 1% level. The omitted category is widowed females with a lower university degree residing in a rural region. Control variables as in Tables 1. and 2. Specifications (a) are as in Tables 1, 2 and 3, and include county home-ownership, individual home-ownership and mortgage variables. Specifications (b) exclude individual-level home-ownership and mortgage variables, and (c) exclude county home-ownership rate.

References

Abbring, J. H. and Heckman, J. J. (2007). Econometric Evaluation of Social Programs, Part III: Distributional Treatment Effects, Dynamic Treatment Effects, Dynamic Discrete Choice, and General Equilibrium Policy Evaluation. In: J.J. Heckman and E.E. Leamer (eds.), *Handbook of Econometrics*, Edition 1, Vol. 6B. Elsevier.

Acemoglu, D., and J. Angrist (2000). How Large Are the Social Returns to Education: Evidence from Compulsory Schooling Laws. In *NBER Macroeconomics Annual 2000*, 9–74.

Angrist, J. (2013). The Perils of Peer Effects. NBER Working Paper No. 19774.

Angrist, J. and J.-S. Pischke (2008). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Ashraf, N., Karlan, D. and Yin, W. (2006). Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines. *Quarterly Journal of Economics*, 121(2), 635–672.

Battu, H., Ma, A. and Phimister, E. (2008). Housing Tenure, Job Mobility and Unemployment in the UK. *The Economic Journal*, 118, 311–328.

Blanchflower, D. G. and Oswald, A. (2013). *Does High Home-Ownership Impair the Labor Market?* NBER Working Paper No. 19079.

Blundell, R., Costa Dias, M., Meghir, C. and Van Reenen, J. (2004). Evaluating the Employment Impact of a Mandatory Job Search Program. *Journal of the European Economic Association*, 2(4), 569–606.

Costain, J. S. and Reiter, M. (2008). Business Cycles, Unemployment Insurance, and the Calibration of Matching Models. *Journal of Economic Dynamics and Control*, 32(4), 1120–1155.

Coulibaly, B. and Li, G. (2006). Do Homeowners Increase Consumption After the Last Mortgage Payment? An Alternative Test of the Permanent Income Hypothesis. *Review of Economics and statistics*, 88(1), 10–19.

Coulson, N. E. and Fisher, L. M. (2009). Housing Tenure and Labor Market Impacts: The Search Goes on. *Journal of Urban Economics*, 65, 252–64.

Crépon, B., Duflo, E., Gurgand, M., Rathelot, R. and Zamora, P. (2013). Do Labor Market Policies have Displacement Effects? Evidence from a Clustered Randomized Experiment. *Quarterly Journal of Economics*, 128(2), 531–580.

Di Tella, R. and MacCulloch, R. (2005). The Consequences of Labor Market Flexibility: Panel Evidence Based on Survey Data. *European Economic Review*, 49(5), pp. 1225–59.

Dynan, K. (2012). Is a Household Debt Overhang Holding Back Consumption? *Brookings Papers on Economic Activity*, 299–362.

Flatau, P., Forbes, M., Hendershott, P. H. and Wood, G. A. (2003). *Homeownership and Unemployment: The Roles of Leverage and Public Housing*. NBER Working Paper No. 10021.

Green, R. and Hendershott, P. (2001). Home-ownership and Unemployment in the US. *Urban Studies*, 38, 1501–1520.

Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics*, 112(2), 443–478.

Mian, A. R., Rao, K. and Sufi, A. (2013). Household Balance Sheets, Consumption, and the Economic Slump. *Quarterly Journal of Economics*, forthcoming.

Mian, A. and Sufi, A. (2010). The Great Recession: Lessons from Microeconomic Data. *American Economic Review, Papers & Proceedings*, 100(2), 51–56.

Mian, A. R. and Sufi, A. (2012). *What Explains High Unemployment? The Aggregate Demand Channel*. NBER Working Paper No. 17830.

Moulton, B. R. (1986). Random Group Effects and the Precision of Regression Estimates. *Journal of Econometrics*, 32, 385–397.

Munch, J. R., Rosholm, M. and Svarer, M. (2006). Are Homeowners Really More Unemployed? *The Economic Journal*, 116, 991–1013.

Munch, J. R., Rosholm, M. and Svarer, M. (2008). Home Ownership, Job Duration, and Wages. *Journal of Urban Economics*, 63, 130–145.

Nickell, S. J. (1998). Unemployment: Questions and Some Answers. *The Economic Journal*, 108, 802–816.

Oswald, A. J. (1996). A Conjecture on the Explanation for High Unemployment in the Industrialized Nations: Part I. University of Warwick Working Paper No. 475.

Oswald, A. J. (1999). The Housing Market and Europe's Unemployment: A Non-technical Paper. University of Warwick.

Scanlon, K. and Whitehead, C. (2004). *International Trends in Housing Tenure and Mortgage Finance*. Council of Mortgage Lenders, London.

Stephens, M. Jr. (2008). The Consumption Response to Predictable Changes in Discretionary Income: Evidence from the Repayment of Vehicle Loans. *Review of Economics and Statistics*, 90(2), 241–252.

Van Leuvensteijn, M. and Koning, P. (2004). The Effects of Home-ownership on Labour Mobility in the Netherlands. *Journal of Urban Economics*, 55, 580–596.

Chapter 4
Welfare State and Life Satisfaction:
Evidence from Public Health Care

Jani-Petri Laamanen

Welfare State and Life Satisfaction: Evidence from Public Health Care

By Kaisa Kotakorpi† and Jani-Petri Laamanen‡

†*University of Tampere, Finland*   ‡*University of Tampere and FDPE*

Final version received 31 March 2008.

We examine the link between the welfare state and citizens' subjective wellbeing by using evidence from public health care services. By combining local-level data on public health care and individual-level data on life satisfaction, we show that relatively high expenditures in health care have a positive effect on individuals' life satisfaction in our data. We find some evidence for an 'ends-against-the-middle' equilibrium in provision of public health care, where middle-income individuals prefer higher public expenditure than low-income or high-income individuals. Further, our results indicate that the welfare benefit of public provision depends on individual political orientation.

INTRODUCTION

The value of many of the services typically associated with the welfare state is difficult to measure, as many of these services have public goods characteristics. Further, even if some publicly provided services are essentially private in nature (such as in the case of health care and education), individuals' valuation for them cannot in most cases be observed directly as they are often offered free of charge or at heavily subsidized prices by the government.

Frey and Stutzer (2005) provide an extensive discussion of the problems associated with attempting to value public goods by traditional revealed preference and stated preference methods, and they argue for the use of individual-level life satisfaction data for this purpose. This method also has potential as a way of evaluating the welfare effects of the public provision of private goods. Adopting the life satisfaction approach is also parallel with the more general view, expressed for example by Layard (2006), according to which life satisfaction, or happiness, should be considered as the objective for public policies. The objective of the current paper is to examine the welfare effect of publicly provided health care by combining individual-level data on self-reports of life satisfaction by Finnish respondents of the World Values Survey for the year 2000, with local-level data on public health care services.

Paradoxically, in a cross-country study of 40 nations in 1980–90, Veenhoven (2000) finds no link between the size of the welfare state and the level of wellbeing[1] of a country's citizens. Further, in another cross-country study of 74 countries, Bjornskov *et al.* (2007) find a *negative* effect of higher government consumption spending on life satisfaction, interpreted as reflecting the tendency of rent-seeking politicians to overspend, for example, in order to maximize re-election probabilities. However, it is likely that individual welfare is affected primarily by the level of service provision—that is, aspects such as the coverage, range and quality of services—and not by government expenditures *per se*. Therefore, in order to obtain a clearer picture of the welfare effects of government activities, it seems necessary to construct measures more closely related to specific aspects of service provision. Indeed, Di Tella *et al.* (2003) use individual-level data from 13 countries, and find a positive welfare effect of a more specific measure of the welfare state, namely the income replacement rate of unemployment benefits.

We aim to examine the effect of public services on wellbeing in more detail by using local-level data and concentrating on a single sector. The World Values Survey includes a question on the respondents' place of residence, and we are therefore able to match the data with measures of the level of public health care services at the municipal level. The municipality level is the level at which primary health care is mainly organized in Finland, and it is also the lowest level of local government. In the analysis, instead of simply using data on raw expenditures in public health care, we aim to construct an expenditure measure potentially more closely related to the level of service provision or service quality. We do this by estimating the amount of excess resources spent on public health care, given the number of treatment days and visits to public health centres in each municipality. We argue in Section I that this approach is appropriate given the institutional context of the study. *A priori* it is of course not clear whether excess resources lead to quality improvements, or whether they are pure waste or reflect a regional clustering of severe health problems. It turns out that neither waste nor health problems can provide an explanation for our results: our results indicate that, controlling for individual characteristics and local conditions, relatively high expenditures in public health care have a positive effect on life satisfaction in our data.

A further contribution of our paper is that we illustrate how subjective wellbeing data can be used to provide direct tests of theoretical hypotheses about how the welfare effects of public provision should vary among different groups in the population.

First, we find some evidence for an 'ends-against-the-middle' equilibrium (Epple and Romano 1996) in the provision of public health care, where middle-income individuals prefer higher public expenditure at the margin than do low-income or high-income individuals. The intuition is that health care is a normal good, and hence demand (for public and private health care taken together) is increasing with income. However, the tax price of public health care also increases with income. Hence the preferred level of public provision can be a non-monotonic function of income when a private alternative is also available.

Second, we examine whether the welfare effect of public health care differs according to whether the individual classifies himself as being right-wing or left-wing on the political spectrum. We find that those who classify themselves as right-wing derive a lower benefit from the public provision of primary health care, but perhaps surprisingly, a higher benefit from special health care. Finally, we examine whether the welfare benefit is inversely related to an individual's belief in a just world, as suggested by Benabou and Tirole (2006). However, we find no evidence for this in our data.

Our paper is closely related to a number of studies that examine the welfare effect of public goods (or public bads) using happiness or life satisfaction data. Frey *et al.* (2004) and Frey and Stutzer (2005) assess the value of security by using regional-level data from the UK and France on terrorist attacks, together with individual-level data on life satisfaction. Similarly, Powdthavee (2005) has analysed how living in a high-crime area affects wellbeing. Van Praag and Baarsma (2005) use individual-level life satisfaction data to evaluate the welfare effects of a negative externality, namely airport noise from Amsterdam airport. Welsch (2002) has conducted a cross-country study of the effect of urban air pollution on average life satisfaction. Further, Alesina *et al.* (2004) find a negative effect of economic inequality on happiness, combining individual-level data on happiness and income inequality measured at the country level. Their study is closely related to ours, as redistribution is one of the central activities of modern governments. To our knowledge, no previous study has used local-level data on public services, together with life satisfaction data, to evaluate the welfare effects of public service provision.

The rest of the paper proceeds as follows. In Section I, we discuss the data and the empirical methodology used. The basic results of the paper are presented in Section II. Section III discusses how the results vary among different groups in the population. In Section IV, the results and their robustness are discussed in more detail. Section V concludes.

I. Data and Methodology

Empirical model of wellbeing

We use Finnish life satisfaction data from the World Values Survey (WVS) conducted in the year 2000.[2] The data are from approximately 1000 respondents.[3] We run the following ordered probit regression:

$$(1) \qquad s_{imr} = x'_{imr}\beta + z'_{mr}\gamma + \delta_r + \varepsilon_{imr},$$

where s_{imr} is the self-reported life satisfaction of individual i in municipality m in province r. We interpret self-reported life satisfaction as being an indication of the latent continuous variable, individual wellbeing, that we cannot observe directly. x_{imr} is a vector of individual-level explanatory variables, and z_{mr} is a vector of municipality-level variables, including the measures of the level of health care services that are our main focus. Further, δ_r is a province dummy, and ε_{imr} is an error term. Because some of the regressors are measured at the municipal level while observations are at the individual level, we need to account for the possibility that the error terms of individuals in the same municipality may be correlated (see Moulton 1986). We therefore use robust standard errors that allow for such correlation.

As our dependent variable, we will use individuals' answers to the World Values Survey question 'Overall, how satisfied are you with your life these days?', where individuals evaluate their satisfaction with life on a ten-point scale. The World Values Survey data also include a question where individuals are asked to evaluate how happy they are. However, happiness is much more imprecisely measured in the data than life satisfaction, as it is only evaluated on a four-point scale. Concentrating on life satisfaction data is in line with the earlier related literature on European countries (see, for example, Alesina *et al.* 2004; Di Tella *et al.* 2001, 2003).

The use of subjective wellbeing data in formal econometric analysis often raises concerns among economists. A first question is whether the statements given by respondents in life satisfaction surveys are a valid measure of individual wellbeing. There is evidence that these statements are strongly correlated with other, perhaps more objective indicators of wellbeing, such as reports by one's friends and family members, or physiological measures of stress.[4] Second, we use ordered probit in our estimations, which requires that the life satisfaction measure needs to be ordinally comparable between people: that is, we require that a person who reports a life satisfaction score of 10 is in some objective sense better off than a person who reports a score of 9.[5] One concern is whether this assumption is valid. The use of subjective wellbeing data is discussed and defended, among others, by Alesina *et al.* (2004) and Di Tella *et al.* (2003). Di Tella *et al.* (2003) point out that even if cross-person comparisons of the life satisfaction measure are not perfectly reliable, this is less of a problem in large samples and given that these variables are used only as dependent variables, as we do in the present study.

Ferrer-i-Carbonell and Frijters (2004), on the other hand, emphasize the importance of controlling for individual-level unobserved factors by using panel data methods. By estimating an individual fixed-effects model, it would be possible to control for time-

invariant unobservables correlated both with subjective wellbeing and the regressors of interest. However, individual-level panel data are rarely available.[6] With our data, it is possible to control for individual-level variables that have been controlled for in previous studies on life satisfaction where no panel data have been used. In addition, we conduct robustness checks where we control for a wide variety of additional individual-level factors that are relevant from the point of view of our research question, notably factors that may be correlated with both life satisfaction and our health care variables of interest—see Section IV for more details.

Regarding the estimation of the welfare effect of publicly provided goods, using subjective wellbeing data can circumvent some of the problems of traditional stated preference and revealed preference methods.

First, stated preference methods ask individuals to state their valuation for a given good. Such surveys are subject to problems associated with the failure of affective forecasting (individuals fail to anticipate adaptation to changing circumstances) as well as focusing illusions (exaggerating the importance of the current focus of one's attention) (Kahneman and Sugden 2005). There may also be a bias if individuals respond strategically to stated preference surveys.

If subjective wellbeing data can be regarded as a valid measure of utility, as argued above, subjective wellbeing surveys will measure experienced utility directly, therefore eliminating the need for forecasting future welfare. This type of data will also be less affected by focusing illusions, at least if the question on life satisfaction is asked at the beginning of the survey.[7] Further, it is unlikely that individuals will behave strategically in giving their life satisfaction response, as they do not know the precise purpose for which the responses will be used.

Second, revealed preference studies are based on the idea that when choosing between different bundles of public and private goods (such as housing), individuals make a trade-off, and thereby reveal information about the value they place on public goods. However, this method relies on the assumption that the relevant private goods' markets are in equilibrium. The life satisfaction approach does not require such an assumption to be made. For more detailed comparisons of the alternative methods, see, for example, van Praag and Baarsma (2005), Frey and Stutzer (2005) and Kahneman and Sugden (2005).

Measuring the level of public health care services

The respondents in our data come from approximately 140 different municipalities. The fact that the data include information on the municipality where the respondent lives makes them well-suited for our purposes, as the municipality level is the level at which public primary health care is mainly organized in Finland.[8] The municipality level is also the lowest level of local government: in the year 2000, there were 452 municipalities, and the average number of inhabitants in a municipality was approximately 12,000 (with a median of 5000).

There is universal access to public health care, and it is provided either free of charge or at a nominal user fee by the municipalities; that is, municipalities have a legal requirement to provide the health services needed by inhabitants. Therefore munici-palities cannot choose which services to provide and for whom. However, this requirement should best be seen as a determinant of service volume, not as a determinant of service quality, which is much harder to monitor. In Finland, as in the other Nordic countries, there is a high degree of fiscal decentralization, municipalities have a high degree of tax autonomy, and they face hard budget constraints (see, for example, Lotz

2006). Due to the financial autonomy of the municipalities, they have direct control over how much resource they will devote to service provision. There may therefore be considerable differences in how the services are provided and in service quality between municipalities. The Finnish system therefore provides an ideal setting to study the issues that we are interested in. For many health services (notably in primary health care and certain types of surgeries in special health care) there is also a private alternative available. Approximately two-thirds of all health care services in Finland are provided by the public sector.

In the analysis, instead of simply using data on raw expenditures to analyse the effect of public health care on welfare, we aim to construct an expenditure measure more closely related to the level of service provision or service quality. Given the institutional framework outlined above, the relevant policy instrument from the point of view of the local government is expenditures conditional on service volume. We therefore estimate the amount of excess resources spent on public health care, given the number of treatment days and visits to public health centres and hospitals. To do this, we use data on local governments' net expenditures on health services in each respondent's municipality. Total service charges paid by users, as well as revenue from selling services to other municipalities and to firms, have been netted out in order to obtain a measure of expenditures used for the benefit of inhabitants. The health care expenditure data, as well as the data on the number of treatment days and visits, have been obtained from official databases of Statistics Finland.

Using a standard method, we estimate two simple health care cost functions.[9] First, we run a population-weighted[10] regression of the log of net expenditures in primary health care on the log of treatment days in primary care and the log of visits to health centres. Second, we run a similar regression for special health care. The cost function estimates are reported in the Appendix. We take the residuals from these regressions and use them as explanatory variables in our life satisfaction estimations. The residuals measure the relative deviation of expenditures from costs that would be predicted by service volume, and are therefore a measure of 'excess' expenditures in health care. It can be noted that our results are robust to other ways of measuring excess expenditures—see Section IV for more discussion.

A priori it is of course not clear whether the residual in the estimated cost functions reflects health care quality, or whether it is pure waste or reflects a regional clustering of severe health problems. It is important to recognize, however, that we do not need to make any assumptions about this. Rather, our results will indicate whether these excess costs are beneficial for welfare. This issue is also related to earlier literature using residuals from estimated cost functions to evaluate the efficiency of health care service producers: in that literature, high residuals have been interpreted as inefficiency (see, for example, Giuffrida and Gravelle 2001). Similarly, in public discussion, high public expenditures are often interpreted as waste. We believe that, rather than being axiomatically unhelpful, extra money has potential to improve aspects of care quality such as effectiveness, patient satisfaction or access to care. These positive outcomes can be attained by using monetary resources to, for example, hire more or better-suited personnel, train personnel, buy better equipment or improve facilities.[11] Naturally, a part of our measure of excess spending is indeed likely to be due to negative factors. However, according to our results, a still greater part is welfare-improving. Neither waste nor severe health problems can provide an explanation for our results—these factors cannot possibly explain why we observe a positive association between excess expenditures in health care and life satisfaction.

Significant cost differences between public service providers are a prevalent feature in many countries. Our results provide direct information on the welfare effects of these differences in the Finnish context, but we believe that the results are likely to be relevant also for other countries, particularly for those with a similar health care system.

Control variables

We include individual-level controls similar to those used in previous studies on life satisfaction or happiness. First, we control for the effect of net income. In the World Values Survey, individuals are asked to place themselves into one of ten income groups according to annual net household income. We turn this into a continuous income measure (in thousands of euros) by using the midpoint of each income interval.[12] From our point of view it is important that the question asks about net income, and the income measure therefore takes into account taxes that are paid in order to finance public services (including health care).

Further, we control for age, gender, employment status, education, marital status, number of children and religious belief. Due to the richness of the data, we are able to include additional individual-level controls that describe whether the individual has a permanent personal relationship (for those who are not married or cohabiting with a partner), has an active social life, or has retired because of incapacity. A respondent is coded as having an active social life if she reports spending time with friends, colleagues, people from the same church or people from other organizations weekly or near weekly. Being retired and under the age of 58 is regarded as being retired because of incapacity, since retiring before this age is only possible for those with physical or mental ill-health.[13]

As for regional-level controls, we will include province dummies and control for municipality-level variables which may be correlated with the level of public services and life satisfaction.[14] Average unemployment and the log of average income (in thousands of euros) are included, as they have been found to be significant explanatory variables in previous life satisfaction studies (see, for example, Di Tella *et al.* 2003; Clark and Oswald 1994), and the macroeconomic situation may also affect a local government's budgeting decisions. Further, better public services may lead to higher house prices, which would partially offset the effect of public services on life satisfaction (see, for example, van Praag and Baarsma 2005). In order to capture the full effect of public health care, we therefore control for the level of house prices in each municipality. Further, in order to control for possible covariation of health care expenditures with expenditures in other public services that may also affect welfare, we control for per capita expenditures of the municipality in other sectors (such as social services, education and culture, and administration). The municipal-level variables have been obtained from separate registers of Statistics Finland, and they therefore include information on all residents of the municipality (and are not calculated from the WVS sample).

II. RESULTS

Our results from estimating equation (1) are presented in Table 1. Let us first turn to the results regarding the health care variables. Excess expenditures in primary health care have a significant and positive effect on the life satisfaction of the average citizen. This result is qualitatively extremely robust, and holds for different ways of measuring excess costs (see Section IV for more discussion). There are clear reasons why better public

TABLE 1
THE DETERMINANTS OF LIFE SATISFACTION[a]

Dependent variable: life satisfaction		
Personal characteristics		
Income	0.0097**	(0.0040)
Age	− 0.0422*	(0.0250)
Age squared	0.0004	(0.0003)
Male	− 0.0016	(0.0698)
Labour market status		
Self-employed	− 0.2094	(0.1349)
Retired	0.2701	(0.2383)
Housewife	0.1779	(0.1826)
Student	− 0.0454	(0.1771)
Unemployed	− 0.2858**	(0.1448)
Education		
Basic	− 0.5108***	(0.1542)
High school	− 0.4301***	(0.1528)
Vocational college	− 0.3610**	(0.1506)
Lower university degree	− 0.4855***	(0.1760)
Marital status		
Living together	− 0.2562**	(0.1094)
Widowed	− 0.4617**	(0.2011)
Divorced	− 0.4795**	(0.1927)
Separated	− 0.0144	(0.3183)
Never married	− 0.5745***	(0.1411)
Personal relationship	0.3210	(0.2022)
Number of children	0.1554***	(0.0540)
Number of children squared	− 0.0207***	(0.0060)
Religious	0.1724**	(0.0849)
Social life	0.3067***	(0.0790)
Incapacity retirement	− 0.5965**	(0.2848)
Health care variables[b]		
Excess expenditures in primary health care	0.4147**	(0.1688)
Excess expenditures in special health care	0.1157	(0.2878)
Municipality-level control variables		
Log of average income	− 0.1211	(0.5417)
Unemployment rate	− 0.9818	(1.7440)
Log of house prices	− 0.2442	(0.1916)
Log of net expenditures in soc. services & admin.	0.2043	(0.2777)
Log of net expenditures in education & culture	0.2539	(0.2111)
Net expenditures in other sectors	− 0.1916	(0.3430)

Notes
[a]Ordered probit regression including province dummies, $N = 847$. Robust standard errors have been used.
[b]See Section I for the definition of excess expenditures.
The omitted category is employed, married females with a higher university degree.
*,**,***Denote significance at the 10%, 5% and 1% levels.

health care may have a positive impact on welfare, even beyond those individuals who currently use public health care services. First, better public services are a form of social insurance. Second, there may be positive externalities from a well-functioning public health care system.[15]

However, expenditures in special health care seem not to have a significant impact on the wellbeing of the average citizen. There are a number of potential explanations for this result. First, the welfare effect of special health care varies widely across different groups in the population: in the next section, we show that the welfare effect of special health care *is* significant and positive for middle-income individuals, but is insignificant for low-income individuals. Second, and related to the first point, the quality of primary health care is likely to be a more salient issue for most individuals, as primary health care is used much more frequently by the average citizen.[16] Third, due to the organization of the Finnish health care sector, the municipalities are likely to have better control over primary health care than they do over special health care, which is provided by hospital districts owned jointly by the member municipalities. Therefore the ability of municipalities to transform additional expenditures into quality might be more limited in special health care. Finally, visits to hospitals are likely to be more heterogeneous than visits to health centres in primary health care. Therefore it may be the case that if a municipality has incurred high costs in special health care, the cases that have been treated may have been more severe or complex than elsewhere.

The coefficients of municipal level control variables turn out to be insignificant. Most of the signs, however, are in line with theoretical predictions. For example, local house prices have a negative effect on wellbeing. Also, the coefficient of average income is negative, possibly because the variable captures the effect of income of the comparison group of citizens (see also Ferrer-i-Carbonell 2005).

The coefficients on per capita expenditures in social services and administration, as well as in education and culture, are positive. On the other hand, the coefficient on per capita expenditures in other sectors is negative.[17] These additional expenditure variables simply serve as controls, and have not been scaled to take into account the amount of output produced. Therefore the result that they are not found to have an effect on welfare is in line with findings from previous literature, as explained in the Introduction. Similarly, it is interesting to note that if instead of our measure of excess expenditures we use per capita raw expenditures in health care to measure the level of health care services, these do not have a significant effect on the welfare of the average citizen. Again, this is in line with previous literature, and highlights the need to construct more specific measures of the level of service provision.

Turning next to the effect of individual-level factors on life satisfaction, our results seem to make sense and they are broadly in line with earlier literature. This gives us confidence that the utility functions that we have estimated are well defined. Personal income has a significant and positive effect. We also tried including the square of income, or using the log of income, but the squared term was not significant and the linear specification gave the best fit. The results are unaffected by the choice of functional form.

Of other personal characteristics, age appears to have a U-shaped effect on life satisfaction,[18] a result familiar from many earlier studies. Unemployed individuals are significantly less satisfied with their life than those who are employed. Further, being married or having a high education has a positive effect on wellbeing. In many studies, women have been found to be more satisfied with their lives than men, but in the Finnish data gender seems to have no effect. Having children has a highly significant but nonlinear effect on life satisfaction in our data. Religious individuals, as well as those with an active social life, are also more satisfied with their life. Finally, retirement because of incapacity has a significant and negative effect on life satisfaction.

To get a feel for the magnitude of the effect of health care expenditures on welfare, we can calculate an average individual's predicted probabilities of reporting high levels of

life satisfaction for different expenditure levels. Higher excess expenditures in primary health care increase the probabilities of reporting satisfaction levels 9 and 10 (while reducing the probabilities of all other categories). An average individual's probabilities of reporting these levels when excess expenditures equal zero (that is, when the municipality is operating exactly at the level of expenditures predicted by service volume) are 27.5% and 6.6%, respectively. An increase of one standard deviation in our excess expenditure variable increases the probabilities to 29.2% and 7.6%. Therefore an increase of one standard deviation increases the probability of reporting a 9 or a 10 by 2.7 percentage points altogether.

To further assess the significance of these effects, we calculate variations in expenditures needed to compensate for the effects of certain individual characteristics. The most obvious reference point would seem to be household income. However, it has been well documented in earlier literature that the estimates of the effect of personal income on life satisfaction are likely to be seriously biased downwards due to reasons such as endogeneity and measurement error—see, for example, Powdthavee (2007) for a detailed discussion of problems associated with measuring the marginal utility of income using data from surveys on subjective wellbeing. Due to such biases, comparing the effect of health care expenditures to the effect of income would lead to a serious overestimation of the welfare effect of health care.

Instead, we use a personal characteristic that is definitely exogenous and likely to be accurately measured, namely the respondent's age. According to our estimates, individuals are most satisfied with their lives at the age of 18, whereas satisfaction is lowest at the age of 49. To compensate for the decline in wellbeing for the 49-year-old by changing excess expenditures in primary health care, a municipality initially operating at zero excess expenditures would need to increase its expenditures 1.67-fold.[19] This is clearly remarkable in size compared to the standard deviation of relative excess expenditures (0.289), a result partly reflecting the sizeable effect of age on life satisfaction. To take another individual characteristic as an example, an expenditure increase of 69% would be needed to compensate for personal unemployment. We can conclude that, despite having statistically significant potential to promote wellbeing, expenditure increments in primary health care would need to be large to make a substantial difference compared to the effects of individual characteristics. It should be kept in mind that these results have been calculated for the average citizen. As we will see in the next section, the effects of (primary and special) health care are more pronounced for some population groups than for others.

III. EXTENSIONS

Differences between income groups

In the current and the following subsection we will analyse how the welfare effect of public health care varies among different groups in the population. First, we focus on the theoretical result of Epple and Romano (1996), that valuation for publicly provided private goods such as health care is in equilibrium likely to be highest for individuals with income just below the mean.

More specifically, Epple and Romano (1996) analyse a good that is provided by the public sector, but public provision can be supplemented by private purchases. They show that in this dual provision game, there are two types of equilibria. This occurs because there are two opposite forces present: the tax price of public health care increases with

income, and therefore individuals with highest income always prefer a zero level of public provision, as the private alternative is cheaper for them.

However, the increasing tax price is countered by the fact that health care is a normal good, so that the demand for health care (for public and private provision taken together) increases with income. The type of equilibrium that emerges depends on the relative magnitudes of the price and income elasticities (see also Kenny 1978). If the income elasticity of demand is lower than the absolute value of the price elasticity, we get the first type of equilibrium where the most preferred level of public provision is (weakly) decreasing in income. In this case the voter with the median level of income is pivotal. However, if the income elasticity is higher than the price elasticity, we get another equilibrium where the most preferred choice increases with income for those with income below the mean (whereas individuals with income above the mean still prefer a zero level of public provision due to the presence of a private alternative). In this equilibrium, therefore, the choice of government expenditure corresponds to the preferences of a lower-than-median-income household, and there is a coalition of middle-income households that prefer higher public expenditure at the margin, whereas a coalition of high- and low-income households prefer a reduction. Epple and Romano (1996) argue that preferences for health care are such that this type of an 'ends-against-the-middle' equilibrium is the more likely outcome. However, which equilibrium emerges in reality is ultimately an empirical question.

There are a number of previous studies, mainly using stated preference methods, on how the valuation for health care varies between income groups. The typical finding is that the valuation for health care or a specific treatment is monotonically increasing with income—an alternative that is not possible in the Epple and Romano model, which concerns public provision in the presence of a private alternative. This apparent inconsistency can be explained by two factors. First, many studies simply ask about respondents' valuation for health care (as in Whynes *et al.* 2003) and do not differentiate between public and private services. These studies therefore show that the demand for public and private health care services *taken together* is increasing with income, a finding that is consistent with the Epple and Romano model. Indeed, the normality of health care is one of the key assumptions behind the model. Second, studies that specifically ask subjects to state their valuation for public health care have typically not allowed for a nonlinear effect of income—see, for example, Mataria *et al.* (2004) and Pavlova *et al.* (2004). No empirical studies that we know of have directly examined the Epple and Romano model or used subjective wellbeing data to assess the welfare effect of health care.

To examine which of the two types of equilibria implied by the Epple and Romano model is consistent with our data, we construct dummies for low-income, middle-income and high income-individuals. In order to ensure consistency with the definitions of Epple and Romano, we define as middle-income individuals those individuals with income (approximately) in the second quartile.[20] We then run a regression similar to that in (1), adding terms where we interact our measures of excess expenditures in public health care with the low-income and high-income dummies. The results are presented in Table 2. We report only the coefficients on health care expenditures and the interaction terms, as the results for all the other variables are virtually identical to the basic results presented in Table 1.[21]

The significance of the coefficient on expenditures in primary health care is improved from the basic regression: this now measures the welfare effect of the public provision of primary health care for middle-income individuals. Further, it is interesting to note that also the welfare effect of special health care is significant (at the 10% level) for middle-income individuals, even though it was not significant for the average citizen. All of the

JANI-PETRI LAAMANEN

TABLE 2

DIFFERENCES IN THE WELFARE EFFECT OF PUBLIC HEALTH CARE: INCOME[a]

Dependent variable: life satisfaction		
Excess expenditures in primary health care	0.7268***	(0.2556)
Excess expenditures in special health care	0.7656*	(0.4106)
Low-income × Excess expenditures in primary health care	− 0.4477	(0.3959)
High-income × Excess expenditures in primary health care	− 1.1026*	(0.5680)
Low-income × Excess expenditures in special health care	− 0.8253*	(0.4518)
High-income × Excess expenditures in special health care	− 0.5089	(0.5273)

Notes
[a]Ordered probit regression including individual- and municipal-level controls and province dummies, $N = 847$. Robust standard errors have been used.
The omitted category is middle-income individuals.
*,***Denote significance at the 10% and 1% levels.

interaction terms are negative, and two of them are also statistically significant at the 10% level. There is therefore some evidence in our data that low-income and high-income individuals would prefer a lower level of public health care expenditures than would middle-income individuals. In particular, high-income individuals appear to derive a lower welfare benefit from primary health care, and low-income individuals from special health care, as compared with middle-income individuals. This heterogeneity can explain the fact that we did not find significant results on the welfare effect of special health care in the previous section, where the parameter on health care expenditures was constrained to be the same for all income groups.

Further, the overall welfare effect of public health care for low-income individuals is given by the sum of the expenditure coefficient and the coefficient of the relevant interaction term (similarly for high-income individuals). It can be noted that the overall effect for low-income individuals in special health care, as well as the overall effect for high-income individuals in primary health care, is negative (though the effect is not statistically significant). Further, if we test the significance of the overall effect of public health care expenditures on the welfare of the low-income group and high-income groups in those sectors where the overall effect is positive,[22] then these effects are not significantly different from zero.

Recall from Section II that individual wellbeing is increasing with income: high- or middle-income individuals are (*ceteris paribus*) better off than low-income individuals. An interesting question is whether this gap can be reduced by public provision of health care. Our results suggest that since the welfare effect of the public provision of special health care is lower for low-income individuals than for middle-income individuals, special health care does not seem to be a particularly effective instrument from the point of view of redistributing welfare. Regarding primary health care, it appears to have potential for diminishing the wellbeing gap at least between middle-income and high-income individuals.

Our results therefore give support for the 'ends-against-the-middle' equilibrium of the public provision game. Even though the evidence is not fully conclusive as some of the interaction terms are not significant, it should be noted that we find no evidence that supports the alternative equilibrium of the Epple and Romano model, where valuation for public health care should be monotonically decreasing with income.

Differences based on political ideology

It is also interesting to examine how the welfare effect of public health care varies according to individual attitudes that are likely to be related to attitudes towards the public sector. We therefore analyse whether the welfare effect of public health care varies according to whether the individual classifies himself as left-wing or right-wing on the political spectrum. Alesina *et al.* (2004) have conducted a similar analysis of the welfare effect of inequality, finding that in Europe, left-wingers are much more bothered about inequality than right-wingers. In a similar vein, we would expect left-wingers to have a higher taste for redistribution, and therefore to benefit more from public services, to the extent that those services can be regarded as a form of redistribution. Our discussion in the previous subsection suggests that this may be the case regarding primary health care, but not necessarily regarding special health care.

In the World Values Survey, individuals are asked to state whether they consider themselves to be left-wing or right-wing on a ten-point scale (where 1 = left-wing and 10 = right-wing). We split the spectrum in half and classify as 'left-wing' those reporting numbers between 1 and 5, and the rest are classified as 'right-wing'.

We examine the question of whether right-wingers derive a lower welfare benefit from public health care provision by again including the relevant interaction terms in our basic regression. The results are summarized in Table 3. The results indicate that the welfare effect of public provision of primary health care is positive and significant for left-wingers. Right-wingers, on the other hand, derive a significantly lower welfare benefit from primary health care, and the overall effect for right-wingers is not significantly different from zero (the *p*-value is 0.47). Interestingly, however, the welfare effect of special health care appears to be higher for right-wingers than for left-wingers. An explanation may be that public special health care is not as redistributive as public primary health care, as was also suggested by our analysis in the previous subsection. For many types of severe illnesses that are treated in special health care, the private market is not well developed and the insurance aspect of public provision is likely to be important.

Finally, we have examined the related question of whether the welfare effect of public health care is higher for individuals who have a belief in a just world (BJW). Benabou and Tirole (2006) have suggested that such a belief is likely to be related to attitudes towards the public sector: roughly speaking, those who believe that the poor are poor because of bad luck are likely to support a larger public sector than those who believe

TABLE 3
DIFFERENCES IN THE WELFARE EFFECT OF PUBLIC HEALTH CARE: POLITICAL IDEOLOGY[a]

Dependent variable: life satisfaction		
Excess expenditures in primary health care	0.8687**	(0.3473)
Excess expenditures in special health care	− 0.1873	(0.3594)
Right-wing	0.1168	(0.0714)
Right-wing × Excess expenditures in primary health care	− 0.7345**	(0.3613)
Right-wing × Excess expenditures in special health care	1.0615**	(0.4746)

Notes
[a]Ordered probit regression including individual- and municipal-level controls and province dummies, $N = 847$. Robust standard errors have been used.
The omitted category is individuals who classify themselves as left-wing.
**Denotes significance at the 5% level.

that the poor are poor because they are lazy. A belief in a just world may also be related to political attitudes, so that those who believe in a just world are more likely to be right-wing (though in our data the correlation between the dummies for right-wing political attitudes and BJW is only 0.14).

The survey data that we use include a question: 'Why are there people in our country who live in poverty? Please state the most important reason.' The answer categories are (1) 'they are unlucky'; (2) 'they are poor because of laziness and lack of willpower'; (3) 'they are poor because the society treats them unfairly'; (4) 'it is an inevitable part of current developments'; and (5) 'none of the above'. We have classified individuals who have given answer (2) as those who believe in a just world, and the rest as not having such a belief. We have then included the relevant interaction terms in our basic regression. However, we find no differences in the welfare effect of public health care provision between individuals who believe that poor people are lazy, and those who do not have such a belief. A possible reason is that public health care is used not only by the poor, but also by higher-income individuals: indeed, our results indicate that middle-income individuals prefer a higher level of public health care services than low-income individuals.

IV. DISCUSSION

In this section, we discuss a number of factors that might potentially cause problems for our analysis, and argue that our results are robust to these considerations.

First, one might worry that some characteristics of individuals in our sample may affect excess health care spending. This brings about a problem if these characteristics are not controlled for and are also correlated with individual life satisfaction. In our regressions, we have controlled extensively for individual characteristics that have been found important in earlier literature. We believe that the potential problem is further attenuated by the fact that characteristics are measured at the individual level, and individuals cannot affect the municipal-level health care variables directly. However, since we cannot entirely exclude the possibility of a correlation between previously uncontrolled individual characteristics and health care spending, we carry out some robustness checks. Because individual characteristics may affect spending mainly through the electoral process, the appropriate strategy is to control for additional variables related to individuals' political attitudes and specifically attitudes toward public services. We believe that this should be a highly effective strategy, as we can directly control for something that is the potential cause of a bias. Our estimates of the effect of excess expenditures would be biased upwards if individuals more satisfied with their life are prone to choose local governments more favourable to high health care spending. Conversely, if high satisfaction with life is related to a preference for low spending (for example because healthy individuals would favour lower health care expenditures), our estimates would be biased downwards.

Our data contain several variables suitable for testing the robustness of our results from this viewpoint. We have tried controlling for answers to various questions related to individual attitudes and personality traits. These include whether society should provide individuals with basic needs (including health care), whether society should reduce income inequality, how much the individual trusts the health care system, whether the individual is willing to do something to improve the living conditions of ill and disabled people, whether the individual classifies herself as being left-wing or right-wing, and which political party the individual would vote for or favours. We have included these variables both separately and simultaneously by including a dummy for each answer category.

For the most part, our results are practically unaffected by the inclusion of these additional control variables. Estimates of the coefficient of excess expenditures in primary health care range from 0.3975 to 0.4591, with p-value always below 0.02. Compared to the coefficient from the basic specification (0.4147), these changes are negligible. Further, it is interesting to note that in most cases the additional controls improve the results by increasing the size of the coefficient and, with some control variables, the significance of the coefficient improves from 5% to the 1% level. The coefficient of excess expenditures in special health care remains positive but statistically insignificant in all specifications.

When it comes to our results regarding differences between population groups, including additional control variables one at a time typically increases the significance of the main coefficient of excess expenditures in special health care, as well as the coefficient of the interaction term for low-income individuals, from 10% to the 5% level. Also, the sizes of the coefficients are increased. Thus the effect of special health care becomes more pronounced for middle-income individuals, while the effect for low-income individuals remains negative and very close to zero. Results concerning the income-group-specific effects of primary health care remain mostly unchanged. In only two cases, a statistically significant coefficient becomes insignificant. Inclusion of individual political orientation on the left–right spectrum alone makes the interaction term between high-income and primary health care statistically insignificant (p-value 0.116). Similarly, when individual political party preference is included alone, the coefficient of special health care for middle-income individuals loses its significance (p-value 0.110). However, the qualitative conclusion—that there is an ends-against-the-middle pattern in the valuation for public health care—remains unchanged in all specifications. Further, including more of the additional controls simultaneously, and thus controlling for political preferences and personality traits more carefully, restores the significance of the coefficients of interest and in some cases even improves significance.

In the light of the above discussion, the conclusions of our analysis are highly robust to the inclusion of individual-level variables that may be related to individual life satisfaction as well as to electoral outcomes.

Another concern also related to politics is that our measure of excess expenditures might actually characterize the relative strength of different groups in municipal decision-making bodies. If individuals care about the politics itself or about some other outcomes of political decision making that are correlated with health care spending, our variables may capture these effects. To check if this is the case, we have calculated variables measuring the relative number of seats in municipal councils (local parliament) held by different political parties. Also, we have calculated a Herfindahl index from these to account for the concentration of political power. To take into account the possible effect of congruence between individual party preference and the composition of the local parliament, we have constructed a dummy which equals one if the party an individual supports is the biggest group in the municipal council. Including these variables in life satisfaction regressions, together with dummies indicating individual party preference, results only in some minor changes in our qualitative results. In the regressions concerning differential effects for different income groups, the interaction term for primary care of high-income individuals loses significance (p-value 0.205). In turn, the interaction for special care of low-income individuals becomes statistically significant at the 5% (instead of 10%) level. In the analysis of differences between left-wing and right-wing individuals, adding the controls improves the significance of the effect for left-wing individuals from 5% to the 1% level.

JANI-PETRI LAAMANEN

A third potential cause of concern is residential sorting. To some extent, individuals' choices of residential location may be affected by health care spending. In this case, people who prefer high (low) spending levels would be prone to move to municipalities with high (low) expenditures. However, it is clear that the source of bias in our estimates would again be a correlation between individual attitudes and spending levels. Above, we have shown that controlling for relevant individual personality traits and opinions does not affect our key results.

Our results are also robust to different ways of measuring the level of health care service provision. We have experimented with using expenditure per treatment day or per visit as the measure of health care service provision, and our results remain unchanged. Further, when using estimated excess costs as the measure of health care service provision, our results are broadly robust to different specifications of the health care cost function. As was explained in Section I, the results that we have reported were obtained with the commonly used log-log specification for the health care cost function. We have experimented with a wide variety of other specifications: a linear cost function, a cost function that has been scaled by population of the municipality, and a cost function where population has been included as a separate control variable. We have also tried leaving out the capital city of Helsinki from the analysis, as it can be expected to differ from other towns in a number of ways that may be relevant. Our results are broadly robust to all of these changes: in particular, the p-value of the primary health care expenditure coefficient is less than 0.06 in all the specifications that we have tried. Also, most of the results regarding differences between population groups are qualitatively robust to these changes. One difference is that for some specifications, the valuation for special health care appears to be monotonically increasing with income (though this pattern is never statistically significant). This can again be explained by the fact that the private market for services in special health care is not very well developed.[23]

Finally, it should be noted that treatment days and the number of visits to public health centres and hospitals are potentially endogenous to the health care expenditures of the municipality. To the extent that lower expenditures lead to lower quality, this might induce individuals to shift towards using private sector health services, which leads to an increase in expenditures per treatment day or visit. This causes a downward bias in our estimate for the welfare effect of health care expenditures, and our estimates can therefore be regarded as the lower bounds for these effects.

V. Conclusions

We have examined the impact of publicly provided health care services on individuals' subjective wellbeing. Our findings suggest that higher relative expenditures in local service provision have a positive influence on the life satisfaction of the average citizen. Further, we have examined how this effect varies among different population groups. Our results provide support for the 'ends-against-the-middle' hypothesis (Epple and Romano 1996), that individuals' preferred level of public provision of private goods such as health care may be non-monotonic with income. Further, welfare effects of health care spending depend on individual political orientation.

Taken together, our results show that devoting more resources to the provision of public services has a positive effect on wellbeing, although this effect may vary between population groups. This conclusion may seem to be in contrast with some earlier studies, which have not been able to find a positive effect of government expenditures on welfare.

However, we have argued that when examining whether a more extensive welfare state increases wellbeing, it is important to use a measure more closely related to specific aspects of service provision: it is likely that individual welfare is affected primarily by the level of services financed by government expenditures, and not by government expenditures *per se*. In addition, it is important to note that if the welfare effects of public services vary between population groups, this may hinder finding positive and significant effects if we only examine the population taken as a single group.

Our results also have interesting implications for the popular debate on public provision of private goods, particularly the issue of whether high expenditures in the public sector are necessarily a sign of inefficiency. Our results indicate that (at least a part of) relatively high costs in public health care are welfare-improving, and therefore cannot be pure waste; rather, high costs are likely to be at least partly reflected in better quality of service. An interesting issue for further research is to try to disentangle these effects, for example by finding suitable instruments for health care quality. The method used in our study could also be useful for determining the welfare effects of other publicly provided services.

APPENDIX

The estimated cost function for the public provision of primary health care is reported in Table A1. We have regressed the log of net expenditures in primary health care on the log of visits and treatment days. As we have separate information for visits in dental care (which is organized as part of primary health care in Finland), we include this as a separate variable in order not to place unnecessary restrictions on the parameters of the cost function.

Similarly, the estimated cost function for special health care is reported in Table A2.

TABLE A1
COST FUNCTION: PRIMARY HEALTH CARE[a]

Dependent variable: log of net expenditures (in 1000 euros)		
Log of visits	0.5924***	(0.1066)
Log of visits (dental care)	0.2581***	(0.0946)
Log of treatment days	0.2189***	(0.0509)
Constant	− 2.7622***	(0.2596)

Notes
[a]Population-weighted OLS, $N = 153$. Robust standard errors have been used.
***Denotes significance at the 1% level.

TABLE A2
COST FUNCTION: SPECIAL HEALTH CARE[a]

Dependent variable: log of net expenditures (in 1000 euros)		
Log of visits	0.6998***	(0.1167)
Log of treatment days	0.2912**	(0.1226)
Constant	− 0.6212***	(0.2112)

Notes
[a]Population-weighted OLS, $N = 153$. Robust standard errors have been used.
,*Denote significance at the 5% and 1% levels.

© The London School of Economics and Political Science 2008

NOTES

1. The terms life satisfaction and wellbeing are used interchangeably in this paper.
2. Even though the WVS was also conducted in 1995 and 2005, we are unfortunately only able to utilize data for the year 2000. Data from other years is hard to combine with our data, as some key variables (such as personal income) are measured differently in the different waves of the survey. The year 2000 is the only year with a measure of net income. Further, accounting practices of Finnish municipalities underwent a major change in 1996 and therefore the cost variables that are central to our analysis would not be comparable between the year 1995 and the latter years.
3. In the analysis, we include respondents in economically active age (15–74). The elderly are likely to be treated within elderly care, which is organized separately from public health care.
4. See, for example, Di Tella and MacCulloch (2006) and Konow and Earley (2008) for a more detailed discussion and further references, as well as Kahneman and Sugden (2005) for a critique of subjective wellbeing data as a measure of true utility.
5. However, we do not assume cardinality of the life satisfaction measure, which would be required if ordinary least squares (OLS) were used in estimation.
6. For related studies where individual-level panel data has been used, see Winkelmann and Winkelmann (1998) on life satisfaction (using the German Socio-Economic Panel (GSOEP)), Hamermesh (2001) on job satisfaction (using the GSOEP and the US National Longitudinal Survey of Youth), and Clark et al. (2005) on satisfaction with one's financial situation (using the European Community Household Panel).
7. This is the case in our data, as the life satisfaction question appears as one of the first questions in the survey that we use (as question number 10 out of a total of 114 questions).
8. Some small municipalities provide primary health care jointly with other municipalities. In special health care, services are provided by hospital districts, which are formed on average from 20 municipalities.
9. Our procedure for estimating the health care cost function corresponds to the deterministic cost frontier method, which has commonly been used for evaluating performance in industries such as health care—see, for example, Giuffrida and Gravelle (2001) for a discussion on this and other related methods.
10. We use population weighting in order to account for the greater random variation of costs in small municipalities.
11. Exploring which specific aspects of care are welfare-improving is beyond the scope of this paper, but some notes can be made concerning the process through which spending is transformed into quality. In a simple probit model of the probability of being retired due to incapacity, excess expenditure variables for both primary and special health care enter negatively (p-values 0.082 and 0.001, respectively). Controlling for age, age squared, gender and province of residence weakens significance (p-values 0.122 and 0.069). Seventy-seven observations are dropped since there are no respondents retired due to incapacity in some provinces (these provinces have higher excess spending than the others, on average). Together with the fact that retirees, due to incapacity, are more likely to induce higher than lower costs, which would bring about a positive correlation, these results indicate that excess spending improves health outcomes of the inhabitants. Although relevant Finnish evidence is rare, some studies suggest that patient satisfaction can be increased by means that require monetary resources (see, for example, Aalto et al. 2008 and Hietanen et al. 2007).
12. In order not to lose the observations in the highest income group, which is open-ended, we use the lower bound for this group together with a dummy that eliminates the effects of the measurement error created. All our results are thus unaffected by using the lower bound (rather than some higher number) as the income level for this group.
13. This is an attempt to partially control for personal health status. Unfortunately, we do not have direct information on the respondents' health status.
14. Finnish municipalities are divided into 20 different provinces (out of which our sample includes 19). Our specification therefore assumes that the municipalities within each province do not differ in any relevant

respects other than the municipal-level factors that we include as controls. This is, in our view, a reasonable assumption, given the rich array of controls used and the small size of Finnish provinces.

15. For a recent contribution on the role of the public sector in providing insurance, see Pirttilä and Tuomala (2007). For a discussion of externalities, insurance and other rationales for public sector involvement in the market, see, for example, Atkinson and Stiglitz (1980, pp. 5–8).

16. The average per capita number of visits to health centres (primary health care) in our sample is 6, whereas the average citizen visited a hospital (special health care) only once in the year 2000. Also, the average number of days spent in treatment in primary health care is 1.5 times as high as the number of days spent in hospitals.

17. This variable contains net costs from other sectors besides those that have been included separately. Some of the services in these sectors are sold instead of offered free of charge (e.g. water supply or waste management) and this variable can therefore be negative (and has thus not been included in logarithmic form like the other expenditure variables).

18. The significance of the linear term is improved by including the squared term, and the terms are jointly significant (p-value 0.08).

19. Recall that our estimated cost functions are in logarithmic form, and excess expenditures therefore measure the relative deviation of expenditures from those predicted by service volume. The welfare effects reported here therefore relate to relative expenditure changes.

20. Individuals in our data are divided into ten groups according to income. We define groups 1–3 as low-income, 4–6 as middle-income, and 7–10 as high-income. Our qualitative results are robust to other sensible divisions as well.

21. We have not included dummies for the low-income and high-income groups, as the effect of personal income is already controlled for. If the dummies are included in addition to the linear income variable, this does not affect the statistical significance of the interaction terms reported in Table 2. The dummies themselves are not statistically significant.

22. That is, we test the linear restriction that the sum of the coefficient on primary care expenditures and the coefficient on the interaction term for the low-income group is zero, and similarly for the high-income group in special health care. The p-value for the former test is 0.35 and for the latter 0.59.

23. The reasons for this are likely to originate on the supply side of the market and are therefore exogenous to the Epple and Romano model, which explains public provision of private goods by demand factors only.

REFERENCES

AALTO, P., KARHE, L., KOIVISTO, A.-M. and VÄLIMÄKI, M. (2008). The connection between personnel resources with work loading and patient satisfaction on in-patient wards. *Journal of Nursing Management*, **15**, 187–206.

ALESINA, A., DI TELLA, R. and MACCULLOCH, R. (2004). Inequality and happiness: are Europeans and Americans different? *Journal of Public Economics*, **88**, 2009–42.

ATKINSON, A. B. and STIGLITZ, J. E. (1980). *Lectures on Public Economics*. Singapore: McGraw-Hill.

BENABOU, R. and TIROLE, J. (2006). Belief in a just world and redistributive politics. *Quarterly Journal of Economics*, **121**, 699–746.

BJORNSKOV, C., DREHER, A. and FISCHER, J. A. V. (2007). The bigger the better? Evidence of the effect of government size on life satisfaction around the world. *Public Choice*, **130**, 267–92.

CLARK, A. E. and OSWALD, A. J. (1994). Unhappiness and unemployment. *Economic Journal*, **104**, 648–59.

——, ——, ETILÉ, F., POSTEL-VINAY, F., SENIK, C. and VAN DER STRAETEN, K. (2005). Heterogeneity in reported well-being: evidence from twelve European countries. *Economic Journal*, **115**, C118–C132.

DI TELLA, R. and MACCULLOCH, R. J. (2006). Some uses of happiness data in economics. *Journal of Economic Perspectives*, **20**, 25–46.

——, —— and OSWALD, A. J. (2003). The macroeconomics of happiness. *Review of Economics and Statistics*, **85**, 809–27.

——, —— and —— (2001). Preferences over inflation and unemployment: evidence from surveys of happiness. *American Economic Review*, **91**, 335–41.

EPPLE, D. and ROMANO, R. E. (1996). Public provision of private goods. *Journal of Political Economy*, **104**, 57–84.

FERRER-I-CARBONELL, A. (2005). Income and well-being: an empirical analysis of the comparison income effect. *Journal of Public Economics*, **89**, 997–1019.

—— and FRIJTERS, P. (2004). How important is methodology for the estimates of the determinants of happiness? *Economic Journal*, **114**, 641–59.

FREY, B. S., LUECHINGER, S. and STUTZER, A. (2004). Valuing public goods: the life satisfaction approach. CESIfo Working Paper no. 1138.

JANI-PETRI LAAMANEN

FREY, B. S. and STUTZER, A. (2005). Happiness research: state and prospects. *Review of Social Economy*, **62**, 207–28.

GIUFFRIDA, A. and GRAVELLE, H. (2001). Measuring performance in primary care: econometric analysis and DEA. *Applied Economics*, **33**, 163–75.

HAMERMESH, D. S. (2001). The changing distribution of job satisfaction. *Journal of Human Resources*, **36**, 1–30.

HIETANEN, P. S., ARO, A. R., HOLLI, K. A., SCHRECK, M., PEURA, A. and JOENSUU, H. T. (2007). A short communication course for physicians improves the quality of patient information in a clinical trial. *Acta Oncologica*, **46**, 42–8.

KAHNEMAN, D. and SUGDEN, R. (2005). Experienced utility as a standard of policy evaluation. *Environmental and Resource Economics*, **32**, 161–81.

KENNY, L. W. (1978). The collective allocation of commodities in a democratic society: a generalization. *Public Choice*, **33**, 117–201.

KONOW, J. and EARLEY, J. (2008). The hedonistic paradox: is homo economicus happier? *Journal of Public Economics*, **92**, 1–33.

LAYARD, R. (2006). Happiness and public policy: a challenge to the profession. *Economic Journal*, **116**, C24–C33.

LOTZ, J. (2006). Local government organization and finance: Nordic countries. In A. Shah (ed.), *Local Governance in Industrial Countries*. Washington, DC: World Bank.

MATARIA, A., DONALDSON, C., LUCHINI, S. and MOATTI, J.-P. (2004). A stated preference approach to assessing health care quality improvements in Palestine: from theoretical validity to policy implications. *Journal of Health Economics*, **23**, 1285–311.

MOULTON, B. R. (1986). Random group effects and the precision of regression estimates. *Journal of Econometrics*, **32**, 385–97.

PAVLOVA, M., GROOT, W. and VAN MERODE, G. G. (2004). Willingness and ability of Bulgarian consumers to pay for improved public health care services. *Applied Economics*, **36**, 1117–30.

PIRTTILÄ, J. and TUOMALA, M. (2007). Labour income uncertainty, taxation and public good provision. *Economic Journal*, **117**, 567–82.

POWDTHAVEE, N. (2005). Unhappiness and crime: evidence from South Africa. *Economica*, **72**, 531–47.

——— (2007). How much does money really matter? Instrumenting for income in life satisfaction equations. Mimeo, University of London.

VAN PRAAG, B. M. S. and BAARSMA, B. E. (2005). Using happiness surveys to value intangibles: the case of airport noise. *Economic Journal*, **115**, 224–46.

VEENHOVEN, R. (2000). Well-being in the welfare state: level not higher, distribution not more equitable. *Journal of Comparative Policy Analysis*, **2**, 91–125.

WELSCH, H. (2002). Preferences over prosperity and pollution: environmental valuation based on happiness surveys. *Kyklos*, **55**, 473–94.

WHYNES, D. K., FREW, E. and WOLSTENHOLME, J. L. (2003). A comparison of two methods for eliciting contingent valuations of colorectal cancer screening. *Journal of Health Economics*, **22**, 555–74.

WINKELMANN, L. and WINKELMANN, R. (1998). Why are the unemployed so unhappy? Evidence from panel data. *Economica*, **65**, 1–15.

WORLD VALUES SURVEY (2000). Finnish data. Version 4. Gallup Finland & The Church Research Institute (conductor of survey). Finnish Social Science Data Archive (distributor), Tampere, 2003.

Chapter 5
Worker Turnover, Structural Change, and Interregional Migration: Evidence from Finland[*]

Jani-Petri Laamanen[†]

Abstract

We study the effects of regional labour market conditions on interregional migration using province-level panel data on bilateral migration flows and disaggregated labour market flows. Our results indicate that hires from unemployment and job separations leading to unemployment have sizeable effects on migration. The effects of hires from and separations to other labour market states, while statistically significant, appear smaller. Further, our results suggest that interindustry and interfirm shifts in employment are immaterial for migration. Taken together, interregional migration is largely affected by regional differences in unemployment and the employment opportunities available for unemployed workers.

JEL codes: J61, J63

†   FI-33014 University of Tampere, Finland; E-mail: jani-petri.laamanen@uta.fi; Phone: +358 50 318 5993, Fax: +358 3 2599 008

5.1   Introduction

The determinants of interregional migration have been the subject of numerous empirical studies in recent decades. Regardless of whether micro-level or macro-level data are used, measures related to the labour market are routinely included as explanatory variables in migration models. Despite the improved availability of detailed regional-level labour market data and contrary to various theoretical arguments, the majority of these studies use the regional unemployment rate as the key labour market variable. Arguing that other aspects of local labour markets may also be of relevance to potential migrants, some authors have used additional variables to explain interregional migration flows. Examples include Carlsen *et al.* (2006), who include measures of labour market tightness, and Hämäläinen and Böckerman (2004), who include excess job reallocation and churning variables. Exceptions of this kind are still rare, so that the the relationship between local labour market conditions and migration remains unclear. However, it can be concluded that the additional aspects considered in the aforementioned studies are important.

This article aims at contributing to the literature by considering new variables to characterize local labour markets and to explain interregional migration flows in Finland from 1988 to 1996. Our data are well suited for studying the determinants of migration for various reasons. Our data are on bilateral migration flows so that we are able to control both the source region and destination region characteristics. Further, we observe gross rather than net migration flows. Gross flows are much larger than net flows; therefore, our migration measure does not exclude a large part of the moves, like many of the earlier studies do. Our key explanatory variables come from a linked employer—employee data set with a sufficient degree of disaggregation. Using disaggregated job and worker flow data, we construct measures closely related to those local labour market phenomena, which, according to search-theoretic considerations and evidence from previous empirical studies have potential to influence migration decisions but have not been used previously. More specifically, we explore the roles of worker turnover (hires and separations) and change in the structure of employment as well as 'churning' as determinants of migration. Importantly, hires and separations are disaggregated by the associated labour market transition of the worker. This enables us to explore the potentially heterogeneous effects of hires from and separations into unemployment and other hires and separations. Industry-level information on job and worker flows allows us to measure structural change by job reallocation across industries and across establishments. By using panel data on bilateral gross migration flows between provinces, we are able to simultaneously identify the effects of source and destination regions' characteristics on migration between the two regions. We address the potential endogeneity of our labour market variables by using GMM estimation.

Our results show that hires and job separations in regions have sizeable effects on migration. Hires by local employers hinder out-migration and increase in-migration whereas job separations increase out-migration and decrease in-migration. The effects of these labour market flows are found to be strongly heterogeneous with respect to the source labour market status of the hired workers and destination status of the separated workers. Specifically, our results reveal that a hire from or a separation into unemployment has, in general, a larger effect on migration than a hire from or a separation into employment or outside the labour force. These novel results are both in line with previous micro-level evidence on individual migration propensities and in accordance with theoretical predictions on the link between local labour markets and mobility. We also find that simultaneous hires and separations have a negative effect on in-migration, possibly reflecting increased competition for jobs. Even though we observe a lot of changes in the structure of employment, that is, employment shifts between industries and between establishments within regions, we find no evidence that this affects interregional migration.

The article is organized as follows. In Section 5.2, we briefly discuss the background and sketch the theoretical underpinnings of the article. Section 5.3 describes the data and defines the variables. In Section 5.4, we introduce the methodological strategy. Section 5.5 presents the empirical results and discussion, and Section 5.6 concludes.

5.2   Background and Theoretical Considerations

The theoretical treatment of interregional migration dates back to the human capital framework by Sjaastad (1962). A similar line of thinking is adopted in the classic two-region model of migration by Harris and Todaro (1970). The central feature of the framework is that by moving or staying, individuals maximize the expected return on their human capital. The expected return in a region depends on local labour market opportunities, consisting of the probability of being employed and the wage level. The human capital approach has since been widely used in theoretical and empirical studies relating local labour markets and migration.

Following the example of the theoretical Harris-Todaro model, most empirical migration studies have used the unemployment rate and wage as the variables characterizing local labour market conditions. In some papers, the net growth of employment has also been used. However, the level of unemployment and the net change in employment hide a considerable amount of dynamics in the labour market; see Davis *et al.* (2006). Generally speaking, a lot of simultaneous creation and destruction of jobs, or hiring and separation of workers, is observed. Related to migration, Fields (1976, 1979) already has pointed out that the unemployment rate and the net employment growth are not necessarily sufficient to capture the features of labour markets that are relevant to individuals who make migration decisions. This means that they fail

to correctly measure the labour market opportunities for potential migrants, which could explain the mixed empirical results on the effect of the unemployment rate in the literature; see e.g. Furceri (2006), Hatton and Tani (2005), McCormick and Wahba (2005) and Parikh and Leuvensteijn (2003). Further, the unemployment rate not only measures labour demand but may also capture individual migration propensities, a point made explicit by Pissarides and Wadsworth (1989). Indeed, studies using micro-level data have shown that personal unemployment increases the propensity to migrate; see e.g. Antolin and Bover (1997), Böheim and Taylor (2002) and Nivalainen (2004). Jackman and Savouri (1992) point out that this is because unemployed workers are more active in job search. This may partly explain the finding of some studies using regional-level data (e.g. Jackman and Savouri, 1992; Etzo, 2011) that regions with high unemployment experience greater out-migration.

Because of the problems associated with using the unemployment rate as the only labour market variable (in addition to the wage level) in the migration equation, some studies have adopted different theoretical approaches and used alternative explanatory variables. In particular, job search and matching models have proved useful in analysing migration and in finding variables to characterize local labour markets in empirical analyses. Jackman and Savouri (1992) present a theoretical matching model, in which some workers search for and are matched to jobs that are not located in their home region. Inasmuch as these matches lead to interregional mobility of workers, migrations may be seen as outcomes of successful interregional job searches. This is in contrast to the Harris-Todaro model where workers cannot search for jobs that are located outside their home region. Clearly, however, distant search should be allowed in a realistic spatial model of modern labour markets (for a theoretical model of optimal search and mobility, see Molho, 2001). Based on their theoretical model, key explanatory variables of bilateral migration flows in Jackman and Savouri's (1992) analysis are the vacancy and unemployment rates. These are included for both the source and destination region of each migration flow. Empirical results for the effects of these variables are in line with the matching model: regions with relatively high vacancy rates and low unemployment rates experience less out-migration and more in-migration. The assumption of interregional job search that gives rise to interregional worker mobility has also been made in some studies estimating regional matching functions (for an overview of this literature, see Petrongolo and Pissarides, 2001).

Some other empirical studies, while adopting the human capital model as the theoretical framework, include alternative labour market indicators in a model of interregional migration flows. Carlsen *et al.* (2006) measure local labour market tightness by the rate at which local unemployed workers aged 25–59 years exit from unemployment. The authors find a statistically significant positive effect on in-migration to Norwegian counties for this variable. Hämäläinen and Böckerman (2004) do not attempt to measure local job-finding probabilities directly, but they suggest that regions

with more internal reorganisation in the labour market are more attractive to workers. They use labour market flow data to calculate measures of excess job reallocation and churning. These variables are shown to be positively associated with net in-migration to Finnish regions.

All the studies mentioned above, which have used alternative labour market variables in the migration equation, have focused on the dynamic features of labour markets. It appears that textit{changes} in local labour markets are intimately related to migration flows. However, there is no prior evidence on the effects of hires and job separations on migration.[1] It seems that the use of any of the aforementioned alternative explanatory variables (including those used in this article) could have been justified by either the human capital model or by search theory. However, as discussed earlier, search theory is apparently the more realistic framework to study interregional migration. Thus, we discuss our hypotheses and interpret our results in the light of search and matching.

We assume that individuals are also able to search for employment outside their home region. With some probability, an interregional match of a worker and job leads to migration of the worker (workers may also commute). Therefore, positive changes in demand for labour by firms in a region result in in-migration through hires of workers from other regions. In turn, out-migration is hindered by demand increases because, as a result, more local job seekers are likely to find a local match. Carlsen *et al*. (2006) have also stressed the importance of hiring. However, our perspective is different from theirs in two respects. Firstly, we assume that individuals first search an area and then move rather than move to be able to search in another region. It follows that the overall amount of hiring in local firms rather than merely the hiring of local workers matters. Secondly, we take into account on-the-job search and job search by those outside the labour force. Whereas many empirical migration studies focus on the effect of employment opportunities for the unemployed, we believe that moves related to job changes should not be ruled out *a priori*.[2] This implies that migration may be expected to depend on all hires rather than hires of unemployed workers only. It should be noted that the difference between the number of local hires and the total number of hires is likely to be small, whereas including hires of employed workers (and those outside the labour force) in addition to hires of unemployed workers is more consequential due to the large number of job-to-job transitions.

Whether the worker will be recruited from the same region or from another region is likely to depend on the characteristics of the job. In particular, jobs in different sectors may attract in-migrant workers to a different extent. Especially when hires take place in sectors with a tight local labour market situation, firms may resort to recruiting workers from other regions. Further, Anderson and Burgess (2000) have

---

1   Exceptions are two early papers by Fields (1976, 1979). He finds that hires and separations are highly correlated with interregional net migration flows.

2   Van Ommeren *et al*. (1999) have discussed these moves in the light of search theory.

argued that employers may have a preference for employed applicants and that the source of a new hire is important. This means that jobs taken by unemployed workers are different from jobs taken by job switchers, or that hires from unemployment and hires from employment reflect different conditions in the labour market.[3] Also, because the propensity to migrate depends on a worker's labour market status, as noted earlier, hires from unemployment and hires from employment may have differential effects on interregional migration. Recognizing the potentially different effects on migration of different kinds of jobs created gives rise to our key hypothesis, which is that the labour market indicators used in an empirical migration model should be disaggregated to take into account the potential heterogeneity of effects.

The extent to which hires by firms lead to recruitment of workers from another region depends on search activity by both local workers and by workers from the other region. Relatively active search by local workers increases the likelihood of local matches, whereas active search by the other region's workers increases the likelihood of interregional matches. It can thus be hypothesized that fierce competition for jobs locally or, put differently, an active local pool of job seekers decreases in-migration and increases out-migration. The latter effect arises because the local market is less tight from the perspective of local workers and because active local job seekers are likely to search actively outside their home regions as well. In a basic matching function, the number of unemployed is used as a proxy for the number of workers who are potentially matched to available jobs. Jackman and Savouri (1992) point out that unemployment duration may affect individual search effort, and they include the share of long-term unemployed as an additional regressor in their migration equation. We share the view that local unemployment alone is an insufficient measure of the competitive situation in a region. It is clear that the local competition for jobs is most strongly increased by negative labour demand changes that lead to job separations. Newly laid-off workers are likely to search relatively actively and thus decrease the opportunities available for in-migrant workers. In turn, the theoretical effect of voluntary separations (quits) on migration is more ambiguous because quits may be associated with job changes (due to labour supply or demand changes) or, for example, retirement. Altogether, separations are likely to change the competitive situation in the local labour market and thus are potentially an important factor in explaining interregional matches and migration.

As in the case of hires, the relevance of separations may differ according to the sector and the resulting labour market transition. In particular, separations in sectors where workers are concurrently hired by other firms may lead to job-to-job changes rather than to hiring of workers from outside the region. In turn, sectoral shifts in employment are likely to be associated with in-migration because local newly separated workers may not

---

3   Yashiv (2008) surveys studies of US labour market dynamics. The author makes a clear distinction between flows with different sources of hires and destinations of job separations and reports that the flows differ in the way they behave over the business cycle.

be qualified for the available job opportunities. Robson (2009), who studies empirically the effect of structural change on regional labour market performance, lists interregional migration as one of the adjustment mechanisms to sectoral shifts in employment. In addition to sector, the resulting labour market status of a separating worker is also likely to matter. Separations that lead to unemployment are often due to lay-offs and thus increase local competition for jobs, whereas separations leading to other labour market states (employment, out of the labour force) may not be that influential.

In our empirical analysis that follows, we estimate a model of interregional migration. The key explanatory variables reflect the labour market changes (hires and separations) that are potentially important in the light of the above discussion. In particular, we are interested in differential effects of different kinds of hires and job separations. The importance of the source labour market status of a hire and the destination of a separation is tested. Further, we test the hypothesis that local changes in the structure of employment (i.e., employment shifts between sectors and firms) lead to increases in local skills mismatch that induce work-related in-migration and stimulate out-migration.

The above discussion suggests that the labour market conditions of both the source region and the destination region affect migration through interregional search behaviour. With our data, we are able to simultaneously include variables concerning the source region and the destination region of each migration flow, which will add accuracy to our results. Further, we observe the gross migration flows between regions, which are much larger than the net migration figures used by many other studies.

5.3   Data and Empirical Specification

To model interregional mobility, we use data on gross bilateral migration flows of people of economically active ages (15–74)[4] between 19 Finnish provinces in the years 1988 to 1996.[5,6] This data provided by Statistics Finland covers all registered moves, that is, every change in the registered place of living of Finnish inhabitants. Because, according to the law, every Finnish inhabitant is obliged to have a registered address, and inhabitants are only eligible for the public services in their home municipality, the data are of high accuracy and likely to capture virtually all residential moves. Altogether, the data consist of 342 units of observation, which are the province pairs. In the analysis, we are able to use a total of 2,736 observations.

---

4   Individuals aged 15–74 years are the population for which the official employment statistics were calculated in our sample period. We use this age group to have measures on the same subpopulation in both sides of the regression equation.

5   For descriptive statistics of the variables in the data, see the appendix.

6   Due to the special character of the region and due to lack of some data, we exclude the autonomous island of Åland from the analysis.

The key set of explanatory variables is obtained from a linked employer—employee data set.[7] The data include province-level gross rates of job creation and destruction as defined in Davis *et al.* (1996) and worker flows into and from employment. We can separate between the flows from and into unemployment and the flows from and to other labour market states. Further, flow data are disaggregated by industry, which allows us to measure the degree of employment shifting between industries and between establishments within the industries. The combination of job and worker flows also allows us to calculate a variable reflecting excess worker turnover (churning).

The empirical migration equation to be estimated is

$$(1) \qquad \ln\left(\frac{M_{ijt}}{Pop_{it}}\right) = \alpha_{ij} + x'_{it}\beta + x'_{jt}\gamma + \delta_t + \epsilon_{ijt}$$

where M_{ijt} is the number of migrants aged 15–74 years from source to destination province, Pop_{it} is the population (in thousands) aged 15–74 years in the source province, and x' is a vector of explanatory variables. The indices i, j, and t stand for source province, destination province, and year, respectively. Parameters of the model are α, the fixed effect of a particular pair of provinces; β and γ, the coefficient vectors and δ_t, the fixed effect for year t. ε is the error term.

The elements of the vector x' that are of interest in this article are the variables characterizing local labour markets of the source and destination provinces. In addition to the conventional labour market variables such as the unemployment rate, the wage rate, and the net change in the number of jobs, we calculate variables that describe the dynamics of the labour markets in two aforementioned separate dimensions. In the following detailed introduction of our variables, we leave out the indices to simplify the expressions. All the variables are calculated for each of the provinces and for every year. Time-invariant factors of provinces and province-pairs such as geographical distance are captured by the fixed effects α.

Our regional labour market flow data cover business sector establishments, excluding farming, public sector, and social and personal services. We will use the following labour market flow variables in our baseline specifications: hires (H); hires from unemployment (HU); hires from other sources, that is, from employment and outside the labour force ($HOTH$); separations (S); separations into unemployment (SU); separations into other destinations ($SOTH$); and industry-specific hires and separations (H_k and S_k, where k denotes the industry).[8] It is interesting to note at this point that a modest negative correlation exists between hires from unemployment (HU)

---

7   We would like to thank Petri Böckerman for making this data available. For a description of the data, see Böckerman and Maliranta (2001). An earlier paper using variables from the same employer—employee data to study migration is Hämäläinen and Böckerman (2004).

8   The definitions of the variables can be found in the appendix.

and hires from other sources (*HOTH*), as well as between the separation variables *SU* and *SOTH*.[9] This is in line with the idea that we will get measures of rather diverse labour market phenomena by disaggregating the flows by source and destination. Because our hire and separation variables exclude some industries, we also control for the total net change in employment in these industries (*NETO*). The flows, as well as most other variables (those representing absolute numbers), are divided by population (aged 15–74 years) in the province.

As in the majority of earlier studies, we include the unemployment rate (*U*) and the wage level (*WAGE*) (in 1,000 euro) in our models. The wage level is calculated as total wage income divided by the number of wage earners in the province. We control for the industrial composition of jobs in provinces and the differences in industrial composition between the source region and the destination region. For this, we include the shares of jobs in agricultural (JOB_{pri}), industrial (JOB_{ind}), and construction (JOB_{constr}) sectors in the source and destination region. Industrial dissimilarities between the two regions are captured by an index parallel to the 'comparability index' in Jackman and Savouri (1992). More specifically, the dissimilarity index (*DIS*) is calculated as the sum of squared differences (between the source and the destination) in the shares of 14 industries.

To sufficiently control for other factors possibly affecting migration, we include an extensive set of province-level variables. For the source province, we control for demographic characteristics of the population. These controls include the number of population aged 15–74 years with different education levels (five categories), the number of inhabitants in different age groups (six categories), the number of children, the number of elderly people (above 74 years), and the number of retired people aged 15–74 years, all divided by the total population aged 15–74 years. For both source and destination province, we control for the share of owner-occupied housing, the number of newly enrolled university students (divided by population aged 15–74 years), the share of population living in municipalities classified as urban areas, and the share of population living in municipalities classified as densely populated areas. Average house prices and rents (per square metre) are also included for both provinces to account for differences in living costs. To account for the possible substitutive role of commuting in migration decision, we include the ratio of the number of employed inhabitants to the number of jobs, which measures the net interregional commuting of workers. In all regressions, we also include year dummies.[10]

---

9   The correlation coefficients are -0.27 (*HU* and *HOTH*) and -0.46 (*SU* and *SOTH*).

10   For the sake of brevity, the results concerning these additional controls are not reported in the tables but are briefly commented in the text whenever noteworthy.

JANI-PETRI LAAMANEN

5.4   The Estimation Method

In some of the more recent studies of interregional migratory flows, it has been noted that not only labour market conditions affect migration, but the reverse may also be true; see e.g. Furceri (2006) and Hämäläinen and Böckerman (2004). Theoretically, this may occur, for example, if a positive exogenous in-migration shock increases labour supply in the region. This could lead to an increase in the number of jobs, a reduction or an increase in the unemployment rate, and a change in the wage rate. For this reason, possible endogeneity of regressors needs to be taken into account in the statistical analysis. Further complication arises from the possible dynamic nature of migration. In some earlier studies, past migration flows are found significant in explaining subsequent flows. In a fixed-effects panel setting, including the lagged dependent variable is likely to bias the estimates. We solve these problems by exploiting the panel nature of our data and using the dynamic panel data GMM method by Arellano and Bond (1991) in estimating equation 1.[11] In the method, the equation is first-differenced and then estimated using the generalized method of moments (GMM). Arellano and Bond (1991) have shown that in a panel context two or more periods lagged values of the endogenous regressors and the dependent variable can be used as instruments. Regressors that are predetermined rather than endogenous need to be lagged one period. To keep the instrument matrix reasonable in size, our instrumenting strategy is to use only the two-period lagged values of endogenous regressors and only one-period lagged values of predetermined regressors.[12] We further restrict our set of instruments to include only the twice-lagged dependent variable and the lagged labour market variables, that is, the hiring and separation variables (lagged twice), net change in employment in other industries (lagged twice) and the unemployment rate (lagged once). This leads the number of instruments in our models to range from 68 to 124, which we consider sufficiently large but not excessive. It should be noted that all the variables included in the instrument matrix serve as instruments for all the regressors. Instead of relying only on lagged regressors as instruments, we have added some 'genuine' instruments. These are investments divided by gross domestic product and exports divided by turnover in the firms located in the province. These variables are one-period lagged and included for both the source region and destination region. We believe that these instruments are valid in the sense that they are likely to affect the local labour market events but are not linked to out-migration or in-migration directly.

Using the described method to estimate an equation in first differences has been shown to have some potential weaknesses. Arellano and Bover (1995) and Blundell and Bond (1998) have suggested that a problem of weak instruments may be present,

---

11   All models are estimated using the Stata *xtabond2* module created by David Roodman (2003).

12   Windmeijer (2005), among others, considers restricting the number of instruments as an advantageous strategy.

especially if the time series of regressors are highly persistent. To correct for this, the authors propose an alternative method (system GMM) where additional moment conditions are introduced. For our data, we believe that weakness of instruments is not a major concern because the variables of interest, in particular the flow variables, are not very persistent over time. Moreover, according to simple correlation coefficients, twice-lagged level variables have better predictive power for first-differenced variables than twice-lagged differences have for level regressors. Therefore, we find the difference GMM method suitable, and preferable, for our aims. In particular, we use the two-step GMM estimator because, with an appropriate finite sample corrected variance estimate, it has been found to be superior to the one-step estimator; see Windmeijer (2005).

5.5   Results

The effects of hirings and job separations
The findings of the analyses are reported in Table 1, where we present three alternative specifications of the migration equation. Specification (a) includes the total hiring and total separation variables. In specifications (b) and (c), these variables are decomposed by labour market status of workers and by industry, respectively.

The impression from the results in column (a) is that there is something wrong with the model specification. None of the coefficients shown in the column is statistically significant. The probable explanation for this is found by looking at the results from specification (b), where worker flows are disaggregated by source (HU, $HOTH$) and destination (SU, $SOTH$) labour market status. We can see that the effects of flows differ markedly in size depending on the source and destination labour market status. Thus, in specification (a), incorrect restrictions on the parameters are imposed. We believe that this is the primary reason for the statistical insignificance of the coefficients in that specification. In (b), most coefficients are of the expected sign, and many of them are significantly different from zero.

Results from specification (b) show that hiring from unemployment (HU) has a sizeable and statistically significant negative effect on out-migration. In contrast, the coefficient of hires from other sources ($HOTH$) is statistically insignificant. Therefore, hiring hinders out-migration insofar as unemployed workers are hired. We do not find statistically significant effects on out-migration for worker separations (SU and $SOTH$). It thus seems that the transition of workers from employment to nonemployment does not increase local competition for jobs enough to induce out-migration.[13] Further, job changes (to local and nonlocal jobs) by the workers in local firms, captured by $SOTH$, are not associated with out-migration. We are able to identify the result commonly found in earlier literature, which is that more out-migration occurs when the rate of

---

13   The relatively generous unemployment benefits in the beginning of the unemployment spell may discourage search activity by newly unemployed workers (captured by SU).

Table 1. Determinants of interregional migration

	(a)		(b)		(c)	
Source province (out-migration)						
H	3.05	(3.63)	—		—	
HU	—		-31.24***	(10.51)	—	
$HOTH$	—		0.14	(1.71)	—	
H_{ind}	—		—		-1.27	(2.33)
H_{constr}	—		—		-1.74	(5.07)
H_{serv}	—		—		4.63*	(2.81)
S	3.50	(9.12)	—		—	
SU	—		-3.98	(4.99)	—	
$SOTH$	—		-0.19	(2.36)	—	
S_{ind}	—		—		-1.47	(3.01)
S_{constr}	—		—		5.97	(7.92)
S_{serv}	—		—		-0.26	(5.21)
$NETO$	-17.54	(14.45)	-6.24**	(3.14)	-1.20	(3.85)
U	14.18	(10.52)	11.38***	(2.66)	-0.84	(2.58)
$WAGE$	0.01	(0.16)	0.07	(0.06)	0.01	(0.04)
JOB_{pri}	0.08	(11.62)	8.62	(5.52)	0.61	(4.66)
JOB_{ind}	5.61	(5.85)	2.08	(3.92)	-0.28	(3.13)
JOB_{constr}	-10.51	(9.33)	6.26	(5.02)	4.13	(5.55)
Destination province (in-migration)						
H	2.37	(2.52)	—		—	
HU	—		17.76**	(8.6)	—	
$HOTH$	—		-0.48	(1.35)	—	
H_{ind}	—		—		1.76	(1.85)
H_{constr}	—		—		5.47	(4.22)
H_{serv}	—		—		4.03*	(2.09)
S	-3.11	(3.55)	—		—	
SU	—		-17.25***	(3.41)	—	
$SOTH$	—		-4.59**	(1.83)	—	
S_{ind}	—		—		-2.44	(2.29)
S_{constr}	—		—		-5.66	(6.32)
S_{serv}	—		—		-8.54**	(4.34)
$NETO$	-6.45	(4.36)	2.96	(2.8)	1.73	(2.95)
U	4.14	(4.21)	-2.80	(2.24)	-1.86	(1.82)
$WAGE$	0.08	(0.09)	0.07*	(0.04)	0.06**	(0.03)
JOB_{pri}	-2.38	(5.81)	-0.50	(2.68)	-5.09*	(2.61)
JOB_{ind}	-0.41	(5.11)	1.01	(2.8)	-4.71**	(2.37)
JOB_{constr}	6.60	(6.16)	6.34	(4.31)	1.47	(4.84)
DIS	-9.72	(27.9)	-8.19	(14.59)	-1.97	(9.4)
$lagdep$	-0.28	(0.24)	-0.01	(0.11)	-0.03	(0.09)
No. of instruments	68		96		124	
Hansen	0.637		0.853		0.031	
AR(2)	0.536		0.562		0.382	

Notes: Two-step GMM with Windmeijer correction. Dep. var.: log migration rate (per 1,000 persons) to the destination province. N = 2,736. Additional controls (see Section 3) and year dummies included. Robust standard errors in parentheses. * denotes significance at 10% level, ** at 5% level and *** at 1% level.

unemployment is higher. Net change in employment in other sectors ($NETO$) decreases out-migration.[14]

The results from specification (b) concerning in-migration underline different features of the labour market than the ones that are important for out-migration. Most importantly, the unemployment rate of the province does not enter significantly, even though the coefficient has the expected sign. Worker flows seem to matter more. Worker flow from unemployment to employment (HU) enters positively and significantly. Hires from other sources (employment and outside the labour force) ($HOTH$) do not seem to have a role in attracting migrants. This result, together with the result that job separations leading to other labour market states than unemployment does not induce out-migration, may indicate that long-distance residential changes associated with job changes are rare. In model (b), both flows from employment to unemployment (SU) and to other destinations ($SOTH$) have statistically significant negative effects on in-migration, although the effect of the flow to unemployment is distinctively larger. Further, higher wages are associated with higher in-migration. Measures of industrial structure and its dissimilarity between the source and destination region enter without statistical significance, although the coefficient is negative, as expected.[15]

The results from specification (c), where worker flows are disaggregated by industry, imply that sectoral boundaries do not generate the aforementioned heterogeneity in the effects of labour market flows. Our data contain an industry classification of seven industries.[16] For specification (c), we have aggregated these to form flow variables for three main industries: mining, manufacturing and energy (ind); construction ($constr$); and services ($serv$). This classification corresponds to the industry classification of our industry share variables (JOB_{ind} and JOB_{constr}). As in specification (a), most of the coefficients are insignificant. Also the diagnostics of the model are troublesome because the Hansen test of overidentifying restrictions rejects the hypothesis that the instruments are not correlated with the errors of the model. We believe that these problems are, as in specification (a), due to misspecification of the model and the resulting weak explanatory power of the regressors.

Thus far, our results highlight the role of a large pool of unemployed and limited labour market possibilities for them as factors that induce out-migration. In turn, people move to regions with abundance of labour market possibilities for unemployed

---

14   An analysis with industry-disaggregated $NETO$ reveals that the (mostly negative) net changes in agricultural employment are behind this result. Thus, decreases in agricultural employment have had a role in driving people away from some provinces.

15   The results from OLS and within estimations of model (b) can be found in the appendix. Many of our key results are robust to assuming exogeneity of regressors and even to excluding the fixed effects.

16   The industries are 1) mining, manufacturing, and energy, etc; 2) construction; 3) trade; 4) hotels and restaurants; 5) transportation, etc; 6) finance; and 7) real estate, business services, etc. Experimenting with models that separated between all seven industries or with alternatively classified variables did not result in different conclusions than those obtained with specification (c).

job seekers and low flow out of employment (especially into unemployment). It seems that interregional matching of unemployed workers and jobs is more important than interregional job changing.

The effects of structural change, reallocation, and net employment changes

We estimate three additional model specifications, in which we allow the net change in employment to have a differential effect from simultaneous hires and separations. By doing this, we assess the robustness of our earlier findings and test our second hypothesis, which is that changes in the structure of employment in regions are related to interregional migration. To measure structural change, simultaneous hires and separations are decomposed into three parts: employment shifts between sectors, employment shifts between firms, and simultaneous hires and separations within firms.

Although all job openings may be available for potential in-migrants, hires in excess of simultaneous separations may be more effective in attracting workers from other regions because they are associated with an increase in total employment. The results of specification (b) in Table 1 also hint that it can be the net change that matters most as we found that the coefficients of the destination province's flow from unemployment to employment and flow from employment to unemployment are almost equal. This would mean that simultaneous hires and separations may leave in-migration unaffected. To assess the validity of this idea, we include separate measures of net changes in employment and simultaneous hires and separations simultaneously in our estimations.

In the job and worker flow literature, simultaneous hires and separations are routinely divided into two components: excess job reallocation and excess worker turnover; see Burgess *et al.* (2000) and Davis *et al.* (1996). Excess job reallocation can be further decomposed into two components: reallocation across industries and reallocation within industries; see Davis *et al.* (1996), Davis and Haltiwanger (1992) and Dunne *et al.* (1989). In the calculation of our measures of net employment changes and simultaneous hires and separations, as well as the measures of structural change, we follow the conventions of the earlier literature.

First, we define the total excess reallocation (ER). ER is the number of hires with simultaneous job separations, and it equals the smaller of hiring (H) and separation (S).[17] In other words, ER is the total number of separation-hire pairs that are not needed to attain the net employment change at the level of the (local) labour market as a whole. Notice that every hire and job separation is captured by either ER or the net change in employment (NET) because $H + S = |NET| + 2 \times ER$.

The next step is to decompose excess reallocation (ER) into three parts: employment shifts between industries (ERB), employment shifts between establishments (within the industries) (ERW), and establishment-level excess worker turnover ($CHUR$).

---

17   To simplify the discussion at this point, we talk about absolute numbers of hires, separations etc. instead of scaled flows. In all analyses, the flow variables are divided by population aged 15–74 years, as mentioned earlier. See the appendix for the formulas used to compute the variables.

Industry-level data on job flows allow us to use the procedure introduced by Dunne *et al.* (1989) to calculate *ERB* and *ERW*. Finally, it is straightforward to calculate the establishment-level excess turnover, because $ER = ERB + ERW + CHUR$.[18,19]

To provide some further intuition behind the variables *ERB*, *ERW*, and *CHUR*, notice that every hire counted in the variable *ER* has a counterpart, that is, a separation of a worker. This separation can occur in the same establishment, in another establishment within the same industry, or in another industry. It feels natural to hypothesize that there are differences between these three situations. For example, if contemporaneous separation only occurs in some other industry, the skills of the newly separated worker may not be suitable for the new job opening. In this case, a job seeker from another region may be a good candidate for the job. In contrast, a situation with a contemporaneous separation in the same industry may involve less demand for the skills of other regions' candidates. The third category of excess reallocation is likely to be different from the other two, but it is not perfectly clear what kind of situations it may characterize. Simultaneous hiring and separation in the same establishment cannot involve hiring and separation of the same worker. Rather, these situations may reflect voluntary quits, lay-offs, retirements, or other such occasions and their replacements. The question of how, if in any way, this affects migration, is left as an empirical question to be answered by our results.

In the first two specifications of Table 2, we do not make a distinction between the three types of excess reallocation to be able to make a simple distinction between net employment change and simultaneous hiring and separations. In the first specification (d), our labour market variables are thus *NET* (net employment change) and *ER* (total excess reallocation). In the next specification (e), we calculate the net employment change and the excess reallocation variables separately for the flows with unemployment as the source/destination (*NETU*, *ERU*) and the flows with other labour market states as the source/destination (*NETOTH*, *EROTH*). Specification (f) serves as our most direct test for the effects of change in employment structure. We make a distinction between excess job reallocation and churning, as well as between job reallocation across industries and within industries. Thus, our explanatory variables in (f) include *ERB*, *ERW*, and *CHUR*.

Model (d) suffers from similar problems as model (a) earlier. We do not find any statistically significant results concerning our variables of interest. It can be noted,

18   The links between the measures in the earlier literature and our variables are straightforward: The sum of *ERB* and *ERW* multiplied by 2 is the excess job reallocation. Similarly, the two excess turnover measures in Dunne *et al.* (1989) are *ERB* and *ERW* multiplied by 2. *CHUR* is the excess worker turnover (or churning) divided by two. For more discussion and interpretations of these measures, see Davis *et al.* (1996) and Dunne *et al.* (1989).

19   Compared with net changes in employment, job reallocation and churning constitute a considerable share of the labour market flows in our data. Similar observation is made in studies documenting labour market flows; e.g., Davis *et al.* (1996), Davis *et al.* (2006), Burgess *et al.* (2000).

Jani-Petri Laamanen

however, that the coefficients of net employment change variables (NET) have the expected signs for both the source and the destination province. Model (e), where net employment change and excess reallocation are included separately for flows concerning unemployment ($NETU$, ERU) and other labour market states ($NETOTH$, $EROTH$), works better. For out-migration, none of the labour market flow variables appears

Table 2. Determinants of interregional migration

	(d)		(e)		(f)	
Source province (out-migration)						
NET	-2.47	(6.9)	—		-0.64	(2.49)
$NETU$	—		-1.40	(4.97)	—	
$NETOTH$	—		-1.03	(2.46)	—	
ER	9.30	(9.49)	—		—	
ERU	—		3.72	(10.87)	—	
$EROTH$	—		2.74	(2.68)	—	
ERB	—		—		6.74	(5.06)
ERW	—		—		-2.58	(3.86)
$CHUR$	—		—		4.18	(5.11)
$NETO$	-19.76	(13.6)	-7.03**	(3.35)	-9.57**	(4.61)
U	17.10	(11.5)	8.66***	(2.49)	5.19*	(2.77)
$WAGE$	-0.01	(0.15)	0.00	(0.06)	-0.19**	(0.08)
JOB_{pri}	0.68	(9.27)	9.91*	(5.61)	6.22	(5.54)
JOB_{ind}	4.04	(5.68)	7.12**	(3.25)	3.80	(3.67)
JOB_{constr}	-4.41	(11.31)	8.31	(5.64)	6.76	(6.17)
Destination province (in-migration)						
NET	2.82	(2.96)	—	(3.1)	2.91*	(1.50)
$NETU$	—	(3.33)	14.10***	(1.54)	—	(4.72)
$NETOTH$	—	(4.35)	3.29**	(8.1)	—	(2.35)
ER	-1.57	(3.96)	—	(2.00)	—	(3.45)
ERU	—	(0.09)	-15.00*	(2.72)	—	(2.85)
$EROTH$	—	(5.62)	-4.99**	(1.92)	—	(2.38)
ERB	—	(5.21)	—	(0.04)	-3.47	(0.05)
ERW	—	(6.35)	—	(2.67)	-2.80	(3.02)
$CHUR$	—	(29.95)	—	(2.82)	-0.70	(3.23)
$NETO$	-6.17	(0.27)	-1.65	(4.22)	-5.79**	(4.71)
U	4.18		-1.03	(15.36)	3.68	(13.57)
$WAGE$	0.09		0.09**	(0.13)	0.05	(0.11)
JOB_{pri}	-1.68		-0.81		1.60	
JOB_{ind}	0.02		0.12		2.02	
JOB_{constr}	4.91		5.70		5.74	
DIS	-13.58		-25.24		-1.48	
$lagdep$	-0.25		-0.02		-0.09	
Number of instruments	68		96		96	
Hansen	0.732		0.619		0.022	
AR(2)	0.480		0.813		0.271	

Notes: Two-step GMM with Windmeijer correction. Dep. var.: log migration rate (per 1,000 persons) to the destination province. N = 2,736. Additional controls (see Section 3) and year dummies included. Robust standard errors in parentheses. * denotes significance at 10% level, ** at 5% level and *** at 1% level.

significant. However, net employment change variables (*NETU* and *NETOTH*) have negative coefficients (as expected), whereas excess reallocation variables (*ERU*, *EROTH*) have positive coefficients. As in model (b), net employment change in other sectors (*NETO*) has a negative effect, and unemployment rate has a positive effect on out-migration. The measures of net employment changes and excess reallocation do not separate between hires and separations. Therefore, the earlier result from model (b), that hires from but not separations to unemployment are important for out-migration, is not captured by model (e).

With regard to in-migration, model (e) roughly reproduces the results of model (b). Net change in employment caused by hires from and separations into unemployment (*NETU*) enters positively and with a strong significance. However, contrary to the suggestion of model (b), excess reallocation (*ERU*) also matters. This variable has a negative coefficient (with a significance level of 10%), which means that, ceteris paribus (e.g., with equal net change in employment), a simultaneous increase in both hiring from unemployment and separations into unemployment discourages in-migration. One interpretation for this is that one newly unemployed worker has a greater negative effect through increased competition for jobs than the positive effect that one hire of an unemployed worker has through increased job opportunities. For worker flows between employment and labour market states other than unemployment, we find similar results, but the coefficients are smaller in size. Net employment change due to hires from and separations into employment and outside the labour force (*NETOTH*) has a positive effect on in-migration. However, simultaneous hires and separations (*EROTH*) has a hindering effect. A large share of this reallocation is likely to be due to employed workers changing jobs. A possible interpretation would be that active on-the-job search and resulting job switching increases the competition for jobs, keeping the workers of other regions out of the local market.[20]

Model (f), to some extent, supports our notion related to model (c) that structural change is not an important factor in determining interregional migration. Although some of the estimated coefficients are now significantly different from zero, the Hansen test indicates that the instruments are not valid. We, again, interpret this as a sign of insufficient explanatory power of the variables included. This may be, in part, because boundaries between different industries or between establishments do not significantly hinder recruitment. Indeed, Bjelland *et al.* (2008) have documented that a very large share of workers switching jobs is also changing industry. Our finding is also in line with the result of Robson (2009), which shows that although structural change does somewhat affect regional labour market performance, these effects are small. It is possible that the structural change causing problems in the matching of skills and jobs in the labour market is likely to be more about boundaries between different

---

20   The results from OLS and within estimations of model (e) can be found in the appendix. Many of our key results are robust to assuming exogeneity of regressors and even to excluding the fixed effects.

Jani-Petri Laamanen

occupations. However, our data do not allow occupational disaggregation, so that our way of measuring structural change may be imperfect.

In none of the models (a)–(e) do we find a statistically significant effect for the lagged migration variable. The coefficient may be biased because, even if not true for the key variables of interest, lagged levels are poor predictors for differences of this variable. The explanatory power of our instruments is not markedly increased by including more lags or by explaining levels of the migration variable by lagged differences (i.e., by using the system GMM rather than difference GMM estimator). Increasing moment conditions makes the coefficient of past migration positive but not significant, and the conclusions regarding the labour market variables remain qualitatively similar. Since we are not mainly interested in the effect of past migration, we have conducted these additional analyses predominantly to assure ourselves of the robustness of the other results.

To assess how important different labour market characteristics have been in determining the migratory flows in our data, we can compare the coefficient estimates with the actual variation in the corresponding variables. We do this for all coefficients that appeared statistically significant in specification (b) or in specification (e). For out-migration, a change of one standard deviation in a province's unemployment rate has a larger effect than a standard deviation change in any other variable in both specifications. In specification (b), this effect is 4.6 times as large as the effect of one standard deviation change in the hires-from-unemployment variable (HU). The effect of the variation in net change in employment in other sectors ($NETO$) is considerably smaller. We can therefore say that variation in the unemployment rate is the most important labour market determinant of out-migration. The rationale for this is easily given because unemployed individuals have a relatively high propensity to move, as discussed earlier. It is also conceivable that the unemployment rate, to some extent, captures a province's labour market possibilities. Related to the results from specification (e), it should also be noted that the effect of variation in industry shares (JOB_k) appears large.

With regard to in-migration, labour market flows have relatively large significance as determinants. Even though the coefficients of the flows with unemployment as the source or destination (HU and SU) are markedly larger than the coefficients of the other flows ($HOTH$ and $SOTH$), their economic significance is reduced by their relatively small variation. The effect of a one standard deviation change in the employment-to-unemployment flow or the unemployment-to-employment flow roughly corresponds with the effect of a one standard deviation change in the flow from employment to employment and outside the labour force. Therefore, each unemployment-to-employment and employment-to-unemployment transition has a larger effect on in-migration than any other transition, but due to the relatively small variation in the former flows, their economic significance remains limited. An interesting observation related to specification (e) is that excess reallocation of workers through employment

and outside the labour force (*EROTH*) is the most important flow measure, while net employment change from unemployment (*NETU*) is almost as important. Excess reallocation through unemployment (*ERU*) and net change in employment from other labour market states (*NETOTH*) contribute much less to in-migration. It should also be noted that variation in wage contributes strongly to the variation in in-migration.

In addition to our results concerning the labour market variables of interest, we find some statistically significant results for the other control variables not included in the results tables. In reporting these results, we focus on those coefficients that appear significant in both models (b) and (e). Firstly, a higher share of owner-occupied housing in a province is associated with lower out-migration. This result is intuitive and in line with micro-level evidence of lower geographical mobility of homeowners; for Finnish results, see Nivalainen (2004) and for a survey, see Dietz and Haurin (2003). Secondly, the share of early retired inhabitants has a positive effect on out-migration, probably because retired individuals lack ties to the local labour market and are therefore freer to move to other locations. Again, in accordance with micro-evidence (Nivalainen, 2004), we find that the more highly educated the population in a province is, the higher is out-migration. An exception to this is that a higher share of the population in the highest educational category (upper university degree or more) is associated with lower out-migration. Housing prices in the source province enter models (b) and (e) negatively and statistically significantly (at the 5% and 10% level, respectively). This result is in contrast with our expectation, but a rationale for it can be easily given. Because housing prices concern the actual transactions of houses and apartments, they are likely to be responsive to changes in housing demand. Thus, out-migration for other than housing market reasons may depress local housing prices. It is conceivable that we identify this reverse link because we may not have a proper instrument for housing prices in our set of instruments. It should be noted here that in specification (b), we find a negative and significant (at the 10% level) effect on in-migration for the rental price variable, which is likely to better capture the variation in housing costs between provinces. The fact that our findings related to the control variables are in line with theoretical expectations and micro studies increases our confidence in the results.

Discussion

Studies using labour market flow variables to explain interregional migration flows are rare, but some comparisons to earlier research can still be made. Most importantly, our result that hires of unemployed workers decreases out-migration and increases in-migration is in line with the results of Carlsen *et al*. (2006). In that study, the authors find that the probability of a region's unemployed leaving unemployment is positively associated with in-migration. However, other labour market flows are not included in their model, so our analysis is more comprehensive in this respect. We found that variables measuring job-to-job transitions of workers are not associated with

interregional migration flows. This result is in line with the finding that residential changes are not linked to job changes by van Ommeren *et al*. (1999).

Interestingly, our results concerning the effects of excess reallocation deviates from the results obtained by Hämäläinen and Böckerman (2004), who use very similar data from almost the same time period. They find that excess job reallocation and churning have positive effects on net migration (mainly through reductions in out-migration), whereas we find that excess reallocation decreases in-migration. On the other hand, our model (b) suggests that hires from unemployment reduce out-migration while the impact of other hires, as well as separations, do not differ from zero. This indicates that the effect of any excess turnover is negative, a result in line with Hämäläinen and Böckerman (2004). There are several differences between their study and ours, so that the results are not directly comparable, and there are many possible interpretations of the discrepancy. Hämäläinen and Böckerman (2004) study migration between smaller regional units, they are unable to simultaneously control the characteristics of source and destination regions, their labour market flow variables are calculated as rates (absolute flows divided by employment), and they do not separate between flows from and into unemployment and flows from and to other sources/destinations.

Our most important result is that worker flows from and into unemployment have significant effects which are differential from the effects of other labour market flows. This is in line with the observation of Davis *et al*. (2006) that flows from and into unemployment are different from other flows and with the observation of Anderson and Burgess (2000) that the source of a new hire is of importance. Differential effects may be due to differences in search activity and migration propensities of workers, or due to employers' preference for employed workers. One way to interpret our result is related to social networks and information. Potential in-migrants may lack information on the labour market of the potential destination region. If unemployed workers are weakly attached to the surrounding labour market as well, a large number of hires of unemployed workers in a region may indicate that jobs are available for in-migrants with weaker networks and less information as well. These ideas are in line with such theories of interregional job search that stress the role of geographical distance in interregional information flows.

A further point should be made about the functional form of our migration equation. Semilog function was chosen because it produced the most credible results and, more importantly, because the specification tests rejected both correlation between errors and instruments and second-order autocorrelation in residuals. The interpretation of the semilog models' coefficients is highly intuitive and in line with theoretical ideas. A coefficient of a semilog specification should be interpreted as the marginal effect of one-unit change in the explanatory variable on relative change in the dependent variable. Therefore, for instance, a one-unit increase in per capita hires from unemployment increases per capita in-migration by the same percentage for all

provinces. This means that the effects of the explanatory variables on migration are stronger when the migration is, in the baseline, high between the two provinces. Theoretically, this may be due to a short geographical distance between the two provinces, a strong past migratory link between the provinces, or other factors, such as cultural similarity between the provinces. Factors of this kind are likely to facilitate information exchange between the two provinces and strengthen the role of the other province as the potential destination of moving. Thereby, changes in the labour market (and other) characteristics of the other region are more relevant for potential movers.

5.6   Concluding Remarks

To thoroughly explore the relationship between interregional migration and the state and dynamics of regional labour markets, we used data on bilateral migration flows and disaggregated labour market flows from Finnish provinces from 1988 to 1996. Analysing the data with the dynamic panel GMM method leads to clear conclusions on the labour market reasons for out-migration and in-migration. The general conclusion is that different types of labour market changes have differing effects on migration, so that the use of simple measures of labour market conditions as explanatory variables has the potential to produce misleading results. According to our results, the ability of a region to offer labour market possibilities to unemployed workers is an important factor in holding back out-migration and attracting in-migrants. Our results indicate that the reasons for out-migration are to some extent different from the reasons for in-migration. Hiring from the pool of employed workers and those outside the labour force also attracts in-migrants but to a much lesser extent than hires from unemployment do. However, when a region simultaneously experiences separations of workers from their jobs, this more than offsets the positive in-migration effect of hires. Therefore, simultaneous hires and separations hinder in-migration, whereas net increases in employment encourage it. We found an effect for the unemployment rate, a variable often used to explain migratory flows, but only on out-migration. Our results show that a high local unemployment rate leads to increased out-migration. We discuss that this effect may be due to the higher propensity of the unemployed to migrate, a result previously found in many micro-level studies.

We also tested the hypothesis that the extent of structural change in local labour markets affects interregional migration. However, we found no robust evidence on the effects of interindustry or intraindustry interfirm shifts in employment. Simultaneous hires and separations within firms did not gain statistical significance in our estimations. These findings are in line with a result from studies of job switching: workers who change jobs often cross industry boundaries.

Appendix

Table A1. Variable definitions.

dep. var.	log migration rate (per 1,000 persons; 15 to 74 year olds) to a destination province
H	hires*
HU	hires from unemployment*
$HOTH$	hires from employment and outside the labour force*
H_k	hires in industry k*
S	separations*
SU	separations into unemployment*
$SOTH$	separations into employment and outside the labour force*
S_k	separations in industry k*
NET	H - S
$NETU$	HU - SU
$NETOTH$	$HOTH$ - $SOTH$
ER	$\min(H, S) = 0.5 \times (H + S - |NET|)$
ERU	$\min(HU, SU) = 0.5 \times (HU + SU - |NETU|)$
$EROTH$	$\min(HOTH, SOTH) = 0.5 \times (HOTH + SOTH - |NETOTH|)$
ERB	$0.5 \times [\sum_{k \in K^+} NET_k - \sum_{k \in K^-} NET_k - |NET|]$
ERW	$0.5 \times \sum_k [\sum_{f \in F_k^+} NET_f - \sum_{f \in F_k^-} NET_f - |NET_k|]$
$CHUR$	$ER - ERB - ERW$
$NETO$	net employment change in sectors excluded from job/worker flow data*
U	unemployment rate
$WAGE$	total wage income divided by the number of wage earners
JOB_k	share of jobs in industry k
DIS	$\sum_k (JOB_{ik} - JOB_{jk})^2$ (i = source province, j = destination province)

Notes: K^+ and K^- are the industries where employment is growing and falling, respectively. F^+ and F^- are the establishments where employment is growing and falling, respectively. * = divided by population aged 15–74 years in the province.

Table A2. Descriptive statistics, 19 provinces, 1988–96.

Variable	Mean	Std. Dev.	Min	Max
dependent variable	1.787	0.514	0.144	3.098
H	0.080	0.032	0.033	0.204
HU	0.011	0.006	0.003	0.030
$HOTH$	0.069	0.033	0.019	0.198
H_{ind}	0.027	0.013	0.007	0.072
H_{constr}	0.012	0.005	0.005	0.031
H_{serv}	0.044	0.018	0.018	0.134
S	0.087	0.029	0.039	0.195
SU	0.016	0.007	0.005	0.036
$SOTH$	0.071	0.031	0.023	0.190
S_{ind}	0.029	0.013	0.008	0.066
S_{constr}	0.014	0.005	0.006	0.026
S_{serv}	0.046	0.016	0.026	0.127
NET	-0.007	0.021	-0.056	0.028
$NETU$	-0.005	0.009	-0.027	0.017
$NETOTH$	-0.002	0.015	-0.042	0.027
ER	0.075	0.030	0.033	0.195
ERU	0.010	0.004	0.003	0.019
$EROTH$	0.064	0.031	0.019	0.190
ERB	0.002	0.003	0.000	0.016
ERW	0.039	0.015	0.017	0.093
$CHUR$	0.036	0.014	0.015	0.099
$NETO$	-0.003	0.007	-0.025	0.011
U	0.138	0.076	0.013	0.285
$WAGE$	19.514	2.286	13.120	24.877
JOB_{pri}	0.118	0.047	0.010	0.228
JOB_{ind}	0.229	0.053	0.150	0.346
JOB_{constr}	0.060	0.011	0.039	0.086
DIS	0.012	0.010	0.001	0.053

Jani-Petri Laamanen

Table A3. Determinants of interregional migration

	OLS		Within	
Source province (out-migration)				
HU	-8.44*	(4.96)	-10.19**	(4.37)
HOTH	-0.85	(1.05)	-0.64	(0.86)
SU	-0.22	(3.23)	0.51	(2.67)
SOTH	1.44	(1.52)	-0.15	(1.22)
NETO	-4.03*	(2.13)	-4.74**	(2.07)
U	1.84*	(0.94)	3.60***	(1.14)
WAGE	0.00	(0.01)	-0.01	(0.01)
JOB_{pri}	1.72	(1.05)	0.41	(1.74)
JOB_{ind}	1.36**	(0.63)	0.78	(1.35)
JOB_{constr}	2.88	(2.19)	0.08	(2.62)
Destination province (in-migration)				
HU	11.02***	(3.58)	8.97**	(3.62)
HOTH	0.41	(1.13)	-1.14	(1.01)
SU	-0.60	(2.17)	-6.69***	(2.29)
SOTH	-0.32	(1.40)	-1.48	(1.14)
NETO	5.23**	(2.01)	2.97	(1.92)
U	-0.19	(0.39)	-1.74**	(0.72)
WAGE	-0.01	(0.01)	0.01	(0.01)
JOB_{pri}	-0.41	(0.41)	-0.73	(1.56)
JOB_{ind}	-0.54*	(0.29)	0.30	(1.31)
JOB_{constr}	-0.76	(1.42)	-1.39	(2.10)
DIS	-1.89***	(0.60)	7.65***	(2.70)
lagdep	0.94***	(0.01)	0.03	(0.04)
R^2	0.941		0.970	

Notes: Specification (b) (see Table 1.). Dep. var.: log migration rate (per 1,000 persons) to the destination province. N = 2,736. Additional controls (see Section 3) and year dummies included. Robust standard errors in parentheses. * denotes significance at 10% level, ** at 5% level and *** at 1% level.

Table A4. Determinants of interregional migration

	OLS		Within	
Source province (out-migration)				
NETU	-1.89	(2.41)	-3.40*	(2.01)
NETOTH	-1.39	(1.05)	-0.49	(0.86)
ERU	-5.41	(5.73)	-8.66	(5.68)
EROTH	1.02	(1.25)	0.34	(1.14)
NETO	-3.76*	(2.12)	-4.42**	(2.04)
U	1.51*	(0.88)	3.49***	(1.06)
WAGE	-0.01	(0.01)	-0.01	(0.01)
JOB_{pri}	1.64	(1.02)	-0.20	(1.76)
JOB_{ind}	1.34**	(0.63)	0.84	(1.33)
JOB_{constr}	2.37	(2.08)	-0.43	(2.61)
Destination province (in-migration)				
NETU	2.78	(1.88)	6.54***	(1.70)
NETOTH	0.63	(1.14)	0.33	(0.91)
ERU	7.93*	(4.28)	1.16	(4.31)
EROTH	-0.14	(0.96)	-2.83***	(0.99)
NETO	5.46***	(2.02)	2.63	(1.90)
U	-0.20	(0.38)	-1.54**	(0.71)
WAGE	-0.01	(0.01)	0.01	(0.01)
JOB_{pri}	-0.54	(0.39)	-0.43	(1.55)
JOB_{ind}	-0.55*	(0.29)	0.27	(1.26)
JOB_{constr}	-0.58	(1.42)	-1.43	(2.09)
DIS	-1.77***	(0.59)	7.28***	(2.65)
lagdep	0.94***	(0.01)	0.03	(0.04)
R^2	0.941		0.970	

Notes: Specification (e) (see Table 2.). Dep. var.: log migration rate (per 1,000 persons) to the destination province. N = 2,736. Additional controls (see Section 3) and year dummies included. Robust standard errors in parentheses. * denotes significance at 10% level, ** at 5% level and *** at 1% level.

Jani-Petri Laamanen

References

Anderson, P. M. and Burgess, S. M. (2000) Empirical matching functions: estimation and interpretation using state-level data, *Review of Economics and Statistics*, **82**, 93–102.

Antolin, P. and Bover, O. (1997) Regional migration in Spain: the effect of personal characteristics and of unemployment, wage and house price differentials using pooled cross-sections, *Oxford Bulletin of Economics and Statistics*, **59**, 215–235.

Arellano, M. and Bond, S. (1991) Some tests of specifications for panel data: Monte Carlo evidence and an application to employment equations, *Review of Economic Studies*, **58**, 227–297.

Arellano, M. and Bover, O. (1995) Another look at the instrumental-variable estimation of error-component models, *Journal of Econometrics*, **68**, 29–52.

Bjelland, M., Fallick, B. C., Haltiwanger, J. C. and Mcentarfer, E. (2008) Employer-to-employer flows in the United States: estimates using linked employer—employee data. NBER Working Papers 13867. Cambridge, MA: National Bureau of Economic Research.

Blundell, R. and Bond, S. (1998) Initial conditions and moment restrictions in dynamic panel data models, *Journal of Econometrics*, **87**, 115–143.

Burgess, S., Lane, J. and Stevens, D. (2000) Job flows, worker flows, and churning, *Journal of Labor Economics*, **18**, 473–502.

Böckerman, P. and Maliranta, M. (2001) Regional disparities in gross job and worker flows in Finland, *Finnish Economic Papers*, **14**, 84–103.

Böheim, R. and Taylor, M. P. (2002) Tied down or room to move? Investigating relationships between housing tenure, employment status and residential mobility in Britain, *Scottish Journal of Political Economy*, **49**, 369–392.

Carlsen, F., Johansen, K. and Røed K. (2006) Wage formation, regional migration and local labour market tightness, *Oxford Bulletin of Economics and Statistics*, **68**, 423–444.

Davis, S. J., Faberman, R. J. and Haltiwanger, J. (2006) The flow approach to labor markets: new data sources and micro-macro links. NBER Working Papers 12167. Cambridge, MA: National Bureau of Economic Research.

Davis, S. J. and Haltiwanger, J. (1992) Gross job creation, gross job destruction, and employment reallocation, *The Quarterly Journal of Economics*, 107, **3**, 819–863.

Davis, S. J., Haltiwanger, J. and Schuh, S. (1996) *Job Creation and Destruction*, MIT Press, Cambridge, MA.

Dietz, R. D. and Haurin, D. R. (2003) The social and private micro-level consequences of homeownership, *Journal of Urban Economics*, **54**, 401–450.

Dunne, T., Roberts, M. J. and Samuelson, L. (1989) Plant turnover and gross employment flows in the U.S. manufacturing sector, *Journal of Labor Economics*, 7, 48–71.

Etzo, I. (2011). The determinants of the recent interregional migration flows in Italy: a panel data analysis, *Journal of Regional Science*, 51(5), 948–966.

Fields, G. S. (1976) Labor force migration, unemployment and job turnover, *The Review of Economics and Statistics*, **58**, 407–415.

Fields, G. S. (1979) Place-to-place Migration. Some new evidence, *The Review of Economics and Statistics*, **61**, 21–32.

Furceri, D. (2006) Does labour respond to cyclical fluctuations? The case of Italy, *Applied Economics Letters*, **13**, 135–139.

Harris, J. R. and Todaro, M. P. (1970) Migration, unemployment and development: a two-sector analysis, *American Economic Review*, **60**, 126–142.

Hatton, T. and Tani, M. (2005) International migration and inter-regional mobility in the UK: 1982–2000, *The Economic Journal*, **155**, 342–358.

Hämäläinen, K. and Böckerman, P. (2004) Regional labor market dynamics, housing and migration, *Journal of Regional Science*, **44**, 543–568.

Jackman, R. and Savouri, S. (1992) Regional migration in Britain: an analysis of gross flows using NHS central register data, *The Economic Journal*, **102**, 1433–1450.

McCormick, B. and Wahba, J. (2005) Why do the young and educated in LDCs concentrate in large cities? Evidence from migration data, *Economica*, **72**, 39–67.

Molho, I. (2001) Spatial search, migration and regional unemployment, *Economica*, **68**, 269–283.

Nivalainen, S. (2004) Determinants of family migration: short moves vs. long moves, *Journal of Population Economics*, **17**, 157–175.

Parikh, A. and van Leuvensteijn, M. (2003) Interregional labour mobility, inequality and wage convergence, *Applied Economics*, **35**, 931–941.

Petrongolo, B. and Pissarides, C. A. (2001) Looking into the black box: a survey of the matching function, *Journal of Economic Literature*, **39**, 390–431.

Pissarides, C. A. and Wadsworth, J. (1989) Unemployment and the inter-regional mobility of labour, *The Economic Journal*, **99**, 739–755.

Robson, M. (2009) Structural Change, specialization and regional labour market performance: evidence for the UK, *Applied Economics*, **41**, 275–293.

Roodman, D. (2003) XTABOND2: STATA module to estimate dynamic panel data models. Center for Global Development, Washington.

Sjaastad, L. A. (1962) The costs and returns of human migration, *Journal of Political Economy*, **70**, 80–93.

Van Ommeren, J., Rietveld, P. and Nijkamp, P. (1999) Job moving, residential moving, and commuting: a search perspective, *Journal of Urban Economics*, **46**, 230–253.

Windmeijer, F. (2005) A finite sample correction for the variance of linear efficient two-step GMM estimators, *Journal of Econometrics*, **126**, 25–51.

Yashiv, E. (2008) U.S. Labor market dynamics revisited, *Scandinavian Journal of Economics*, **109**, 779–806.